



**NetApp™**  
Go further, faster

NETAPP UNIVERSITY

## VMware on NetApp Solutions

Student Guide: Revision 1.0  
Date: April 2008  
SALES-ED-ILT-VMNSOLLAB

This page is intentionally left blank.

## ATTENTION

The information contained in this guide is intended for training use only. This guide contains information and activities that, while beneficial for the purposes of training in a closed, non-production environment, can result in downtime or other severe consequences and therefore are not intended as a reference guide. This guide is not a technical reference and should not, under any circumstances, be used in production environments. To obtain reference materials, please refer to the NetApp product documentation located at [www.now.com](http://www.now.com) for product information.

## COPYRIGHT

Copyright © 2008 NetApp. All rights reserved. Printed in the U.S.A. Specifications subject to change without notice.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products or materials described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product or materials does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

## RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## TRADEMARK INFORMATION

NetApp, the NetApp logo, and Go further, faster, FAServer, NearStore, NetCache, WAFL, DataFabric, FilerView, SecureShare, SnapManager, SnapMirror, SnapRestore, SnapVault, Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and SpinStor are registered trademarks of Network Appliance, Inc. in the United States and other countries. Network Appliance, Data ONTAP, ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, gFiler, MultiStore, SecureAdmin, Smart SAN, SnapCache, SnapDrive, SnapMover, Snapshot, vFiler, Web Filer, SpinAV, SpinManager, SpinMirror, and SpinShot are trademarks of NetApp, Inc. in the United States and other countries.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp is a licensee of the CompactFlash and CF Logo trademarks.



## TABLE OF CONTENTS

<b>WELCOME</b>	<b>1</b>
<b>MODULE 1: THE NETAPP VALUE PROPOSITION</b>	<b>1-1</b>
<b>MODULE 2: VMWARE CONNECTIVITY TO NETAPP STORAGE</b>	<b>2-1</b>
<b>MODULE 3: PROOF OF CONCEPT</b>	<b>3-1</b>
<b>MODULE 4: BACKUP, AND RECOVERY</b>	<b>4-1</b>
<b>MODULE 5: VMDK ALIGNMENT</b>	<b>5-1</b>
<b>MODULE 6: REPLICATION AND DISASTER RECOVERY</b>	<b>6-1</b>
<b>MODULE 7: VMWARE ON NETAPP DATA DEDUPLICATION</b>	<b>7-1</b>
<b>MODULE 8: CLONING ON VMWARE</b>	<b>8-1</b>
<b>MODULE 9: GROWING DATASTORES AND VIRTUAL DISKS</b>	<b>9-1</b>
<b>MODULE 10: SIZING STORAGE FOR VMWARE ENVIRONMENTS</b>	<b>10-1</b>
<b>MODULE 11: PERFORMANCE MONITORING</b>	<b>10-1</b>

## **Table of Contents**

### **Tab**

#### **Labs:**

1. Connecting Datastores and Cloning VMs
2. Virtual Machine Backups Using NetApp Snapshot
3. Virtual Machine Recovery from NetApp Snapshot Copies
4. File Level Recovery from NetApp Snapshot Copies
5. Aligning Virtual Disk File Systems
6. Virtual Machine Replication and Disaster Recovery Using SnapMirror
7. FAS Deduplication of NFS and VMware
8. Thin Provisioning & FAS Deduplication of VMFS Datastores
9. Provisioning Datastores & VMs with FlexClone
10. Growing Datastores




Go further, faster™

# VMware on NetApp Solutions

Revision 1.0



NetApp Confidential – Limited Use




## Logistics

- Introductions
- Schedule (start time, breaks, lunch, close)
- Food and drinks
- Restrooms

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 2

## Logistics




 **Safety**

- Alarm signal
- Evacuation route
- Assembly area

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 3

## Safety



## Course Objectives

- Upon completion of the course you will be able to:
  - Architect VMware on NetApp storage solutions
  - Articulate the key differentiators that NetApp products offer to customers with VMware server environments
  - Do VMware on NetApp presentations in front of customers

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 4


## Course Objectives

Upon completion of this program you will be able to:

Architect VMware on NetApp storage solutions

Articulate the key differentiators that NetApp products offer to customers with VMware server environments

Do VMware on NetApp presentations in front of customers



## Course Agenda

- Day 1
  - Morning
    - Module 1—NetApp Value Proposition
    - Module 2—Connectivity
      - Lab 1—Connecting Datastores and Cloning VMs
    - Module 3—POC
  - Afternoon
    - Module 4—Backup and Recovery
      - Lab 2—Virtual Machine Backups Using NetApp Snapshot
      - Lab 3—Virtual Machine Recovery from NetApp Snapshot Copies
      - Lab 4—File Level Recovery from NetApp Snapshot Copies
    - Module 5—Alignment
      - Lab 5—Aligning Virtual Disk File Systems
- Day 2
  - Morning
    - Module 6—Replication
      - Lab 6—Virtual Machine Replication and Disaster Recovery Using SnapMirror
    - Module 7—Deduplication
      - Lab 7—FAS Deduplication of NFS and VMware
      - Lab 8—Thin Provisioning & FAS Deduplication of VMFS Datastores
    - Module 8—Cloning
      - Lab 9—Provisioning Datastores & VMs with FlexClone
  - Afternoon
    - Module 9—Growing
      - Lab 10—Growing Datastores
    - Module 10—Sizing
    - Module 11—Performance

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 5

## Course Agenda






Module 1  
**The NetApp  
Value  
Proposition**



NetApp Confidential - Limited Use

Welcome to VMware on NetApp Solutions, a comprehensive survey of NetApp technologies that position NetApp as the storage leader in the exploding virtualization segment of enterprise computing.




## Learning Objectives

- High-level NetApp value propositions
- Mid-level NetApp value propositions

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 7

## Learning Objectives

This module describes the NetApp value proposition in the VMware virtualization market. When you have completed this module you will be able to explain confidently and succinctly to either a high-level customer audience or a mid-level customer audience the core business reasons why NetApp offers the best virtualized storage solution for their business.



### Why Server Virtualization Matters


- Gartner states that within the next 36 months 90% of all open systems will be virtualized
- NetApp provides virtualized storage solutions for open systems
- Our objective is to gain share in the server virtualization market by displacing incumbent storage vendors

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Why Server Virtualization Matters

Industry analysts, such as Gartner, are guiding toward 90 percent of open systems being virtualized within the next three years. NetApp has by far the best storage in this space and has an opportunity to come out gaining market share as customers implement a virtual infrastructure. It is NetApp customers who are most likely to go down this virtualization path first. That is why it is very important that we keep up with them.

NetApp's objective in the server virtualization market is to gain share by displacing incumbent storage vendors. This course will arm you with important information so that you are comfortable talking at a consultant level with your customers, understanding what the problems are up front, helping them articulate how we solve these issues and even and running through some concepts on your own for those who wish to do so.



### High-Level NetApp Value Proposition

- Reduced cost (see [Oliver Wyman](#), and [Mercer](#) studies)
- Reduced risk
- Faster implementation


© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 9

## High-Level NetApp Value Proposition

As you are focusing on the technical aspects of virtualization, it is important to keep your role as a pre-sales engineer in mind. The NetApp message can be summarized by three high-level message points:

- Reduced cost
- Reduced risk
- Faster implementation





### Mid-Level NetApp Value Proposition

- Increased Storage Utilization & Data Protection
  - Unmatched storage utilization
  - Unmatched data protection
- RTOs reduced to minutes
  - Policy driven backups
  - Recovery within minutes
- Increased Operational Efficiencies
  - Make storage virtualization as easy as server virtualization

**See *VMware on NetApp Technical Diamonds* Presentation**

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 10


## Mid-Level NetApp Value Proposition

The corresponding mid-level message points for the managers and decision makers you are likely to deal with are:

- Increased storage utilization and data protection through unmatched storage utilization and unmatched data protection;
- Restore time objectives (RTO) reduced to minutes; and
- Increased operational efficiencies—basically extending all of the NetApp simplicity and management values into the virtualization environment to make it simpler to run your virtualization environment rather than harder

If you focus on speaking to customers at this level, the rest will come naturally.

We'll talk about each of these points in a little more depth.



## Unmatched Storage Utilization

VMware storage is redundant by design. The upper half of this slide is a representation of a VMware server and four virtual disk files or VMDKs. Those files and their content map directly to the storage array below it. VMware has the ability to copy one of those data files and assign it to a new virtual machine and deploy it from a template or a clone. Twenty, thirty minutes later you have a new server up and running. From an operational standpoint, that's fantastic. Giving the VM administrator the capability to deploy servers as he needs them in 30 minutes is a very good story. As a NetApp SE, you don't want to disrupt that model. But you do want to point out that VMware is putting the same data over and over and over again. Isn't there something we can do to make that more intelligent?

**NetApp deduplication Is The Only Deduplication Technology Available For Production Data**

- FAS Dedupe
  - Reduces storage by 50%
  - Your mileage may vary
  - (40-70% seen)
- Change the cost per GB
  - Customer data is more valuable than software binaries


**NetApp Deduplication is the only dedupe technology for production data sets!**

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 12

## NetApp Deduplication Is The Only Deduplication Technology Available For Production Data

This is a perfect opportunity for deduplication since NetApp deduplication, formerly known as A-SIS, is available for primary data—production data—not just the backup copy. NetApp Deduplication allows you to collapse redundant data. In fact, the more redundant your data, the better deduplication works. Since you have more iterations in desktops than you would in servers, the deduplication ratios are going to be much higher. Deduplication is going to be much more efficient, and save even more money for your customers.

NetApp is the only storage vendor offering deduplication for both production and disaster recovery. Since NetApp's deduplication delivers 50% storage savings, your customer is going to be able to store their production copy and their disaster recovery copy in the same net amount of physical storage that traditional storage would require for just the production copy. NetApp Deduplication also reduces that data that has to be transferred by SnapMirror



## Unmatched Data Protection

- The Problem
  - Double disk failure is a mathematical certainty
  - RAID 5
    - Insufficient protection
  - RAID 10
    - Double the cost
- NetApp RAID-DP™ Solution
  - Protects against double disk failure
  - High performance and fast rebuild
  - Same usable capacity as RAID 10 at half the cost
  - See [TR-3515](#)

	RAID 5	RAID 10	RAID-DP
Cost	Low	High	Low
Performance	Low	High	High
Resiliency	Low	High	High

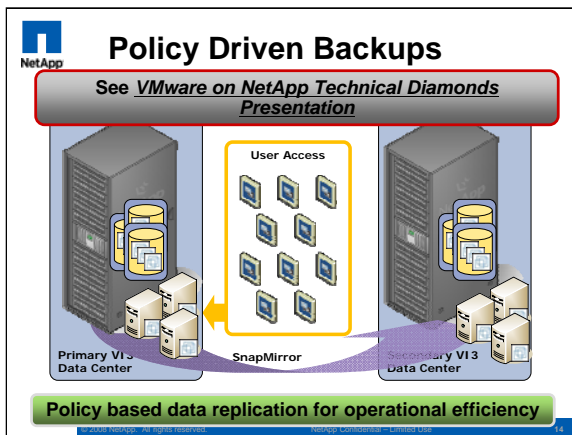
© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 13

## Unmatched Data Protection

And then of course NetApp has RAID DP. Data protection is especially important in a virtualized environment because now you are stacking more servers and more users and more applications on the same physical storage. Only NetApp can give that customer RAID 10-level protection, but at half of the cost of equivalent traditional storage.

In a traditional distributed desktop environment, a failed desktop affects exactly one user. By contrast, if you have hundreds or thousands or tens of thousands of virtual desktops on a centralized device and it fails, you will have tens or hundreds or thousands or tens of thousands of users who are down. A double disk failure would be a really inexcusable reason to send thousands of people home for the day. So a shared virtual desktop environment like this has to be reliable. It has to be much more reliable than the physical infrastructure that it replaces.

RAID 5 data protection isn't sufficient protection because a double disk failure will result in losing your data and all your desktops are going to be down. Other storage vendors will look at RAID 10 as an alternative to this. It takes twice as many spindles to run a RAID 10 array group as it does a RAID 5 array group, so the storage cost is going to go up. But, the whole point of this virtualization exercise is to reduce those costs. By contrast, NetApp's RAID-DP gives you the data protection of RAID 10 at the cost of RAID 5, while providing the highest level of availability for this environment.




## Policy Driven Backups

One of the important distinctions between traditional and virtualized server environments is the use of policy driven backups. Virtualized environments are not going to backup at the virtual machine level like you would with a traditional architecture. Rather, they are going to backup at the Datastore level or even multiple Datastores at the same time. In a virtualized environment, a traditional backup model like tape can result in very very lengthy recovery times.

In a traditional backup environment you are backing up at a server level. If you take a traditional backup model into a virtualized environment, you would be backing up virtual machines individually—pushing a backup client out to every virtual machine with the same number of backup schedules as before and all of the IO issues that would come along with that. When you backup at the Datastore level backup works the same regardless of the number of servers that you have in your environment. And, you don't have to worry about pushing out clients, or about scheduling individual backups. You don't have to worry about the IO limitations of trying to pull all of that data off of your systems every day

This slide depicts policy driven backups between two systems at the datastore level. Users are all hitting the primary side. And this is how quick if we have a failure we can bring everything over to the other side.



## Recovery Within Minutes

- Snapshot copies for VM level recovery
- Immediate VMware backups with Snapshot
  - SnapManager for Virtual Infrastructure
- Full data center RTO from DR site in minutes
  - SRM enhances this solution
  - FCP & iSCSI support
    - NFS planned for a future version

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Life 18/19


## Recovery Within Minutes

This slide tells the NetApp VMware replication story in a nutshell.

- Snapshot for immediate policy driven backups to provide virtual machine level restoration locally.
- SnapMirror to get it offsite—SnapMirror, not SnapVault because SnapMirror allows you to run off the DR side in minutes. SnapMirror also allows you to re-sync if there really is a failure event. And SnapMirror allows you to run anything from one VM at another location, to a datastore, to the entire data center at the second location.

Both of these tools, Snapshot and the SnapMirror can be done with scripts or using the recently released SnapManager for Virtual Infrastructure.

The second piece is VMware Site Recovery Manager or SRM which is a GUI for managing the breaking of mirrors and starting up virtual machines on another location. VMware's SRM works with 11 storage array vendors, so it is not unique to NetApp. Rather, SRM integration is a minimum functionality for playing in the game. Combining SRM with SnapManager, however, is a NetApp value proposition because NetApp's offering ensures integrity of the data at the remote site when you break the mirror, something that Site Recovery Manager does not do alone. When SRM is used with SnapManager, the system administrator is assured when he has to break the mirror in a DR scenario and start anywhere from tens to hundreds of virtual machines that they're going to boot without having to check-disk every block on the storage array? There's no value in having the system back up in five minutes if the next five hours you have to read all your data before it can be used. That's the value of integrating SnapManager with Site Recovery Manager.



### Increased Operational Efficiencies


- NAS treats storage like a resource pool
- Dynamic fluid usage without VMware admin operations
  - SANs require configuration on every server
  - No orphaned storage as with SANs
- Virtualization is experienced by the VMware admin team
  - Check and balance implementation of technology

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 18

## Increased Operational Efficiencies

Most customers provision a fixed sized datastore. For example, they might standardize on 400 gig for every datastore running VMFS over fiber channel. They add virtual machines into a datastore until it is nearly full or until they start to have a performance problem. Then they back out a little bit and you create another datastore and start writing virtual machines into that one, and so on. This practice can become very wasteful very quickly. When you add up 20, 30, 40 gig here and there in each one of these Datastores you can start wasting space very rapidly.

With NAS you can size data storage to exactly what is needed. If it is too big, you can just shrink it. Additionally, when you provision storage in a NAS environment, it is available immediately. What you see is what you get. So if the virtualization team has free space it is because it is free in the volume. When you provision new storage, they realize it immediately. In a SAN environment, there has to be an interaction between the storage administrator and the server administrator. If they want more storage it has to be provisioned as a new LUN.



## NetApp Value Proposition Summary

- High-Level NetApp Value Proposition
  - Reduced cost
  - Reduced risk
  - Faster implementation
- Mid-Level NetApp Value Proposition
  - Increased Storage Utilization & Data Protection
  - RTOs reduced to minutes
  - Increased Operational Efficiencies

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 17

Now that you have completed the first module, you should be able to explain to a high-level customer audience how NetApp’s virtualization technology is the right business decision because of its reduced cost, reduced risk, and is faster implementation when compared to the competition. And you should be able to explain to a mid-level audience how NetApp virtualization solutions make their life easier by increasing storage utilization and data protection, reducing recovery time objectives to minutes, and increasing operational efficiencies.





## Resources

- [VMware on NetApp Technical Diamonds Presentation](#)
- Network Appliance and VMware ESX Server 3.0: Building a Virtual Infrastructure from Server to Storage ([TR-3515](#))




## Resources

- [VMware on NetApp Technical Diamonds Presentation](#)
- Network Appliance and VMware ESX Server 3.0: Building a Virtual Infrastructure from Server to Storage ([TR-3515](#))



## VMware on NetApp Advanced Training: Connectivity

The remainder of this module reviews basic VMware terminology, and then describes VMware connectivity to NetApp storage.



## Learning Objectives

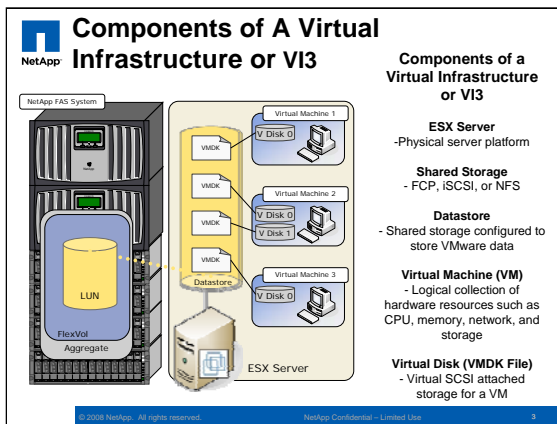
- VMware Virtual Infrastructure components
- VMware Virtual Infrastructure Suite features
- VMware Service Console, VMkernel, VMware File System (VMFS) and VMtools
- VMDK files
- VMware hardware- and software-based connectivity
- VMware storage access concepts

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Learning Objectives

This module reviews basic VMware terminology, and then describes VMware connectivity to NetApp storage. There is also a lab at the end of this module. When you have completed this module you should have a basic understanding of:

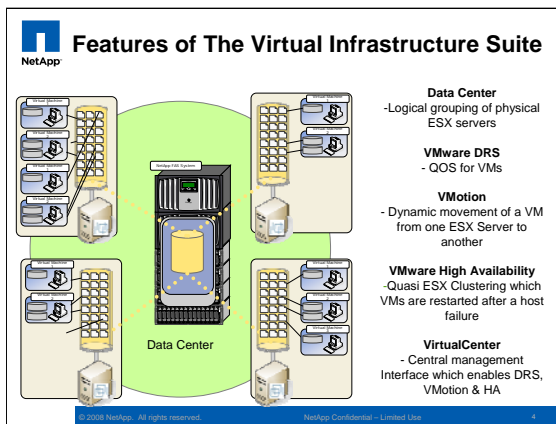
- The essential components of a VMware virtual infrastructure such as datastores, virtual machines, and virtual disks;
- The most important features of the VMware virtual infrastructure suite such as DRS, VMotion, and VirtualCenter;
- The purpose of the VMware Service Console, VMkernel, VMware File System (VMFS) and VMtools;
- The significance of VMDK files.
- VMware hardware- and software-based connectivity; and,
- VMware storage access concepts such as hardware- and software-based multipathing, and VMware path management policies.



## Components of A Virtual Infrastructure or VI3

This slide shows some of the components of a VMware deployment:

- The **ESX server** is the physical server running the VMware operating system. One of the requirements of VMware is a shared storage environment. That could be using the fiber channel protocol, iSCSI or NFS, preferably to a NetApp storage system.
- A **datastore** is the container in which the virtual machines reside. You start with the storage container, which would be a LUN or an NFS volume. Then you either format the LUN with VMFS or mount it as NFS and it becomes a datastore that can house virtual machines.
- A **virtual machine** is a logical collection of emulated hardware resources. It is CPU, memory, disk, and network connections. And, each virtual machine has one or more virtual disks. In a VMware environment, virtual disks are represented as and stored as a file, which we commonly refer to as a VMDK file. Regardless of whether you are going to use fiber channel, iSCSI or NFS, at the end of the day, you have a virtual disk talking to a VMDK file. There is always a one to one correlation between virtual disk and VMDK file, with the exception of RDMs, which we will talk about shortly.



## Features of The Virtual Infrastructure Suite

This slide shows some features of the Virtual Infrastructure suite.

A collection of multiple ESX servers with multiple virtual machines all sharing data is referred to as a Data Center. It is basically a management boundary.

DRS is Dynamic Resource Scheduler and is based on real time load on the systems. It moves virtual machines from server to server to balance workloads. If a server's workload reaches a predetermined threshold, it triggers an alert—a DRS policy—which will institute VMotion.

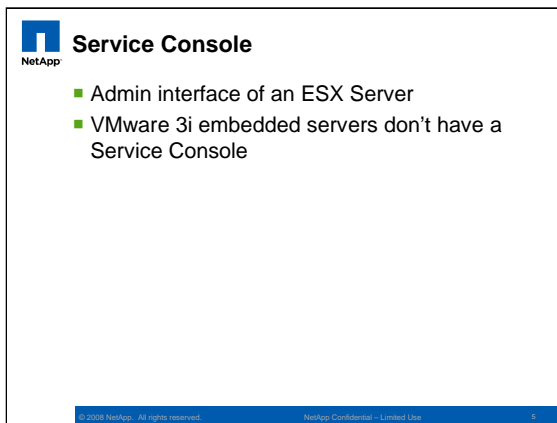
VMotion is real time migration of the running state of a virtual machine from one ESX server to another. In real time, if work load is too high on an ESX server, one or more virtual machines are moved to another server. It can always be done manually. DRS is an extra layer of management that provides automatic management of those virtual machines. The administrator doesn't have to make decisions about where the virtual machines should go. You just create a pool, create virtual machines and let the DRS policies deal with it. This is very common in customer environments. One of the reasons NetApp products have to be VMotion aware is because we can't assume that a particular virtual machine is going to be running on a particular server at any given time.

VMware High Availability is a bit of an overstatement for what this product actually does. In other words, it does not provide stateful failover. In the event that a server in a VMware HA cluster pair fails, affected virtual machines are automatically restarted on its partner, but the failover is not instantaneous and does not maintain the state of running applications and other services. This is a cold restart of the failed virtual machines to automatically restores service.

VMware HA and DRS work together to balance resources. If after a failover, the server taking over for the failed server becomes overloaded, DRS spreads this load around to the other physical servers in a VMware Infrastructure resource pool that have spare capacity.

True HA with instantaneous stateful failover of virtual machine services from one to the other, may be in the next major release of VMware. Keep in mind that it is up to the customer and the VMware SE to decide how much extra overhead to build into a VMware Data Center environment.

VirtualCenter is the management overlay on top of the Virtual Infrastructure. The VI client will let you connect and manage each of these ESX servers individually; but, the major value-added applications—HA, VMotion, DRS, VMware cloning, and a couple of other things—only work through Virtual Center, which runs on Windows as a thick client. Most customers are going to want Virtual Center because at the very minimum, almost everyone will want the capability to dynamically move ESX servers using VMotion.

A slide titled "Service Console" with the NetApp logo in the top left corner. The slide contains two bullet points: "Admin interface of an ESX Server" and "VMware 3i embedded servers don't have a Service Console". At the bottom of the slide, there is a blue footer bar with the text "© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use." and a small number "5" on the right side.

**Service Console**

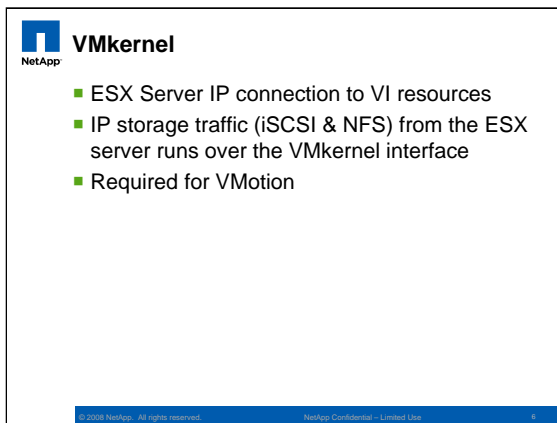
- Admin interface of an ESX Server
- VMware 3i embedded servers don't have a Service Console

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use. 5

## Service Console

The Service Console is the Linux shell ESX server administrator interface. It gets its own IP address. However, VMware 3i servers don't have a service console. VMware 3i is ESX on a chip.. Customers will order a server from a traditional server vendor such as Dell or IBM with the VMware hypervisor installed in firmware the motherboard. The administrator boots to a bios configuration prompt, plugs in a couple of parameters—an IP address, the IP of the VirtualCenter and a couple of other thing—and they have an ESX server. 3i servers don't have a service console.

3i servers don't have a Service Console. This is potentially a problem for NetApp products in the short term, because some NetApp functionality uses the Service Console. For example, OSSV 2.6, which is OSSV for ESX, runs entirely in the Service Console today. If your customer is excited about 3i, it might be a good idea to start qualifying what sort of things they are going to want to do, and whether that is going to work without a Service Console. For the short and medium terms, ESX will be released in two flavors—traditional VMware ESX, which runs on top of Linux, and VMware 3i.

A slide with a blue header containing the NetApp logo and the title "VMkernel". Below the title is a bulleted list of three items. At the bottom of the slide is a blue footer with small white text.

**VMkernel**

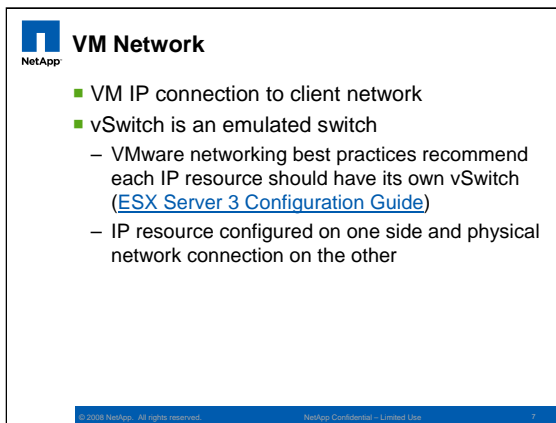
- ESX Server IP connection to VI resources
- IP storage traffic (iSCSI & NFS) from the ESX server runs over the VMkernel interface
- Required for VMotion

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMkernel

The VMkernel is a proprietary IP connection between ESX and other virtual infrastructure resources. Essentially, IP storage traffic from the ESX server runs over the VMkernel interface. The VMkernel is important to NetApp storage because the software iSCSI traffic and the NFS stack of ESX all run over the VMkernel. It is required for VMotion, so most customers will have the VMkernel already configured on their ESX servers.





**VM Network**

- VM IP connection to client network
- vSwitch is an emulated switch
  - VMware networking best practices recommend each IP resource should have its own vSwitch ([ESX Server 3 Configuration Guide](#))
  - IP resource configured on one side and physical network connection on the other

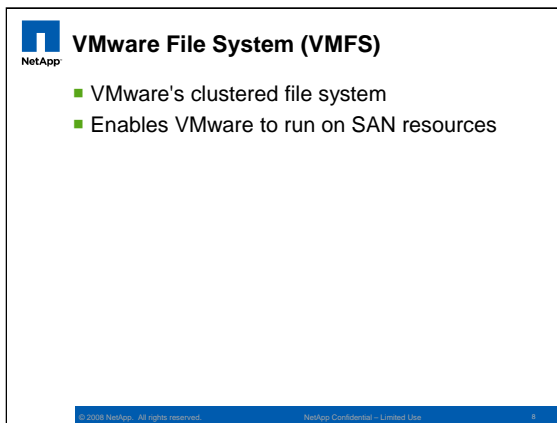
© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VM Network

The VM network, or the virtual machine network, is a separate network that is dedicated for traffic between the virtual machines and the public network. The Service Console has an IP address on the VM network as well as the VMkernel and each virtual machine. You could certainly access all of those resources over one physical connection to an ESX server

VMware best practices say you give each of these resources—the Service Console, VMkernel, or virtual machines etc.—its own vSwitch ([ESX Server 3 Configuration Guide](#)-[http://www.vmware.com/pdf/vi3\\_35/esx\\_3/r35/vi3\\_35\\_25\\_3\\_server\\_config.pdf](http://www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_3_server_config.pdf), page 68). A vSwitch is basically an emulated switch. You configure one of these resources on one side and a physical network connection on the other. VMware best practices say a separate vSwitch and a separate physical connection for each one of these resources. In most deployments, you are talking about a minimum of three connections and that is even before you have any redundancy. Most customers will have many physical connections to many different networks.

Some customers may lack the physical infrastructure or server interfaces to completely separate all network traffic and this is acceptable. Whenever possible separate front-end (VM network) and back-end (VMkernel network) traffic.




**VMware File System (VMFS)**

- VMware's clustered file system
- Enables VMware to run on SAN resources

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMware File System (VMFS)

VMFS is VMware's cluster file system. We have talked about this quite a bit already. It enables VMware to run on SAN resources and to write the files and encapsulate the virtual machine. Note: If you run NFS, then there is not VMFS file system.




## VMTools

- VM drivers and management APIs
- Required for Vmotion
- VirtualCenter uses VMTools to gather information about the virtual machines
- VMware Virtual Machine Snapshot and NetApp Snapshot backup process leverage VMTools to freeze a VM's file system

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMTools

VMTools is a set of drivers and management APIs that are installed on the guest operating system that runs on the virtual machine. Virtual Center gathers information about the virtual machine—such as its IP address—using VMtools. The sync driver makes it possible to freeze a virtual machine's file system. VMware Snapshot uses this capability and NetApp's Snapshot feature in turn leverages VMware Snapshot to freeze the file system before taking a snapshot. It is also required for VMotion because of the syncing and freezing that goes on during the VMotion migration..



**VM Files**

- VM Configuration File (\*.vmx)
  - Text file containing a VMs configuration data
- Virtual Disk Descriptor File (\*.vmdk)
  - Used for ALL VM to storage connectivity
- Virtual Disk Data File (\*.flat.vmdk)
  - Flat file accessed by VMs via scsi commands
- RDM (Raw Device Mapping)--Direct access from VMs to LUNs

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use. 16

## VM Files


VMware creates three files for each virtual machine.

The first VM file is the configuration file, which has the .vmx filename extension. It is a flat text file that contains descriptive information about the virtual machine including it's name, emulated physical resources—such as memory, CPU, disk types, and network connections—how those network connections are configured, and the path to the VMDK—virtual disk— files. All of that information is stored in the .vmx file. This file also stores a path to the virtual disk (.vmdk) files.

And, there are two virtual disk (.vmdk) files. If you browse the datastore with the VI client, only the Virtual Disk Descriptor file is visible. This files has just the .vmdk filename extension and is a descriptor file, not the actual data. It stores the path to the Virtual Disk Data File, as well as meta data about that disk--basically the geometry information and anything else about that virtual disk. T

The second virtual disk file, the Virtual Disk Data file, is hidden and contains the data. It has the same name as the Virtual Disk Descriptor File but has a --flat.vmdk" filename extension. The .vmdk file (the Virtual Disk Descriptor file) is very small and the --flat.vmdk file (the Virtual Disk Data file) is large—the size of the virtual disk data. This data file is not visible when browsing the Datastore using the VI client, but you will be able to see this file from the service console

Lastly, RDM, Raw Device Mapping, is the only storage configuration that doesn't use a virtual disk data file. It does use a descriptor file that contains the path of the physical raw device, but it doesn't use a virtual disk data file. RDM presents a LUN directly to a virtual machine. It still emulates the SCSI protocol, but at the back end instead of having a file that is formatted with the VFMS file system this configuration actually communicates directly to a LUN.



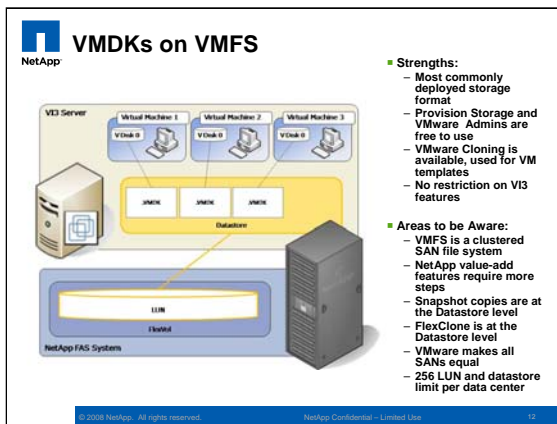
## VMware Storage Designs

- When it comes to storage there are more decisions to be made than just protocol (FCP, iSCSI or NFS)
  - Virtual Disks (VMDKs)
    - Support over VMFS (FCP or iSCSI)
    - Support over NFS (IP)
  - Raw Device Mappings (RDMs)
    - Support over VMFS (FCP or iSCSI)

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 11

## VMware Storage Designs

There are a couple of considerations when you are designing storage for VMware. We know about the three storage protocols—fiber channel, iSCSI and NFS. But, there are two ways to store data—in virtual disk files, or using RDM to present LUNs directly to the virtual machines. When using RDM, the connection to the LUNs can be fiber channel or iSCSI. When you connect to an RDM, VMware uses the ESX IO stack rather than having the iSCSI initiator reside inside the virtual machine. This is a valid configuration for the root drives and everything else, but certainly not a common configuration



## VMDKs on VMFS

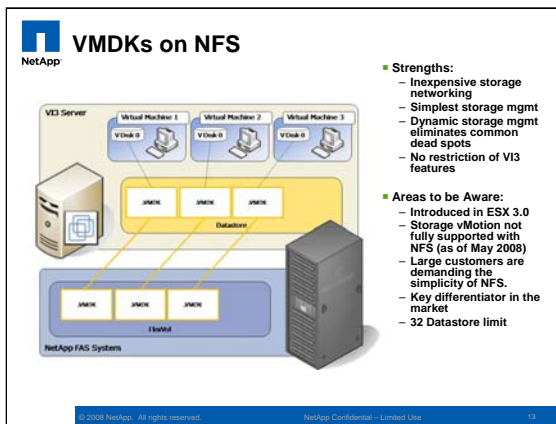
This slide depicts 90-95% of the VMware storage deployments today. It uses VMFS over fiber channel or iSCSI. This design shows a LUN that is formatted with VMFS, with 3 virtual disks. Most existing ESX systems are configured like this because prior to ESX 2.5, this was the only configuration available. NFS and iSCSI are still fairly new—shipping with 3.0. Most legacy deployments look like this.

Most VMware administrators and storage administrators are going to be familiar with this design. The storage administrator presents LUNs which the VM administrator formats and uses for virtual disks and virtual machines. When the Datastore is full, the VM administrator requests another LUN and starts the process again. Administration can be done within the VI GUI. Cloning is done on a per virtual machine level as a copy out. For example, you can make a clone of virtual machine three and the ESX server will copy out all of the files and write them new within the Datastore.

From a NetApp storage management perspective, there is one obvious disadvantage to this configuration—lack of granularity within the Datastore. Provisioning, cloning, FlexClones, and creating and restoring Snapshot copies are all done at the LUN level. If you want to restore only one virtual machine, it is more work because you have to mount a clone, and copy the data.

To some extent, VMware makes all SANs equal in that SAN storage from any vendor will be running VMFS. However, with NetApp SAN storage you get NetApp Snapshot backups and recovery—albeit with a bit more time and effort—and you can do SnapMirror or SnapVault replication.

This configuration is limited to 256 LUNs, and 256 datastores per data center—whether you have 1 ESX server or 10 ESX servers in the data center. A Datastore could be more than one LUN, but in most cases you will have a one to one correlation. If you need more than 256 datastores, you can start another data center. But, you cannot share any of the resources from one data center to another.



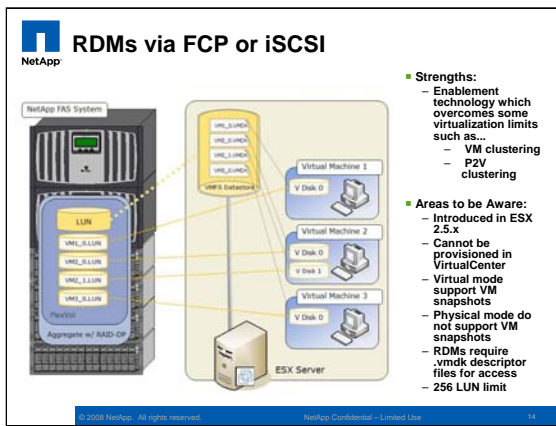
## VMDKs on NFS

There is a subtle difference between this slide, which depicts VMDKs on NFS, and the preceding slide, which depicted VMDKs on VMFS. In this NFS configuration there is no LUN. The NetApp FAS system is presenting a FlexVol directly to the ESX server via an NFS mount. Now the server is writing files directly into the NetApp storage system volume using IP instead of fiber channel.

Storage management from the VMware side works exactly the same way whether using NFS or VMFS. Once the volume is mounted, a VMware administrator probably won't even notice the difference. From the storage administrator's perspective, he no longer has to provision LUNs. Instead, storage can be grown dynamically as needed. In addition, all V13 features such as VMotion, DRS, and HA are still available. If there is any concern about this configuration, it is that NFS support is relatively new. But we are aware of no problems with running VMware storage on NFS.

The maximum number of datastores in an NFS configuration is 32 mounts for 32 Datastores—again at the data center level. All of your servers in the data center are limited to 32 datastores. If you have small servers, start another data center. For most customers, it should not be that big of a problem.

The datastore limit has not been raised in VMware 3.5. NetApp has formally requested that it be increased. VMware is looking at 64 in a future release of ESX, but not in 3.5. Commonly recommended best practices include maxing around 20 virtual machines per Datastore, so take that number times 32 datastores can handle a lot of virtual machines. If you need more than that, just start another data center. Most big customers will have other reasons to run multiple data centers anyway. The only significant disadvantage to multiple data centers is that a server can only be in one data center or another. VMotion cannot move a virtual machine from one data center to another. Most customers are going to be small enough that they will never bump up against this or big enough that they are going to want multiple data centers anyways.



## RDMs via FCP or iSCSI

RDM is raw device mapping uses Fiber Channel or iSCSI connectivity. In this scenario, the FAS System provisions LUNs. One of the LUNs is configured as a VMFS datastore that contains virtual disk pointer files. The other LUNs are connected directly to the virtual machines. Each virtual machine points to a virtual disk pointer file in the VMFS datastore that in turn, redirects to a LUN on the NetApp storage system. This scenario should get higher performance to the virtual disk because the dataflow from the virtual machine to the storage does not include VMFS. You're removing a layer of abstraction, especially in a scenario where you have lots of virtual disks within the same file system. VMware recommends this to high performance customers. When the data is being written out to individual LUNs, we now can use the NetApp side statistics to measure the input and output for each one of those virtual disks individually. So we can learn a lot more about the environment without necessarily having to go into the ESX server itself in this scenario.


The ESX features like VMotion, DRS, HA, all work with RDM too

RDM is required for running MSCS on virtual machines today and we expect it to continue that way. Also when you're using RDM in the virtual mode, you still have the ability to use some of the VMware special features, such as ESX Snapshot and VMotion.

We don't recommend RDM as much as we used to. In ESX 2.5 the NFS option was not yet available. RDM was the only way we could do virtual machine operations from the storage side. In 2.5 RDM was required to FlexClone a virtual machine. One of the major problems with RDM is that the Virtual Center GUI doesn't support it. It can take quite a bit more work.



Lastly, this configuration uses LUNs, so 256 LUNs per Datastore is the max. There are still some pretty cool scenarios if you have the need to very rapidly provision large numbers of virtual machines, or very rapidly clone large numbers of virtual machines. This is a pretty cool solution for that because you can do LUN clones or FlexClones on these LUNs and very quickly spawn off virtual machines.





## Hardware Based Connectivity

- FCP
  - Support for Qlogic and Emulex Adapters
  - Supports boot from SAN
- iSCSI
  - Support for Qlogic only
  - Supports boot from SAN
- Host Attach Kit for ESX 3.X can ease configuration
  - Qlogic and Emulex HBAs only
- Native Multipathing
  - Built into ESX kernel
- Supported Adapters Changes Regularly. Always check the latest support matrices.
  - See [NetApp Interoperability Matrix](#) and [VMware Compatibility Guides](#)

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use. 18


## VMDKs on VMFS

Fiber channel connectivity is certified for Qlogic and Emulex. And ESX does support boot from SAN with fiber channel. NetApp certifies an ESX version, Data ONTAP version and HBA. VMware certifies those same things—ESX version, Data ONTAP version and HBA. VMware also certifies by FAS product family. VMware actually certifies FAS 3000 separately from FAS 6000, separately from FAS 200 and FAS 2000. As you well know, it is the same stuff. We make no support distinctions. The problem is that if there were ever a serious issue, VMware does have the right to make a support distinction about the product family that it is running on.

iSCSI hardware is certified on Qlogic 4052 only, and you can also boot it from the SAN, but Qlogic is dropping the card, with no plans to replace it. Yeah, so that is what it seems to be. Hopefully the software administrator becomes robust enough.

Host attach kit for ESX3 is a NetApp product that works only with Qlogic and Emulex HBAs to automatically set the tunables to the appropriate NetApp and VMware recommended settings. It also does path manager, which we will cover in a couple of slides.

ESX has native multi-pathing. It is active-passive multi-pathing. It can't interpret a NetApp partner path. It sees all paths as equal, and it round robins LUNs across available paths. Universally, half of LUNs are going to be on a partner path in a configuration. Thankfully, we have the auto support now, but it is a problem because if a customer is not aware of it, and they just set it up, half of their LUNs are going to be on a partner path. You can go in and set the primary paths manually or you can use the attach kit which will go and communicate with the filer via R shell and then rebalance everything across primary paths only. For large fiber channel deployments, make sure your customer gets this. The R shell security thing might be a chore, but it is going to save them a lot of trouble.



## Software Based Connectivity


- NFS
  - Built into VMware application
  - Requires VMkernel
- iSCSI software initiator
  - Built into VMware application
  - Sourceforge iSCSI initiator (3.5)
  - Requires VMkernel and Service Console
  - iSCSI session management over service console
  - Native multipathing in 3.5 and later only

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 16

## Software Based Connectivity

NFS is native to VMware. It does require the VMkernel. VMkernel is also required for VMotion. 90% of the time you ask if a customer has VMkernel configured on your servers and they will say yes because they use it for VMotion.

iSCSI is also native to VMware, ESX. It is the Cisco iSCSI software initiator. It also uses VMkernel connection. Same thing. Customer may already have VMkernel, but now link resiliency and all of that may come into play. No native multi-pathing in this software environment.




### ESX Storage Summary

	Hardware		Software	
	FCP	iSCSI	iSCSI	NFS
Boot ESX	✓	✓		
VMDK	✓	✓	✓	✓
RDM	✓	✓	✓	
VMotion	✓	✓	✓	✓
Native Multipathing	✓	✓	✓	
VMkernel			✓	✓

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use. 19

## ESX Storage Summary

Hardware adapters, fiber channel, and iSCSI can boot ESX. Software protocols, iSCSI and NFS, cannot. All can support virtual disk files. All except NFS can support RDMs. VMotion works regardless of the protocol. Native multipathing is available with the hardware-based protocols, and the VMkernel is required for the software-based protocols.



### Hardware Storage Multipathing Summary

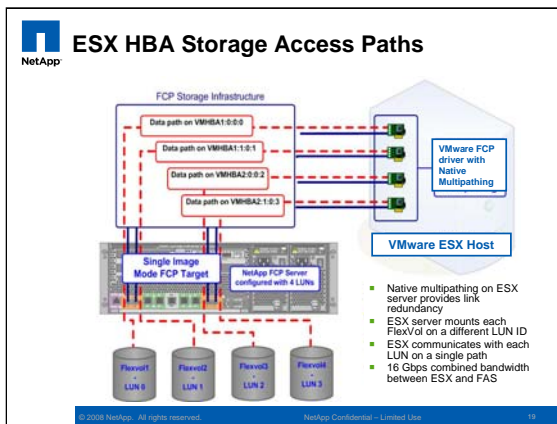
- iSCSI Hardware Initiator:
  - Use native multipathing for link redundancy
  - On FAS use VIFs for link redundancy unless using target HBAs
  - Maximum of 1 Gbps throughput to each storage controller
- FCP:
  - Use native multipathing for link redundancy
  - Maximum of 2/4 Gbps throughput to each datastore
    - Many datastores may be required to fill bandwidth

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 18

## Hardware Storage Multipathing Summary

If you use iSCSI hardware initiators, and assuming that you are going to have two adapters or a dual port card, you can use the native multipathing for link redundancy. If you are using software target on the FAS system, use trunks or use VIFs for link redundancy, unless you are using target HBAs. There is a maximum 1 gigabit per second throughput to each storage controller per ESX server. That is very important.

With fiber channel, use native multipathing for link redundancy. There is a maximum 4 gigabits per second throughput to each serer.




## ESX HBA Storage Access Paths

Let's look at connectivity. Assume that you have an ESX server with fiber channel connections to 4 FlexVols, each containing a LUN configured as a datastore, all stored on a FAS 3000 storage system. You have 4 potential paths, 4 potential targets, and 4 initiators. You can manually set the primary path for each datastore to be its own target and can configure this for the maximum theoretical throughput, with each LUN having its own target port and its own initiator. Provided you have a minimum of at least one datastore for every physical connection, you can go active-active across all of the connections. It is not active-active on a per connection basis. It is still active-passive. One path is active, and everything else is still a passive path. This is all being done with native multipathing on the storage side. Each ESX server mounts each FlexVol on a different LUN ID, single path.

In this scenario, these are all 4 gig connections, total theoretical bandwidth between the ESX server and the storage device is 16 gig. That is pretty cool. The host utilities kit in this scenario would see 4 targets, and it would round-robin the LUNs around the 4 targets.

**NetApp Path Management Policies**

- VMware only allows traffic to a datastore over a single path (regardless of protocol)
- Fixed
  - Recommended with NetApp Storage
  - Resumes using the preferred path when connectivity is restored
    - Prevents LUN from remaining on a proxy path after cluster failback
- Most Recently Used (MRU)
  - Administrator action required to fail back to primary path
  - Required for MSCS VM clusters




© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 20

## Path Management Policies

ESX gives you the option to either use the host attach kit or to manage paths manually. By default, if you don't use the host attach kit, ESX will see all of the paths, both the primary paths and the partner paths and will round-robin select the primary path for each LUN. The odds are that 50% of your LUNs will be on the wrong path. To avoid this problem, you can either set the paths manually or you use the host utilities kit.

Of course, any time your fabric changes, the path management is going to change. If you are manually managing it, you have to look at everything again. If you are running the host utilities kit, you just run the configuration script one more time and it will automatically rebalance paths.

ESX has two path management policies: fixed and most recently used (MRU). Fixed means that if a path fails, the LUN will select another path as a fail over. If the original path becomes available again, the LUN will switch back, which is probably the behavior you want when you have a cluster fail over event. With the MRU setting, you would have to manually move the paths back. NetApp recommends using the fixed setting in all scenarios except when you are using MSCS on your virtual machine in which case MRU is required.



### ESX Network Design for IP Storage

- The Goal: To design a network that is
  - redundant across physical switches
  - uses multiple physical paths simultaneously
  - can scale to additional physical interfaces
- Two high level options
  - Cross-stack etherchannel - one storage subnet, multiple storage IPs, and IP load balancing
  - or-
  - Without cross-stack etherchannel - two storage subnets, multiple storage and host IPs

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 21

## ESX Network Design for IP Storage

When designing an ESX network for IP storage, the goal is a network that is redundant across physical switches and that can use multiple physical paths simultaneously. You also want a network that can scale to additional physical interfaces. There are two high-level ESX network designs that will meet this goal—one with and the other without cross-stack Etherchannel.

The cross-stack Etherchannel design uses one storage subnet, with multiple storage IPs, and IP load balancing

The design that uses Etherchannel without cross-stack uses two storage subnets, with multiple storage and multiple host IPs.

**IP Storage and Cross-stack Etherchannel**

- Multiple storage IPs are required
- ESX host requires one VMkernel port

Active Path  
Passive Path

Primary data path of datastores on IP1  
192.168.1.201  
192.168.1.202

Primary data path of datastores on IP2  
192.168.1.203  
192.168.1.204

ESX1  
Virtual Switch vSwitch1  
VMkernel 1  
192.168.1.201  
VMkernel 2  
192.168.1.202  
vnic1  
vnic2

ESX2  
Virtual Switch vSwitch1  
VMkernel 1  
192.168.1.201  
VMkernel 2  
192.168.1.202  
vnic1  
vnic2

Ethernet Infrastructure

FAS1  
vif1 (IP: 192.168.1.201)  
vif2 (IP: 192.168.1.203)

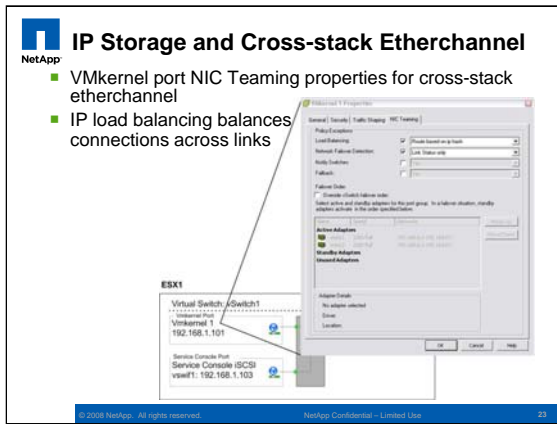
FAS2  
vif1 (IP: 192.168.1.202)  
vif2 (IP: 192.168.1.204)

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 22

## IP Storage and Cross-stack Etherchannel

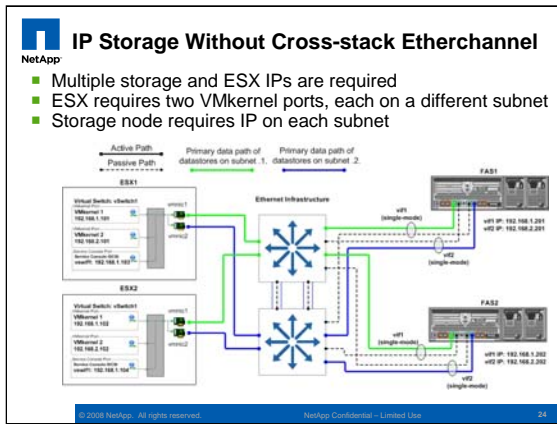
This slide shows IP storage and cross-stack Etherchannel. Notice that multiple storage IP addresses are required but the ESX host needs only one VMkernel port. Multi-mode vifs provide link redundancy on the storage side.





## IP Storage and Cross-stack Etherchannel

When using the cross-stack Etherchannel design, open the NIC Teaming tab on the VMkernel port Properties screen and select **Route based IP hash** load balancing and **Link status only** network failover. The ESX server will then IP load balance connections across the available links.



## IP Storage Without Cross-stack Etherchannel

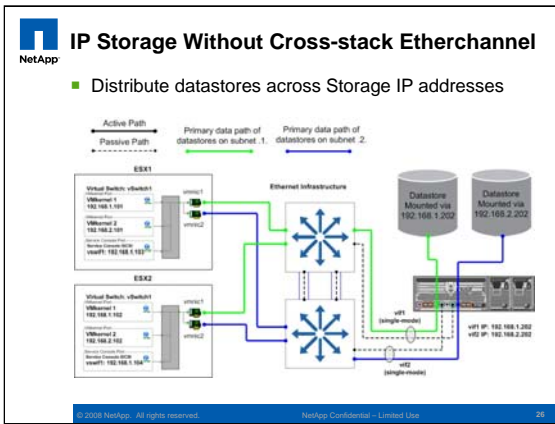
By contrast, the IP Storage design without cross-stack Etherchannel uses multiple storage IPs and multiple ESX IPs. There must be two VMkernel ports on the ESX server, each on a different subnet. In addition, the storage node needs an IP address on each subnet. This design uses single-mode vifs between the storage controllers and the Ethernet infrastructure.

**IP Storage Without Cross-stack Etherchannel**

- VMkernel port NIC Teaming properties for two VMkernel ports
- Connections are manually balanced by selecting a different Active Adapter for each VMkernel port

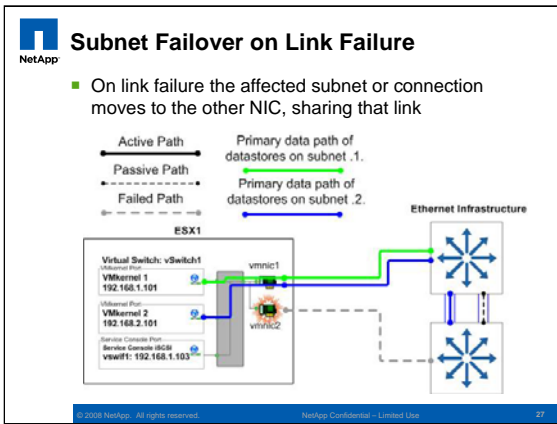
## IP Storage Without Cross-stack Etherchannel

When using the IP storage design without cross-stack Etherchannel, open the NIC Teaming tab on each VMkernel port Properties screen and select the **Link status only** network failover option. Manually balance connections by first selecting **Override vSwitch Failover Order** and by selecting a different Active Adapter for each VMkernel port.




## IP Storage Without Cross-stack Etherchannel

This slide shows IP storage and cross-stack Etherchannel with datastores distributed across storage IP addresses.



## Subnet Failover on Link Failure

On link failure the affected subnet or connection moves to the other NIC, sharing that link



## IP Storage Multipathing Summary

- iSCSI Software Initiator:
  - Service Console must be able to communicate with iSCSI target (HW & SW)
  - ESX 3.0.2 limited to 1 GbE throughput to each controller
  - ESX 3.5 Maximum of 1 GbE throughput to each datastore
- NFS:
  - Maximum of 1 GbE throughput to each datastore

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 28

## IP Storage Multipathing Summary

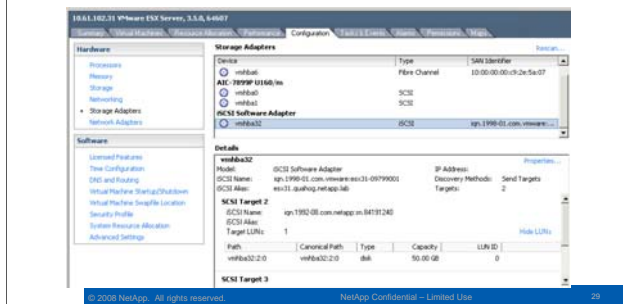
For, iSCSI software initiator and NFS, use Etherchannel trunks and vifs for link redundancy at the storage side. This is assuming the fastest available link is 1 gig to each controller with ESX 3.0.2 or to each datastore with ESX 3.5.

No matter how you configure this, at the end of the day, you only have one active path for each datastore. The maximum amount of throughput you can get to each individual datastore is still one connection—1 gig in this scenario. The aggregate throughput between the server and the storage can be more. But, an individual datastore is still limited to one link.



## ESX 3.5 Software iSCSI Multipathing

- In the past iSCSI hit a performance limit based on Data ONTAP displaying only a single IQN.
- Now one target per IP



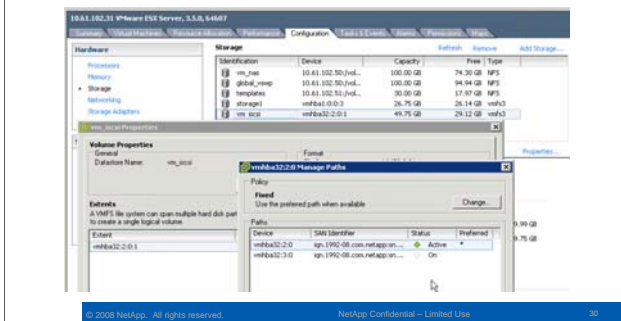
## Software iSCSI Multipathing

When you upgrade the ESX server to version 3.5, you get one target for every IP on the storage device, even though it still only has one IQN. You see multiple iSCSI targets, and when you manage paths, it looks like fiber channel multi-pathing. It works the way it appears to work.



## ESX 3.5 Software iSCSI Multipathing


- Software iSCSI Multipathing has multipathing based on multiple NetApp IPs



## Software iSCSI Multipathing

You use the NIC teaming on the ESX side to get multiple links and then on the target side you round-robin the active path. ESX 3.0 does IP Hash only, but 3.5 has added round-robin. But this round-robin capability is currently tagged as experimental. In VMware terms, experimental is any feature that has not been fully QA'd. For example, when iSCSI shipped in ESX 3.0, it was tagged as experimental because it had not been QA'd. Once iSCSI passed VMware QA, they removed the experimental restriction. They didn't patch anything, they just made iSCSI available for use and agreed to support it.






## Jumbo Frames and 10Gb Ethernet

- Certified for VM network
- Not certified for VMkernel / IP Storage
- NetApp support 10gb Ethernet as a NFS target today

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 31

## Jumbo Frames and 10Gb Ethernet

Support for jumbo frames (extended Ethernet frames that range in size from the standard 1,518 bytes up to 9,000 bytes) has been added in ESX 3.5, but this feature is certified by VMware for virtual machine networks only. It is not certified for NFS, but jumbo frames on NFS should eventually be supported after it goes through QA.




## Resources

- [ESX Server 3 Configuration Guide](#)
- [VMware Compatibility Guides](#)
- [NetApp Interoperability Matrix](#)

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Resources



## Summary

- VMware Virtual Infrastructure components
- VMware Virtual Infrastructure Suite features
- VMware Service Console, VMkernel, VMware File System (VMFS) and VMtools
- VMDK files
- VMware hardware- and software-based connectivity
- VMware storage access concepts

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 33

## Summary

Now that you have completed this module, you should have a basic understanding of:

- The essential components of a VMware virtual infrastructure such as datastores, virtual machines, and virtual disks;
- The most important features of the VMware virtual infrastructure suite such as DRS, VMotion, and VirtualCenter;
- The purpose of the VMware Service Console, VMkernel, VMware File System (VMFS) and VMtools;
- The significance of VMDK files.
- VMware hardware- and software-based connectivity; and,
- VMware storage access concepts such as hardware- and software-based multipathing, and VMware path management policies.



## Lab 1: Summary

- Create and configure a VMFS datastore over iSCSI.
- Connect to an existing volume and use it as an NFS datastore.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 34

## Lab 1: Summary

In this lab you create and configure a VMFS datastore over iSCSI, and also connect to an existing volume and use it as an NFS datastore..



This module covers when and how to do a VMware on NetApp proof of concept during a sales cycle.



## Learning Objectives

- Proof of concept
  - When to do a proof of concept (POC)
  - Common POC agendas
  - VMware IO benchmarking best practices
  - VMware IO benchmark results

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

2

## Learning Objectives

After completing this module, you should have a basic understanding of:

- When to do a proof of concept (POC)
- Common POC agendas
- VMware IO benchmarking best practices that you should follow in a POC, and
- VMware IO benchmark results



## VMware Proof of Concept (POC)

- Why do some customers want to see an IO test?
  - Servers were underutilized and virtualization solved this issue
  - Customers commonly state that FCP utilization is very low
    - 2-3% per port on 4Gb FCP is typical (8-12 MBs)
- Customers want validation that IP protocols are production ready

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

3

### VMware Proof of Concept (POC)

Why do some customers want to see an IO test? Servers were underutilized and virtualization solved this issue. Customers commonly share that FCP utilization is very low—2-3% per port on 4Gb FCP is typical (8-12 MBs). Customers want validation that IP protocols are production ready



## VMware Proof of Concept (POC)

- First rule of Proof of Concept... use reference accounts instead, whenever possible.
  - POCs tend to slow down the sales cycle.
  
- When a Proof of Concept is Necessary:
  - Engage your local VMware expert or NetApp counterpart to assist you;
  - Look for help before you begin to verify your test plans
  - Limit the scope of your testing to real world configurations

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

4

## VMware Proof of Concepts (POC)

The first rule of proof of concept is to avoid them if you can because they slow down the sales cycle. Use reference accounts instead, whenever possible. But, there are sometimes valid reasons to do a proof of concept, and you will have key customers who will require them.

The second rule of proofs of concept is to engage your local VMware expert or NetApp counterpart to assist you because a VMware sale can be complex. At the very least, get a VMware expert to take a look at your plan before you present it to your customer.

Limit the scope of your testing to real world configurations—IO runs with a read and write mix and a realistic block size.





## Common POC Agendas

- Typical storage and application resiliency
  
- VMware IO performance
  - FCP
  - NFS
  - iSCSI
  
- Backup and Restores of VMs

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

5

## Common POC Agendas

Performance benchmarking is going to generate the most work for you in a proof of concept.

Whenever you decide that a proof of concept is necessary, what should it include? The labs associated with the instructor-led version of this course demonstrate how to go through a technical proof of concept with a customer.

There are three types of demonstrations that NetApp customers might want to see in a virtualization proof of concept:

- The first is typical storage and application resiliency--both hardware and software fail over. We will talk briefly about multipathing in a virtualization context, but, otherwise, resiliency issues are not unique to a virtualization environment, so we are not going to talk much about storage and application resiliency in this module.
- By contrast, performance is one of the big issues in most virtualization proofs of concept. A lot of customers say, "You say NFS is great, and I believe you. I can see the management story that you are telling. I just don't know about the performance." Sometimes customers go into virtualization planning without a particular protocol in mind and they simply want to do a bake off. They say, "I want the best performance. I am not really concerned about the features of the different protocols. I want to talk about performance," especially when they want to compare the different protocols to each other or compare NetApp performance with that of another storage vendor's device.
- Backup and restore is a big driver for NetApp in this market. You will see a lot of proofs of concept around backup and recovery. Generally, when you go into a proof of concept, you should first show how the fundamental technology works. For example, you should do backup and recoveries of one virtual machine at a time from the command line on the ESX server so that you can show the customer very explicitly exactly what is happening, rather than running a script and saying, "Okay, now it is backed up." You should break it down, step by step, to a granular level and say, "This is exactly what is happening right now. This is what is happening in the virtual machine and this is what is happening on the storage device." It really helps expose what is different about NetApp's backups compared to other backups and how much more powerful NetApp's backups are. And you will be able to demo SnapManager.



## Environment Setup Is Critical to Success

- All test environments should reflect production in terms of design
  - This is commonly overlooked and leads to inaccurate results
  - VMware deployments consist of:
    - Many physical servers & virtual machines
    - All which access a large datastore
  
- IO tests require at least two ESX servers
  - IO constraints are not observed with single server tests

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

8

## Environment Setup Is Critical to Success

When you are benchmarking speeds and feeds in a virtualized environment, you are measuring IO over a number of physical servers and a greater number of virtual servers running on top of them, not just one object to another object, or just one database server to one storage device. This is very important distinction to keep in mind.

If you run IO tests in a virtualization proof of concept on individual virtual machines, with individual servers, the test results are not going to mean anything. In the real world, no customer runs a single virtual machine on an ESX server, or a single ESX server to the same datastore. If you don't test on production configurations—in terms of design—you are not going to get valid production numbers. This will be especially important when you are comparing FCP to NFS—whether over fiber channel, iSCSI or both. When you run multiple physical servers pointing at that same datastore, the short comings of the FCP protocol compared to NFS really come to light.

It is important that you stress to customers that you have to test with multiple servers and with multiple virtual machines—at least two physical servers with at least two virtual machines each. It simply is not going to be a valid test otherwise. You don't need ten, you don't need twenty, you need two. If you can get three or four, great. Any more than four is just wasting time. Each physical server is going to run two virtual machines. That's all you need. With multiple servers running multiple virtual machines hitting the storage, you'll be able to show the customer very interesting results.



## Storage Design

- Create multiple datastores:
  - One datastore to house operating system VMDKs
  - One datastore for each protocol to be tested
- VMDKs for IO testing reside on these datastores.
- This design allows VMs to use the same IO test on different disks which are connected via different storage protocols

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

7

## Storage Design

Here is the ideal design for doing IO performance testing in a virtualized environment.

Assume that you are going to test all three storage protocols side by side:

- VMFS over fiber channel,
- VMFS over iSCSI and
- NFS.

In this configuration you set up four datastores:

- Create one datastore to house all of the virtual machines—the root drives of all of the virtual machines being tested. You don't care how that is being stored, it could be a local disk.
- Create a separate datastore for each of the three protocols that you want to test—FCP, iSCSI, and NFS—with a virtual disk in each one. You will run all IO through these virtual disks mounted on the virtual machines housed on the first datastore.

This way you can easily compare the performance of the three protocols by sending the same IO load to all three virtual disks, each of which is running a different protocol over exactly the same virtual machine. You can either mount one virtual disk at a time and run the IO performance test through it, or you can mount them all concurrently as separate drive letters and test IO performance side by side.



## Ensure Proper VMDK Alignment

- Ensure a starting partition offset of 32768 (see [TR-3428](#))
- Guest OS VMDK does not require this setting

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

## Ensure Proper VMDK Alignment

One important thing to note: when you do a test or proof of concept, make sure that all virtual disks are properly **aligned** to starting partition offset of 32768 (see Technical Reference-3428). Disk misalignment is not a uniquely NetApp issue. Rather it is a virtualization issue. If the blocks don't line up, disks encapsulated as files running on another file system will generate extra IOs. We cover this disk misalignment issue and how to handle later in this module. At this point, it is enough to note that misalignment can have a huge negative impact on performance. Make sure that when you go into a proof of concept that all virtual disks are properly aligned.

Proper alignment is not as big an issue with **Guest OS VMDK** because of very low IO. If the customer has already built virtual machines that you are testing, you don't need to make them rebuild the virtual machines just because they may or may not be misaligned. You can just create new properly aligned virtual disks and use those to run the test.



## IO Generation Tool Settings

- IOMeter for Windows VMs
  - Supported by both VMware & NetApp
    - VMmark is not a good tool for stressing disk IO systems
    - 400000 byte working set allows data to be served out of FAS memory and to effectively measure the IO limit of the protocol
    - Real world block sizes of 4kb and 8kb
      - NTFS formats at 4kb by default, this is used for most applications
      - SQL & Oracle use 8kb by default

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

### IO Generation Tool Settings.

So what testing tools should you use? Customers will probably ask about VMmark—VMware’s testing suite—but it is not appropriate for several reasons. VMmark is a very large and complex testing suite. It takes a lot of servers, and it takes a lot of time to set up and administer. VMmark is designed for stressing the virtualization platform CPU and memory. It is not built for generating a lot of IO. It will be very rare that you will have a platform that can generate enough IO to stress your storage system by simply using the VMmark test suite. VMmark will generate a lot of CPU workload, lots of memory usage, and a lot of page writes, but it won’t generate a meaningful volume of IO to the disk, which is the whole point of going through this exercise. Instead of VMmark, you should use IO-dedicated tools such as IOMeter for IO benchmarking. IOMeter is very good for Windows virtual machines.



## IO Generation Tool Settings

- Test 1
  - 4kb block, 4kb offset, 60% random, 25% writes, 4 workers, 400000 byte working set
- Test 2
  - 8kb block, 8kb offset, 60% random, 25% writes, 4 workers, 400000 byte working set
  - Working set allows data to be served out of FAS memory and to effectively measure the IO limit of the protocol

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

10

### IO Generation Tool Settings.

One of the things to stress here is relatively small work loads—only four hundred thousand bytes for example. You will be testing one protocol versus another. You don't want the controller, the disks, the model of the storage system, or the number of spindles to become a bottleneck. You want to deal only with moving data between the ESX servers and the storage device. A small workload will ensure that you are not bound at the disk level. Obviously, if you want to stress the system and compare it to another storage device, your workload would need to be the same as used to stress the competing storage system.



## IO Generation Methodology

- Measure IO at the datastore level
  - This is where bottlenecks are found
- Use tools like sysstat and statit to gather IO information from the storage controller
  - [Performance Analysis Fundamentals, Release 7.2](#) class is an excellent resource for more information on these tools
- Test IO on a single protocol at a time
- Increase IO load by starting load on one VM at a time
- Be deliberate—take each step and measure before moving to the next

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

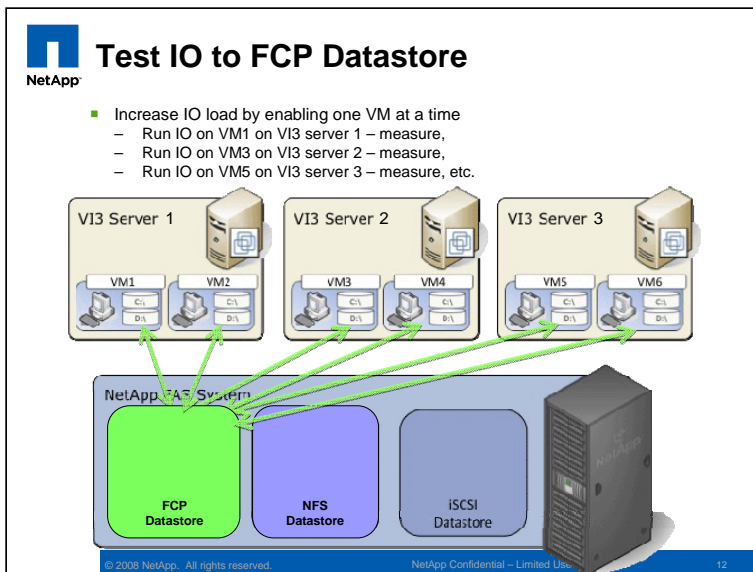
11

## IO Generation Methodology

Here is your IO performance testing methodology:

- Measure IO at the datastore level.
- Test IO on a single protocol at a time.
- Increase the IO load by enabling one virtual machine at a time. and
- Be deliberate—take each step and measure before moving to the next.

Performance Analysis Fundamentals, Release 7.2 class is an excellent resource for more information on IO performance measuring tools.



## Test IO to FCP Datastore

PHASE 1 is to test one protocol at a time. Pick a protocol, and test it out.

Phase 2 is to scale up. The simplest thing to do is just scale up by virtual machine. Kick off the IOMeter job—or whatever you are using to generate the load—one machine at a time. You are going to see dramatically different results from the virtual machine to when you start running multiple virtual machines to that same datastore. Generate IO to the same datastore from multiple sources.

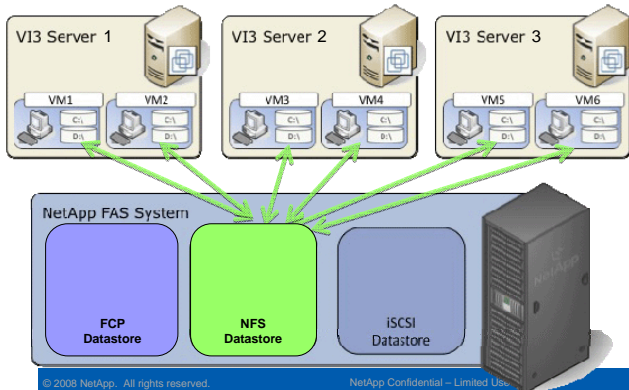
This slide shows a representation of three physical servers, each with two virtual machines. C drive is used by the host OS and D drive is for I/O. On the bottom is the NetApp system. You have three volumes, one for fibre, one for NFS, one for iSCSI. Then when you test, turn on the first virtual machine on the first server and measure throughput. Then start the first virtual machine from the second physical server and measure throughput both per virtual machine and aggregated. And repeat, starting the first virtual machine on the third server, document, and move on. And just repeat, adding the second virtual machine on each server one at a time. And then you stop and make a virtual disk on NFS and you repeat the same process.





## Test IO to NFS Datastore

- Increase IO load by enabling one VM at a time
  - Run IO on VM1 on VI3 server 1 – measure,
  - Run IO on VM3 on VI3 server 2 – measure,
  - Run IO on VM5 on VI3 server 3 – measure, etc.



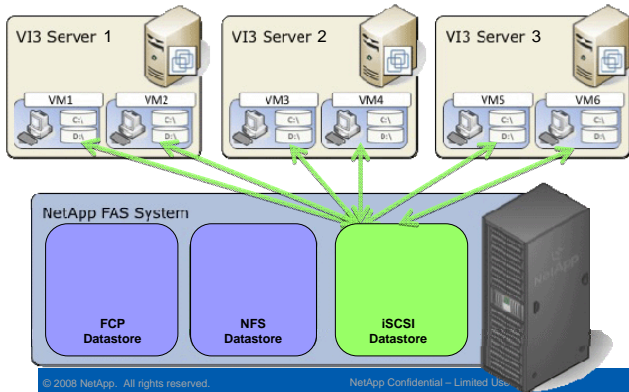
## Test IO NFS Datastore

Next, test IO to the next datastore from multiple sources,



## Test IO to iSCSI Datastore

- Increase IO load by enabling one VM at a time
  - Run IO on VM1 on VI3 server 1 – measure,
  - Run IO on VM3 on VI3 server 2 – measure,
  - Run IO on VM5 on VI3 server 3 – measure, etc.



## Test IO to iSCSI Datastore

And repeat the process again for the iSCSI datastore.



## What Results Will You See?

- The results will be much higher than production workloads
  - Falsely inflates load on resources
  
- When measuring IO at the datastore level...
  - All protocols will be within 10% of each other
    - These results are counterintuitive
  
- When measuring CPU...
  - FCP has little cost
  - iSCSI & NFS has moderate cost

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

15

## What Results Will You See?

Now let's discuss the results you should get from the IO benchmark.

It is good to establish a performance requirement baseline with the customer before you talk about comparing protocols to each other. Because, regardless of the results of your test, a single link, whether it be iSCSI or NFS, is theoretically capable of approximately 80-100 megabytes per second per ESX server. Most virtualization platforms are not going to need that kind of throughput. So it is a good idea to talk about what kind of performance numbers the customer needs before you run all of these tests. Generally, you are going to blow away anything that anyone needs. And, it is easier to set that stage first than have to back track when you realized that IO performance is much better than they need anyway, so there may be no need to run a test.

The results will be much higher than production workloads and could falsely inflate actual load on resources. Now what's always surprising is that all protocols will be within 10% of each other. These results are counterintuitive.

When measuring CPU performance, FCP has little impact, while iSCSI and NFS have moderate CPU performance impact.



## What About Storage Protocol?

“There is no iSCSI vs. FC performance argument any more in this space-- the data is in. IP protocols (iSCSI and NAS) seem to deliver equivalent performance to FC in the vast majority of VMware applications.”

- Chuck Hollis, Vice President of Technology Alliances at EMC (December 2007 Blog)

“Traditionally people have this belief that Fibre Channel has the best performance... with NFS & iSCSI you can get very high performance solutions...”

- Bing Tsai R&D Manager Storage Performance Group at VMware

[http://www.vi3demo.com/vmworld/Bing1\\_2.mov](http://www.vi3demo.com/vmworld/Bing1_2.mov)

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

16

It is now acknowledged by competitors and by VMware that FC, iSCSI and NAS typically deliver equivalent performance in most VMware environments.



## A Joint Test from VMware & NetApp

- Both companies are completing an IO scaling benchmark for customers to reference
  - Testing in final phases
    - Testing was delayed, but is back on track
- Test methodology was the same as discussed in this presentation

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

17

### A Joint Test from VMware & NetApp

NetApp and VMware are completing an IO scaling benchmark for customers to reference. The testing of this benchmark is in its final phases. The test methodology was the same as discussed in this presentation.



## Resources

- [Network Appliance And VMware Virtual Infrastructure 3 Storage Best Practices](#)
- [Performance Analysis Fundamentals, Release 7.2](#)

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use



## VMware POC Summary

- Offer reference accounts
- If you must do a POC, follow the test methodology
- Be prepared to explain the results – establishes VMware expert level status
- Remember – we sell all protocols

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

18

### VMware POC Summary

Wrapping up on proof of concepts: avoid them if you can because they can increase the sales cycle; but, if you are going to run one, make sure you have a test in advance. Get a VMware expert to look at your test.

And don't forget that we sell fiber channel and iSCSI as well as NFS. If the customer just bought a million dollar SAN, just sell them fiber channel to plug into their SAN. In fact, most of the VMware on NetApp deployments to date have been on SAN storage. The great thing about our solution is that we can extend our value proposition—backup and recovery, replication, deduplication, FlexClone—regardless of the protocol. This is why you don't want to go in talking about protocol. Rather, you should go in talking about features. Because you can tell a customer better data resiliency, better data protection, better replication, and so on, and you don't have to talk about protocol. You don't have to make any caveats when you are talking about these because all of these features, all of these value adds are available regardless of the protocol that you are using. It doesn't matter. You can talk about features first. Get them excited first and then get into protocol last.



This module covers NetApp backup, recovery and replication differentiators in the virtualization space.





## Learning Objectives

- VMware backup, recovery, and replication challenges
- VMware on NetApp value propositions
- The advantages and disadvantages of VMFS, NFS, and RDMS in terms of VMware backups
- Using scripts, SnapManger for Virtual Infrastructure, and OSSV to create backup copies
- Using scripts and SnapManger for Virtual Infrastructure for data recovery

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## Learning Objectives

When you have completed this module, you will have a basic understanding of:

- VMware backup, recovery, and replication challenges;
- VMware on NetApp value propositions;
- The advantages and disadvantages of VMfS, NFS, and RDMS in terms of VMware backup, recovery, and replication;
- Using scripts and SnapManager for Virtual Infrastructure to create Snapshot copies;
- Using scripts and SnapManager for Virtual Infrastructure for data restoration.



## VMware Backup and Recovery Challenges

- Shrinking backup windows and growing data sets
- Losing a server has much larger impact than before
- Virtualized servers makes this problem far worse
  - More data behind each server
  - VMware's shared model does not provide enough IO
  - Backup windows have not grown
- No direct Fibre Channel Tape support

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

## VMware Backup and Recovery Challenges

Perhaps NetApp's greatest value to customers is to offer superior solutions for backup, recovery, and replication of mission-critical data. This is especially true in the virtualization space. Almost every VMware customer has a problem backing up virtual storage and even more severe challenges with data restoration and disaster recovery.

VMware's storage model works extremely well for sharing virtual servers and virtual storage, especially when the virtual servers don't have high I/O requirements. But, customers find that after they virtualize, the time it takes to run a full backup is doubling because they often have more data to restore per server than before. Virtualization consolidation results in all the data volumes to be backed residing on a single server. Most customers have 10 or 15 times the data per physical server to backup and that would need to be restored in the event of a failure. They no longer have the bandwidth to backup the data in a reasonable timeframe, and have virtually no chance of meeting restore time objectives. all those lightly loaded systems get busy at backup time.

In addition, tape backup from a virtual environment is a tedious manual process, where the associated tapes are continually growing adding to cost and complexity. Recovery is error prone, and extremely slow. Tape based DR is even more complex. And the VMware environment provides no direct fiber channel tape support—no way to attach a tape backup device directly to an ESX server for backup or replication purposes.

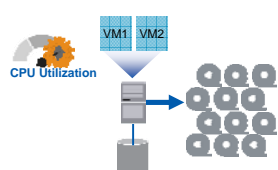
**NetApp** **Instantaneous Backup, Zero Server Impact**

**The Problem**

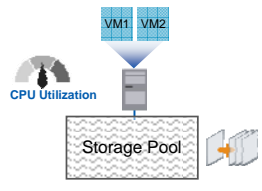
- High server utilization
- No spare cycles for backups
- Tape is slow, complex, & expensive
- Disaster recovery is very difficult

**NetApp Snapshot Solution**

- Servers run apps, not background processes
- Instantaneous backup and recovery
- Low storage overhead
- Application consistent



**Traditional Backup Is NOT Practical**



**Fast, Affordable, and Simple Backup and Restores**

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 4

## Instantaneous Backup, Zero Server Impact

Enter NetApp Snapshot with SnapRestore. NetApp has unique value in this area, with industry-leading capabilities that fit exactly in the VMware environment. The backups are run on NetApp storage, and many copies can be made at any time increment in a matter of seconds. These copies are not full copies of data. They are only tracking changes and are very efficient in terms of overall storage capacity.

Restores can be done instantaneously from any of the copies. Furthermore, the backups are application aware in that they have been coordinated with the application in a known state and cover the application binaries, the logs, and the application data, which in turn enable the restores to be that much faster as well as return to a known state.



## Ask A Few Questions When You Begin

- What is your backup software?
- What is your restore SLA?
- How long does it take to restore a VM? “X” hours?
- How many VMs do you store on a datastore? “Y”
- Does “X” times “Y” hours meet your SLA?

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## Ask A Few Questions When You Begin

Whether or not your customers are aware of these backup and replication challenges, just asking a few questions and talking about the solution can open up a gap that your competitors can't fill.

Start by talking with your customer about the inherent risks of the shared storage model, where a failure doesn't take down a server, it takes down lots of servers. You just have to ask simple questions of your customer, each of which can have significant impact on backup performance:

- What backup software is in place? For example, Tivoli Storage Manager (TSM) backs up in a different way than other backup applications. And you need to know if they are running VMware Consolidated Backup (VCB).
- What is the restore service level agreement (SLA) for all of the servers in this environment—i.e., how much time do you have to complete the restoration?
- How long does it currently take to restore each virtual machine and how many virtual machines are housed on a datastore? The least likely failure scenario is an individual virtual machine. If a datastore for an ESX server fails, multiple VMs will need to be restored.
- If it takes “X” hours to restore a single virtual machine, and you have “Y” virtual machines in a datastore, is it fair to say it will take X times Y hours to restore a datastore?
- Does that meet your SLA?

These questions are not new for NetApp SEs, but in a virtualized environment you can present them with a stronger sense of urgency.



## Define backups as...

- A full copy of the production data
- Previous versions of the data
- Stored on a second set of media
- Stored in a offsite / remote location
- This used to be defined as tape
  - Due to restore RTO times tape is better suited for archival purposes

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## Define backups as...

Another question that you want to ask your customer is how they define a backup? From a NetApp perspective, it is best to a define a backup as a complete copy of the data—a backup history on another set of media that's stored off site. It should be a full copy of the production data including multiple versions—not just the current version and the version before that--stored on a second set of media in another site. It sounds like tape but tape's play is now archive. If you want backup and restore, then it's disk, and particularly, disk that is ready to use. Backing up to a disk pool using TSM, or backing up to any vendor's box from which you can't quickly get production back up and running provides you nothing.



## Datastore Level Is The Single Point Of Failure In VMware

- VMware provides many layers of redundancy:
  - VMotion
  - DRS
  - HA
  - Redundant storage networking
  - Redundant storage switches
  - Redundant storage controllers
- But, there is only one copy of the data on disk somewhere.
- Calculate the RTO for a full datastore loss
  - Multiply restore time for a single VM by number of VMs per datastore

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

7

## Datastore Level Is The Single Point Of Failure In VMware

VMware provides many layers of redundancy, but the datastore remains the single point of failure. In an enterprise-level VMware deployment, with multiple ESX servers in a data center, a great deal of redundancy is built in—DRS, HA, VMotion, as well as redundant power, redundant network connections, and redundant fabric connections. An ESX storage controller also has redundant disks and redundant connections. But, there is still only one copy of the data on disk somewhere. If that one copy of the data on disk somewhere were to become lost or logically corrupted, it becomes your single point of failure in a VMware environment. You need to get the customer to focus on the scenario of losing their entire datastore. How long is it going to take to recover all of the virtual machines in the datastore from their current tape backup storage?



## Agent Based Backup


- Multiple agents: one or more per host
- Server resources: multiple virtual servers using the resources of a single physical server
- Operational dependencies with scheduling and setup
- Backup windows are a challenge to meet
- Backup rate significantly slower than a physical server

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

### Agent Based Backup

Until now, there have been two ways to do a tape backup from an ESX server. The first is to install a backup client inside each virtual machine. But since an ESX server can house many VMs, managing the backup schedules of dozens of servers is a nightmare. As stated earlier, VMware customers find that backup windows are an increasing challenge to meet. Using this configuration, the backup rate is significantly slower than a traditional tape configuration.



**VMware Consolidated Backup (VCB)**

- Backup proxy solution that leverages VMware's snapshot
  - Writes to virtual disk suspended during backup
  - Virtual disk frozen while log file writes are appended
- Two options:
  - Image based
    - Easy VM recovery
  - File based (windows only)
    - Easy file level recovery
- FCP and iSCSI support
  - NFS support planned for future version

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## VMware Consolidated Backup (VCB)

VMware Consolidated Backup (VCB) is the other solution commonly considered to provide tape backup. However, it is not an enterprise-level solution. It is a Windows proxy server that can attach to a datastore and dump its contents to tape. It is used along with backup software such as NetBackup to manage getting the data off of the datastore and onto tape.

VCB takes a VMware snapshot, dumps it to tape, and then deletes the snapshot. The problem is that the virtual disk is frozen until the VMware snapshot is deleted. All writes are going to a log file that will have to be written to the virtual disk as soon as the snapshot is deleted. If you have an 8 hour backup window, the system could be frozen for a significant amount of time as it replays that snapshot log when the tape backup is done.

The other problem with VCB is that it is either a file-level backup solution or a virtual-machine-level solution. You choose when you take the backup. If you want to have both virtual-machine-level recovery and file-level recovery, you have to run your backup job twice.

This solution supports fiber channel and iSCSI today, but not NFS. NFS support is planned for a future release of ESX.

If you are using VMware on NFS, you can use the VCB server as an enabler with UFS Explorer or a similar tool for file level recovery. You can use UFS Explorer with CIFS, for example, to browse the file system. The problem is, if you use VMFS, and the virtual disks are in that VMFS file system, you can't see inside the datastore. VCB, which has a VMFS driver for Windows, could be used as an enabler for that. You could use VCB server to mount the VMFS datastore, and then use UFS explorer.





## Why NetApp for Backup & Recovery?

- NetApp Snapshot technology can perform near instantaneous backups and restores of Virtual Machines
- OSSV can solve the I/O problem for environments requiring a more traditional file based backup solution
  - Only client side agent which reduces amount of data sent over network
  - Block level incremental backups forever
  - NetApp support of OSSV inside of VMs

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

10

## Why NetApp for Backup & Recovery?

NetApp can help address these backup challenges that the customers face in two ways. Snapshot technology either through scripts or using SnapManager, can enable customers to instantaneously take backups of all their virtual machines and provide for recovery of an individual virtual machine, literally within minutes or seconds depending on what storage topology they're using. For the customers that want to keep a more traditional file-based backup, Open System SnapVault provides a solution that's unmatched in the marketplace, in terms of actually being able to deploy a client into each virtual machine and completing a file-based backup that after the first backup is done is incrementals forever. OSSV and/or Snapshot can actually reduce the amount of data coming off of each virtual machine allowing you to back them up in the traditional fashion. Customers on VMware solutions can consolidate more physical servers onto an ESX Server when using OSSV because they can effectively handle the backup for twice as many servers and therefore double their utilization rate.



## Datastore-Level Snapshot Copies

- Policy driven immediate backups of VMs
- Eliminate client costs and scheduling issues
- Provide VM level restores within minutes
  - Unmatched RTOs when compared to tape
- Scripts today, SnapManager in the future

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

### Datastore Level Snapshot Copies

Data ONTAP makes Snapshot copies at the datatstore level—whether you are using scripts, or using SnapManager for Virtual Infrastructure. Backup policies and backup scheduling is all done at the datatstore level. You backup all of the virtual machines in that datatstore when you take a Snapshot. And you schedule backups at the datatstore level—not at the virtual machine level. Therefore, best practice is to group virtual machines in datatstores based on backup policy.

Even though Snapshot copies are at the datatstore level, you still have the ability to do virtual-machine-level and file-level restores locally with Snapshot and FlexClones. In a traditional tape backup model, the customer would have to pull from tape to restore a virtual machine, or even just a file. With NetApp Snapshot copies, virtual-machine-level recovery, and file-level recovery is done using local Snapshot copies. You need to access an offsite replica copy or tape archive copy only in the event that you lose the storage controller or lose the entire volume.

All NetApp backup solutions are available today. It can all be done using customizable scripts, and with VIBE. SnapManager for Virtual Infrastructure will be able to manage NetApp backup solutions in the near future.



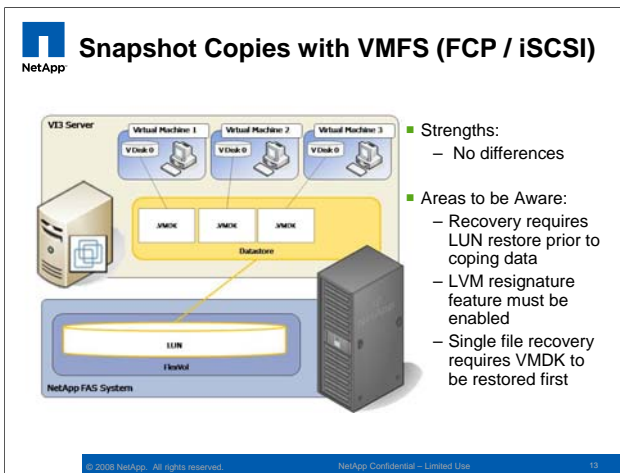
## Storage Savings Over Competitive Solutions

- Storage savings need to include production and backup media
  - Disk and/or tape
- NetApp data deduplication reduces primary storage by 50%
- NetApp data deduplication reduces DR storage by 50%

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

### Storage Savings Over Competitive Solutions

We talk a lot about storage savings in other modules, but, it is important to point out here as well that FAS deduplication can achieve 50% storage savings on both the primary and secondary sites. NetApp storage systems can provide two copies of your data--the production copy and a mirrored DR copy—all for the cost of your original data requirements.

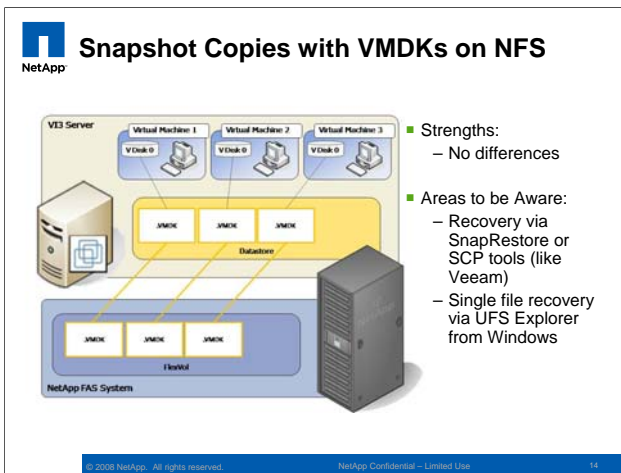


## Snapshot Copies with VMFS (FCP / iSCSI)

Let's look at Snapshot backups as we break them down for different protocols.

We have seen this slide before. It is the traditional VMFS datastore—a LUN that is formatted with VMFS—containing three virtual machines and connected via fiber channel or iSCSI. When you take a Snapshot backup in this configuration, it is obviously at the volume level. Recovering the entire datastore is easy. You revert the LUN to a previous Snapshot.

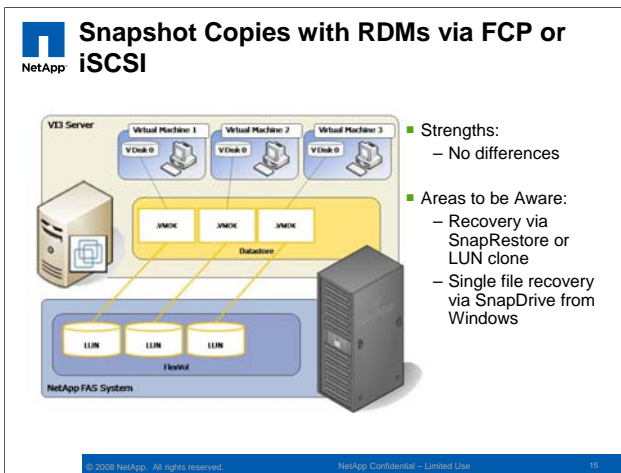
Now, if you just want to recover an individual virtual machine, or an individual file, you shouldn't revert to a Snapshot of the entire LUN. Instead, you mount a Snapshot of the LUN and then copy the VMDK files for the VM you need to recover. If a LUN contains large virtual machines, this process for recovering an individual virtual machine can be a lengthy process. For example, if you have large data base servers that are virtualized, and they have big data drives encapsulated in a virtual disk, recovery will involve a large file-level copy operation so recovery time is something to think about when designing this architecture. There is nothing you can do to make it run any faster. If it is a 100 gig file, and it becomes corrupted, you have to wait while that 100 gig file copies back. Single file recovery requires virtual disks to be restored first.



## Snapshot Copies with VMDKs on NFS

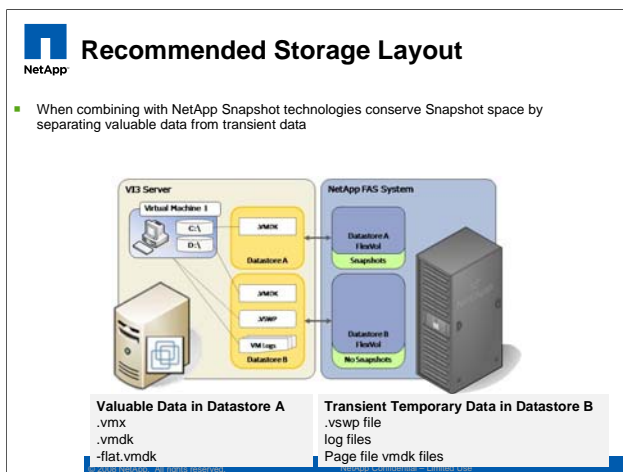
When you run VMware storage on NFS, Snapshot copies are still taken the same way. You still take a Snapshot of the entire volume, so you are still backing up at the datatstore level. But now in the event that you have a problem, not only can you revert to a previous Snapshot of the datatstore, you also have the ability to restore individual virtual machines using SnapManager without the need to copy files or by doing something as simple as a single file SnapRestore of a virtual disk file.

In addition, because the virtual disks are not abstracted within a LUN, you can use third party tools via CIFS or NFS to extract individual files. For example, UFX Explorer is a tool that can read a VMDK file and pull individual files out of it in much the same way that Single File Mailbox Recovery works in Exchange.



## Snapshot Copies with RDMs via FCP or iSCSI

As with VMFS and NFS, a Snapshot from a RDM-based datastore backs up all of the virtual machines in that datatstore. As before, you can do a datatstore-level recovery. And, as with NFS, you also have the capability to recover individual virtual machines. In this case, because this is a LUN and not a file, you can recover a couple of different ways: you can do a single file SnapRestore on the LUN or you can do a LUN clone, mount it, and run from there. Single-file recovery is still available in this environment. You simply mount a clone of the LUN with SnapDrive from another server and drill into that LUN and recover data.



## Recommended Storage Layout

There is a lot of transient data inside a virtual machine—a page file, and temporary directories, for instance. It is a best practice, documented in the Technical References, to separate non-transient “real” data from this temporary transient data into separate datastores, especially if you will be doing any kind of replication or long term archiving of the virtual machine.

In this slide, the virtual machine has two virtual disk drives. C drive contains the operating system and applications. Transient data such as page files and the Vswap and temporary VM logs files that VMware uses are stored on the second virtual disk drive called D drives. Vswap, is a VMware page file that gets created with every virtual machine. It is equal to the size of the memory for that virtual machine. VM log files are associated with VMware Snapshot copies.

Because of the transient nature of the data on drive D, you don’t need to worry about restoring it in the event of a failure. To do a backup, you only need to take a Snapshot of the datatstore that contains the non-transient data, Drive C in this example. This Snapshot can then be replicated, and archived to tape. You don’t need to take a Snapshot of the datatstore that contains drive D, the transient datastore. For ESX servers with multiple virtual machines, you can create two different virtual disks for each VM—one on each datatstore. You then take Snapshot copies only of the datastore that contains non-transient data.

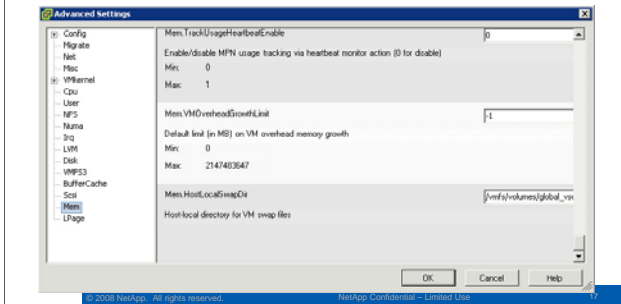
The location of these virtual disk files and the .vswp file is in the .vmx file. If you lose a virtual machine, you will have to rebuild all of this. If you are recovering and you have the .vmx file you just have to have the datatstores and the virtual disks in place. With Windows virtual machines, if you don’t have the drive D available when you start, Windows will automatically move the page file back to C and you would then have to go through a clean up process. If you are in a disaster recovery and you want to get those virtual machines on line quickly, you can just start the VM and let Windows move the page files automatically and then clean it up later.

One of the problems that we have with this is that Windows pages constantly. A VMware recommended best practice for virtual desktops is to make as small a page file as possible and over provision the memory to take advantage of the memory sharing and to try to affect the behavior of the virtual machine.



## Set Global VM Swapfile Location

- A VM swapfile (.vswp) is equal to the size of memory assigned to a VM
- Set a swapfile directory per server in a datastore not backed up or replicated to save storage space in NetApp Snapshot copies



## Set Global VM Swap Files Location

In ESX 3.5 server you can set the global VM swap file directory location for the entire ESX server by setting the value of HostLocalSwapDir. This is in advanced settings on the ESX server configuration. It is the very last option, HostLocalSwapDir, on the memory (Mem) page of the Advanced Settings window.

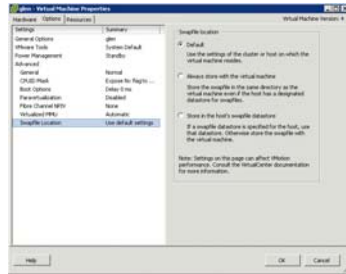
You normally set the Swapfile location on a per ESX server basis. In a VMotion event, the swapfile is recreated by the parameters assigned in that ESX server. It could jump around if you don't have all of your ESX servers set the same. It is important to make sure that this parameter is set the same on all ESX servers.





## Modify VM Swapfile Location

- Swapfile location can now be modified per VM in VI Client
- Used to require VMX editing



## Modify VM Swap Location Per VM

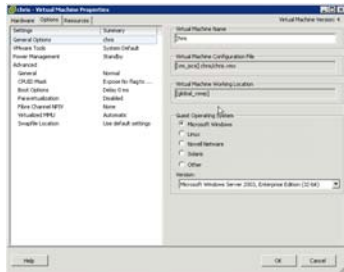
You can also set the Swapfile location at a virtual machine level on the Options tab of the Virtual VI Client inventory Machine Properties window, shown in this screenshot. Your choices are to:

- Accept the default settings for the cluster or host ESX server on which the virtual machine resides; or
- Always store it with the virtual machine, which means it will always be put in the same datastore and directory as the VMX file; or
- Store the swapfile in the host's default swapfile datastore, if one exists, or otherwise, store it with the virtual machine.



## VM Working Location (workingDir)

- VMware snapshot files are stored in the Working Dir (Working Location)
- Still set in the VMX file
  - `workingDir = "/vmfs/volumes/global_vswp"`



## VM Working Location (workingDir)

The Virtual Machine Working Location, also called the `workingDir`, is where VMware stores snapshot files and log files that it creates when it takes a VMware snapshot. The VI Client inventory Virtual Machine Properties window, in ESX 3.5, shown in this screenshot, now gives you the capability to check the `workingDir` setting, but you still change its value by editing the `workingDir` line in the VMX file.

**NetApp Backup Storage Policies**

- Assign a backup policy per datastore
  - One datastore per Flexvol
  - Policies includes snapshots, retention, and SnapMirror

Aggregate w/ RAID-DP

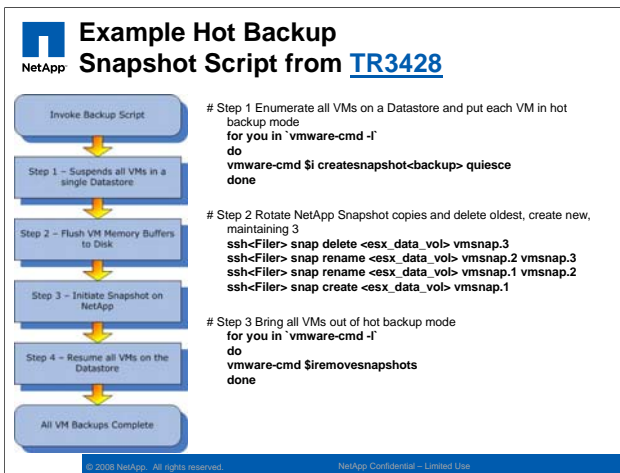
© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Backup Storage Policies

What if you have many datastores? How should you manage backup scheduling? Best practice is to assign backup policies at the datastore level. With Snapshot, the Vibe tool, AutoSnap, and SnapManager replication scheduling and intervals are all done at the datastore level. When you create a virtual machine, you can set the backup policy by assigning VM storage to the appropriate datastore.

In the example shown in this slide, you have two tiers of virtual machines: virtual machines that are backed up every 8 hours, and virtual machines that are backed up every 24 hours. When you create a virtual machine, you decide whether to assign its non-transient storage to a tier-one datastore to be backed up 3 times per day, or to a tier-two datastore to be backed up once per day. And then the other data, temporary and transient data—page files and swap files—will all get assigned to a dedicated datastore, which in this case isn't backed up at all.

Any time that you want a new backup policy, you can create a different datastore and configure the backup and replication policy for that datastore individually. Following this plan, you won't have to modify the backup infrastructure because you added a new virtual machine. You will have customers who decide to create an individual Snapshot of each virtual machine and they will have to change their script to accommodate a new VM, but, our model is to move away from that and deal with everything at the datastore level.



## Example Hot Backup Snapshot Script from TR3428

What happens when we take a backup? This slide shows the backup process distilled down to its simplest mechanism. These are commands that are running in the service console on the ESX server. The tools that we have, like Vibe and SnapManager, use ESX APIs to communicate with the service instead of running these service console commands. But the underlying process is exactly the same.

The first thing that happens is to discover all of the virtual machines in a particular datastore. Next, each virtual machine should be suspended—that is, put in a hot backup mode, which is a lot like the Oracle hot backup mode. Suspend freezes the virtual disk file and starts a log file to which all writes will be appended. This is why there are so many performance problems with the VMware snapshot copies. Once you take your first snapshot, the VMDK file is frozen. No writes will go into that VMDK file until all of the snapshot copies have been deleted. If you have a 5 day old VMware snapshot, the VMDK file has not changed in five days. All of the writes for the last 5 days are sitting in a log file. When you finally delete that snapshot, all of those writes that are sitting in the log file have to be appended to the virtual disk. We have customers that saw they had an old VMware snapshot sitting on their virtual machine and decided to delete it. Their ESX server was frozen for 10 hours while VMware appended all of that data.

Next, after suspending writes to the VMDK file, create a VMware snapshot for each virtual machine in the datastore. If you are looking at this process in Virtual Center, you see a snapshot called backup appear on each virtual machines on the script.

Before each snapshot is taken, the virtual disk is made consistent by VM tools. This is just file system consistency, not application consistency. The virtual machine flushes all of its writes to disk and then takes a snapshot.

Nothing is being written to the virtual disk files at this time. VMware has started up a log file for each virtual disk and is appending writes to it. All of the virtual machines' disks are in a frozen state. Next, Data ONTAP takes a NetApp Snapshot on the volume that contains the datastore. This is just a little rotating Snapshot script.

After you have taken the NetApp Snapshot and have captured the intact virtual disk, reenumerate all of the virtual machines and delete the VMware snapshot on each virtual machine. When you delete the VMware snapshot, each log file gets written into the respective virtual disk. There could be a small performance impact for taking these Snapshot copies if you are under heavy load because the writes will have to get written twice when you delete those Snapshot copies. Ideally you still want to take this backup in a period of low activity. But, in most cases, these Snapshot copies are only going to exist for a minute or so.

This model does not cover application consistency, period. If you want the application to be consistent you have to get the application into a backup mode where its files are frozen as well, and then do the backup. If you are running Oracle in a VM, it is really easy to freeze Oracle. You can write a script that first freezes Oracle and then freezes the virtual machine, makes NetApp Snapshot copies of everything, and then releases all of the freezes. But, with an application that doesn't provide that kind of control, such as Exchange where you can't manipulate it outside of the APIs, you wouldn't be able to write a script to create a consistent state in the application.


Keep in mind that these backups do not protect the running state of the virtual machine. They are protecting the contents of the virtual disks, making sure that the virtual disks are recoverable in a disaster. If you have to restore this Snapshot or fail over in a disaster, you are starting these virtual machines and they will boot from scratch.

What you are doing with this Snapshot is: (1) making sure that all of the file systems are intact; (2), making sure that all of the file systems are marked consistent. When you go to start these virtual machines, you don't have to do a check disk or anything like that. Theoretically, you could take a NetApp Snapshot without first taking a VMware snapshot to freeze the file system. But, if you have 1000 virtual machines and you go to fail over and all your virtual machines run check disk simultaneously to the same storage device, that's going to have a big impact if you need to restore all of these simultaneously.



## Data Layouts Impact Restore Options

- Objects residing on flexible volume can be recovered via SnapRestore
- Compare the impact on a VM when a SnapRestore is run



© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

22

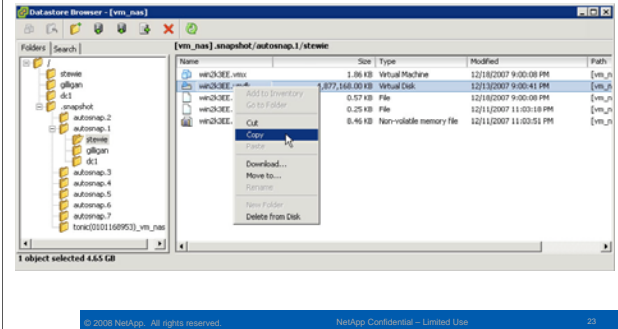
## Data Layouts Impact Restore Options

This slide summarizes the data layout differences between VMFS, RDM, and NFS and what happens when you need to do a virtual-machine-level recovery. In each case, backup is always at the volume level. You are always capturing the entire datatstore. With VMFS you can only recover the entire datatstore. With RDMs and NFS, you can do virtual-machine-level recovery directly from the storage device, but to do virtual machine level recovery, or file level recovery, you need to use a third party tool.



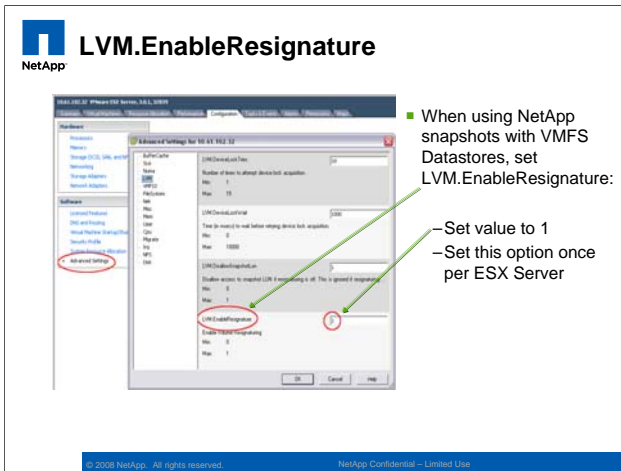
## Restoring a VMDK with VirtualCenter

- Can be used to manually restore a VMDK from a Snapshot



## Restoring a VMDK with VirtualCenter

VMware ESX Server 3.5 went GA in December, 2007. VirtualCenter 2.5 was released at the same time. This screenshot shows VirtualCenter's Datastore Browser. It is a browser interface that enables you to easily copy and paste files, and to manipulate virtual machines. You can use the Datastore Browser to restore a VMDK from a snapshot, for example. This screenshot shows an NFS datastore, but it works equally well for VMFS datastores.



## LVM.EnableResignature

If you mount a Snapshot of a VMFS file system as a LUN clone or as a FlexClone, the production datatstore and its clone have the same volume serial number. In ESX 2.5, this would fail. You had to use another ESX server in order to mount the clone. VMware fixed this issue in ESX 3.0. To mount a datatstore and a clone of the same datatstore on the same ESX server, go to advanced settings under the configuration tab. Select LVM for logical volume manager. Set the value of LVM.EnableResignature to 1. In the event that you mount a clone, the ESX server will automatically change the volume serial number of the clone copy. You set this once per ESX server, and it is permanent. By default this setting is off.



## NetApp Enhanced Data Protection

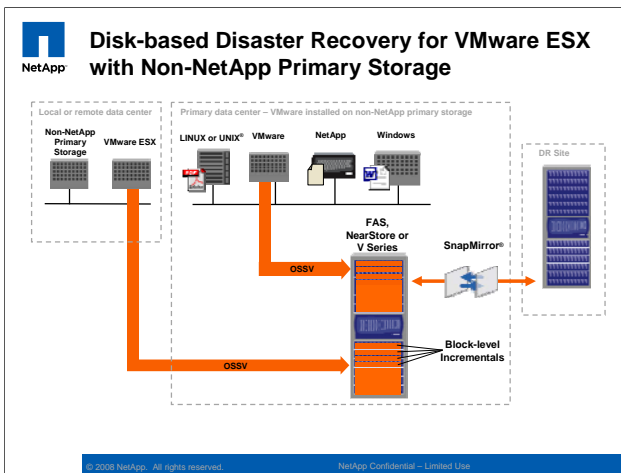
- Improved backup and recovery through enhancements to OSSV in version 2.6
- Improved management through SnapManager for Virtual Infrastructure

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

### NetApp Enhanced Data Protection

In addition to script-based Snapshot copies, and client-based SnapVault, NetApp has added enhanced data protection for VMware infrastructure. OSSV 2.6 has added support for running directly on an ESX server. And SnapManager for Virtual Infrastructure offers a graphical user interface for backing datastores and for restoring datastores and individual virtual machines.





## Disk-based Disaster Recovery for VMware ESX with Non-NetApp Primary Storage

This slide depicts the NetApp disk-based disaster recovery architecture for VMware ESX servers with non-NetApp primary storage. This solution is using OSSV to backup to FAS, Nearstore, or V-Series. When you compare NetApp's deduplicated disk-based solutions to tape for long term archiving, you can achieve 50% space savings. If you are doing weekly full backups to tape, you are storing 100% of your storage content every week on tape. TSM has the incremental forever model that doesn't have to run full backups forever, but, the disk pools for TSM consume at least as much space as a full backup. In a big environment, that could be a huge disk pool. By contrast, NetApp's deduplicated backup copy contains 100% of the data, but consumes 50% of the space.



## Open Systems SnapVault

- Simplifies centrally managed serverless disk-based backup and restore of VMware datastores residing on non-NetApp primary storage
- Block-level “incrementals forever” technology for optimal performance
  - Only sends 4KB blocks that have changed
  - Ideal for limited bandwidth networks
- Cost effective
  - Creates a highly efficient, low impact method for protecting file and application data stores residing on non-NetApp primary storage
  - Limited hardware requirements—no proxy server requirement
- Simple backup and restore management with Protection Manager

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

## Open Systems SnapVault

Open Systems SnapVault is a centrally managed serverless backup engine for Non-NetApp primary storage. By supporting block-level incremental disk-based backups forever, disk utilization and performance is optimal. It offers optional integration with Protection Manager for simplified management.

Open Systems SnapVault 2.2 and later clients can be installed into individual VMware ESX virtual machines.



## OSSV ESX Support


- OSSV 2.6 extends SnapVault's core benefits through integration with VMware ESX
  - VMotion aware
  - Lightweight agent installs in ESX Service Console
  - Single agent per ESX server
  - Licensed per VMware ESX instance
  - Recovery granularity at ESX server datastore or VM level
- Optional integration with Protection Manager for simplified management and discovery specifically for virtualized infrastructure
- See the [demo](#) from VMworld for a look at OSSV 2.6

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## OSSV ESX Support

OSSV 2.6 extends SnapVault's core benefits through integration with VMware ESX, running as a single agent per ESX server, rather than as a client on each VM. It offers optional integration with Protection Manager for simplified management and discovery specifically for virtualized infrastructure. It is VMotion aware, and is licensed per VMware ESX instance. Due to architectural differences Open Systems SnapVault 2.6 cannot support VMware ESX 2.x or VMware GSX versions.

The OSSV 2.6 demo that was shown at VMworld is attached to this module for you to download, if you are interested. It is a 17.8MB in size.



## SnapManager for VI & SnapDrive

The SnapManager for Virtual Infrastructure product addresses the backup problems described in earlier slides by enabling VMware administrators to use the NetApp ONTAP storage-assisted backup capabilities (Snapshot, SnapRestore, SnapMirror).




## SnapManager for VI - Value Proposition

- Robust Business continuity
  - Automated snapshot management of VMs in VI
- Flexible IT infrastructure
  - Rapid Provisioning of VM's in VI
  - Policy based management
- Monitoring, Reporting and Management of VMs
- Allows IT management to maximize their investment
  - VMware's strength is in optimizing server resources
  - NetApp's strength is in optimizing storage resources with features like Snapshot, SnapMirror, SnapVault, Flexclone, and thin provisioning

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

### SnapManager for VI - Value Proposition

Backup management is the number one goal for SnapManager for VI. This first version of SMVI provides a policy-based integrated backup and recovery management tool for virtual machines as well as a cost-effective DR solution. It does all this through an easy-to-use interface (GUI and CLI) which is co-resident with VMware Virtual Center. SMVI allows IT management to maximize their investment in virtualization, and removes the dependencies on the storage administrator to perform data protection and recovery of virtual machines. The value proposition is robust business continuity, automated Snapshot management of Virtual Machines within the Virtual Infrastructure, and flexible IT infrastructure. You should look at SnapManager for VI as a server admin tool specifically for VMware environments.



## SnapManager for VI 1.0 Features

- Datastore-level Snapshot backup
- VM-level restore
- SnapMirror integration
- CLI or GUI options
- FCP, iSCSI, or NFS
- Runs on Windows
- VMotion aware
- Future releases:
  - Storage Provisioning & Cloning
  - Virtual Center required
  - Policy based management

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

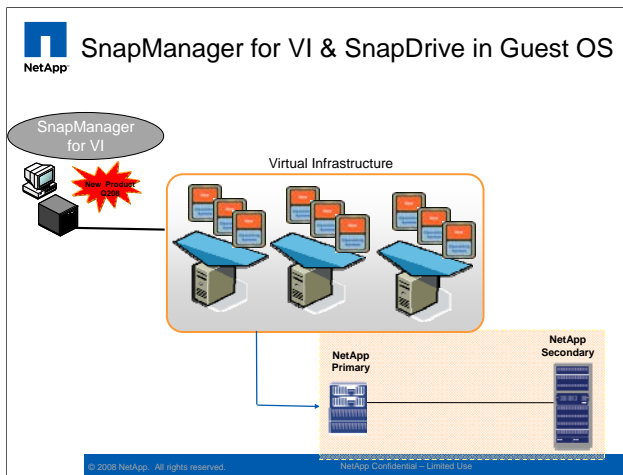
## SnapManager for VI 1.0 Features

SnapManager for Virtual Infrastructure is a management tool that sits on the ESX server, outside of the Virtual Infrastructure. This slide lists the high-level features found in version 1.0 of SnapManager for VI. It delivers Snapshot datastore-level backup management and individual VM-level restore management, SnapMirror integration, and both command line and its own graphical user interface.

In terms of protocols, it supports FCP, iSCSI, and NFS and has the standard Windows support. Future releases will include support for storage provisioning, cloning, VirtualCenter integration, policy-based management, as well as integration with all of NetApp's Storage Management Application Integration products such as Provisioning Manager, and Protection Manager.

SnapMirror support is a check box that asks whether you want to update the associated SnapMirror replica when the datastore is backed up. This provides consistent data for Site Recovery Manager. SRM support is not in and of itself a NetApp differentiator, because there are 14 other storage vendors who are writing hardware modules for SRM. Using Vibe or using SnapManager, we can guarantee the state of the data on the other side. SRM support is definitely very important and it is very good that we have it. Keep in mind that data integrity in this replication and fail over environment is much more important and is the NetApp differentiator.

SnapManager for Virtual Infrastructure doesn't require SnapDrive to operate because all of that functionality already exists within ESX and Virtual Center. Rather it queries the Virtual Center and initiates all of the VMware snapshot commands—the same snapshot commands that we use to take the backups today—and it will initiate all of those backup commands through ESX. NetApp didn't have to write a file system layer or build any mechanisms for freezing the file systems.



## SnapManager for VI & SnapDrive in Guest OS

This slide shows SnapManager for VI running outside the Virtual Infrastructure and SnapDrive running in the guest OS. In addition to the OSSV demo, there is a SnapManager for Virtual Infrastructure demo attached to this training module for optional download .



## Competitive Overview

- No directly competing product, other than scripts, that takes advantage of NetApp technologies
  - Snapshot, SnapMirror, SnapVault, Cloning and more.
- Potential Competitive Threats
  - VMware VCB
  - HP - StorageWorks & Opsware
  - IBM - Virtualization Manager
  - EMC – RepliStor
  - vizioncore - vRanger Pro
  - PlateSpin - PowerConvert
  - Backup vendors: Symantec; CommVault

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use

## Competitive Overview

NetApp is not the only vendor working to integrate its storage software with VMware virtual environments. vizioncore with its vRanger Pro and vReplicator backup and replication software has emerged as one of the leading third-party VMware tool providers. Replication vendors like Double-Take Software have VMware-centric versions of their products, and all the major backup software vendors -- Symantec, CommVault, and EMC, to name a few--have announced integration with VMware Consolidated Backup (VCB), VMware's tool for offloading the backup of virtual machines from the ESX host. There are some products out there that have the potential to pose competitive threats including HP StorageWorks and Opsware, and IBM Virtualization Manager, but no competing product, however, can offer the value adds included in NetApp's suite of offerings--Snapshot, SnapMirror, SnapVault, and FlexClone.





## Competitive Overview

- vizioncore – vRanger Pro
  - Imagelevel backup
  - Integrated with Virtual Center
  - Policy based and VMotion aware
  - Compress and send VM's across WANs
- PlateSpin – PowerConvert
  - Backup and Recovery
  - Single console across MSFT & VMware
  - One-click Failover
  - Auto discovery of VMs
  - Task-Based Wizards and Drag-and-drop Interface

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

## Competitive Overview

vizioncore and PlateSpin are probably the best known of these competitors. Even though they have lot of good features already, and they are ahead of the game with respect to SnapManager for VI, the key differentiator is our NetApp value adds. PlateSpin PowerConvert is a similar kind of product for backup and recovery, with a single console for Microsoft, and VMware, one-click failover, auto discovery of VMs.



## Resources

- [SnapManager for Virtual Infrastructure demo](#) from VMworld 2007
- [OSSV 2.6 demo](#) from VMworld 2007

© 2008 NetApp. All rights reserved.

NetApp Confidential – Limited Use



## Lab Summary

- Lab 2: Make NetApp Snapshot copies of VMFS and NFS datastores.
- Lab 3: Recover VMs from Snapshot copies created in Lab 3.
- Lab 4: Recover a single file created in Lab 3 by connecting to a Snapshot and copying the desired file.

© 2008 NetApp. All rights reserved. NetApp Confidential – Limited Use

In Lab 3 you make NetApp Snapshot copies of VMFS and NFS datastores. In Lab 4 you recover VMs from Snapshot copies created in Lab 3. And in Lab 5 you recover a single file created in Lab 3 by connecting to a Snapshot and copying the desired file create back into the existing Virtual Machine.





This module covers virtual disk alignment best practices that ensure optimum performance of VMware on NetApp storage systems. This is an issue that really should not come up in the sales cycle, but that you should be aware of, especially if you need to do a proof of concept in order to land an account.



## VMDK Alignment Issue

- For optimum I/O performance, the file systems of the VMDK, the datastore, and the storage array must be in proper alignment.
- Default installation values typically result in partition misalignment, and performance degradation.
- The best time to remedy this issue is at initial configuration.

© 2008 NetApp. All rights reserved.


NetApp Confidential - Limited Use


## VMDK Alignment Issue

For optimum I/O performance, the file systems of the VMDK, the datastore, and the storage array must be in proper alignment. Unfortunately, default installation values typically result in misalignment, and IO performance degradation.

**NetApp** **VMDK Alignment Issue**

- By default operating systems reserve space at the beginning of a disk in order to store their boot information—the partition offset.
- To build a VMDK:
  1. Start with a FlexVol.
  2. Create a vmdk file (on either a VMFS or NFS Datastore).
  3. Create a Partition and a Filesystem (within the vmdk).

**NTFS**      Windows.vmdk → 

**FlexVol** 

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMDK Alignment Issue

Anytime you build a server, the installation routine doesn't write its operating system starting at the first block of a disk. It always backs up to a pre-defined amount of space so that it can store information that is needed for storing user data, boot sectors, master boot records, and so on. That isn't a problem on a single machine such as a laptop or a desktop with one drive, but when you create virtual drives on storage arrays from NetApp, EMC, or whomever, the default settings where an operating system starts writing data typically does not align optimally for the file system or the block size of the physical storage object. That's why you indicate the type of LUN when you deploy traditional SAN storage on NetApp, to make sure that when you put Windows on the storage system, it will be aligned efficiently with WAFL.

**NetApp** **Default Partition Offset**

- Both NTFS and WAFL ultimately store blocks.
- The default partition tends to start at an offset of 63 sectors—32256.
- This can cause performance issues!


The diagram illustrates the default partition offset for NTFS and WAFL. It shows two horizontal bars representing data blocks. The top bar, labeled 'NTFS', is offset to the right relative to the bottom bar, labeled 'WAFL'. Vertical lines connect the boundaries of the NTFS blocks to the WAFL blocks, showing that the NTFS blocks do not align with the WAFL blocks. This misalignment is the result of the default partition starting at an offset of 63 sectors (32256).

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Default Partition Offset

The goal is for a 4K logical call to pull only 4K data off the disk. In a virtualized environment, storage is in flat files served by virtual entities. A storage vendor cannot control where data writing will start on an object. The default settings will almost always write data on a virtual disk such that it does not align optimally with blocks in the underlying storage layer that is provided by the storage array vendor. Misalignment can potentially result in the storage system reading two blocks—8K—for every block requested from the virtual disk.

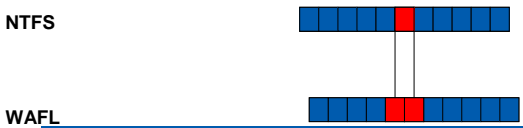


 **Performance Impact Example**

- Read 1 Block within the Virtual Machine
- Requires 2 IO operations

NTFS


WAFL



© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

### Performance Impact Example

Block boundaries are going to be off. Potentially, every I/O call from the virtual machine level could generate two on the storage array.




### Default Windows Partition

- By default Windows OS begins its partition at 32256
- By default customers build VMDKs starting at 32256
- Optimally, Windows VMDKs should begin at 32768

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Default Windows Partition

32256 is the default offset for Windows OS partitions. Unfortunately, if you let the VMware installation CD automatically create a virtual machine and accept the defaults for the virtual disk, it is going to create a partition starting off at 32256, which is not divisible by 4K, the size of blocks stored by WAFL.




## Impact of Misalignment

- Performance impact of misalignment
  - Varies based on workload
  - Smaller files have greater penalty
  - See [Recommendations for Aligning VMFS Partitions](#)
- Data deduplication impact of misalignment
  - Not completely clear
  - A customer experienced 6% savings then 64% space savings after properly aligning virtual disks
- Ask customers to correct their VM templates
  - Make future VMs perform optimally

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Impact of Misalignment

The performance impact of misalignment will vary depending on workload. It could be as much as 100%. The smaller writes are going to generate more 4K reads or writes, which is going to cause a bigger problem. It appears to have a significant impact on deduplication. NC State University ran deduplication and said, “Big deal, 6%, who cares.” And it turned out that every virtual machine in their environment was misaligned. As a test, they created representative virtual machines, properly aligned, and achieved 64% better performance, a 10-fold deduplication performance improvement. The improved performance was so compelling to them that they have been migrating to aligned virtual machines.



**VMDK Misalignment Issue is Vendor Neutral**

- The alignment issue impacts all storage vendors
  - Most customers are unaware of this issue
  - Do not offer your competitor this knowledge (and performance gain)
- Misaligned VMDKs is consistently a top 10 customer support issue, reports VMware

✓ Tip: Create properly aligned additional VMDKs when performing IO benchmarks / POCs

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMDK Misalignment Issue is Vendor Neutral

This issue impacts all storage vendors, not just NetApp. According to VMware, it is a top ten issue with every storage vendor. Customers are calling up VMware and saying, “I bought a really big array and I thought it would go faster.”

This issue is not widely known, and not one you should bring up while you're selling. If you're asked to compare your solution to the competitor, don't volunteer that your solution is properly aligned, and the competitor's may not be. If you need help, call a CSE. If you're going to run it on your own, have your CSE review your plan. Make sure you set up each virtual storage object on which you are going to be doing benchmarking so that it is properly aligned. You will have a 100 percent performance advantage over your competitor whose storage will likely be misaligned from accepting the default settings during operating system installation. If you don't have to do a POC, this issue should become a deployment concern and you should wait until deployment time to let the PS guys figure out how they're going to handle it, manage it, and communicate it. In a bake off, however, keep this knowledge to yourself.



## Creating Partition Alignment

- When creating a new VM, pre-create the VMDK prior to installing OS
- Three Options
  - Create VMDK, connect to existing VM for partition creation
  - Create VMDK, assign to new VM, boot with Windows Preinstall Environment CD for partition creation
- See NetApp TR-3428 for alignment issues description and alignment procedures

© 2008 NetApp. All rights reserved.


NetApp Confidential - Limited Use

## Creating Partition Alignment

When you create a new virtual machine, you first create the virtual disk and partition it before installing the operating system. You have three options for creating this partition with the proper offset for optimum virtual disk performance.

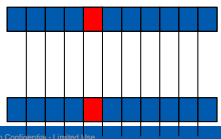
- You can create a new virtual disk on an existing virtual machine, and format it with the proper offset.
- You can create a virtual disk, create a virtual machine, boot off of a bootable CD that has the right tools such as the Windows Preinstall Environment.
- Or, you can distribute pre-aligned virtual disks VMDK files, download them to your ESX server through the service console, extract them and away you go.

This information can be found in is all in TR-3428.

 **VMDK Alignment Solution**

- Still based on blocks
- But, now they line up
- 1 for 1 IO operations!

**NTFS**




**WAFL**

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMDK Alignment Solution

All the blocks line up for optimum performance.



## VMDK Alignment Solution

1. Start with a FlexVol
2. Create a vmdk file (on either a VMFS or NFS Datastore)
3. \*Use diskpart (Windows) or parted (Linux) to fix the offset
4. Create the client filesystem

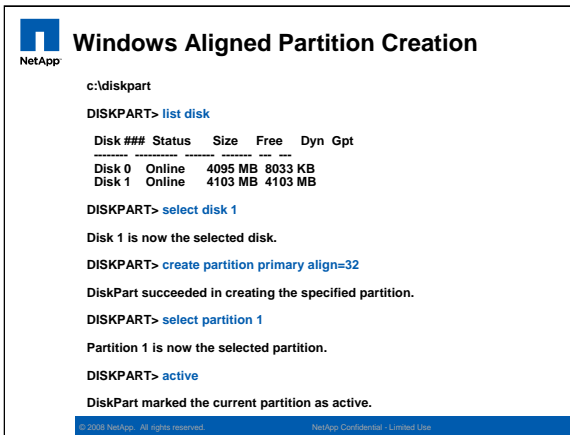
VMFS – or NFS  
datastore

WAFL – 4KB blocks

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMDK Alignment Solution

For example, in a Windows environment, you start by creating a FlexVol which is formatted by WAFL. WAFL has 4K blocks. Then create a vmdk file on either a VMFS datastore on a VMware LUN type, or on a NFS datastore. Use diskpart (Windows) or parted (Linux) to fix the offset. Finally, create the client filesystem.



```
c:\diskpart
DISKPART> list disk

Disk ##  Status  Size  Free  Dyn  Gpt
-----  -
Disk 0   Online   4095 MB  8033 KB
Disk 1   Online   4103 MB  4103 MB

DISKPART> select disk 1

Disk 1 is now the selected disk.

DISKPART> create partition primary align=32

DiskPart succeeded in creating the specified partition.

DISKPART> select partition 1

Partition 1 is now the selected partition.

DISKPART> active

DiskPart marked the current partition as active.
```

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Windows Aligned Partition Creation

Here is the diskpart command used with Windows:

- Select your disk:

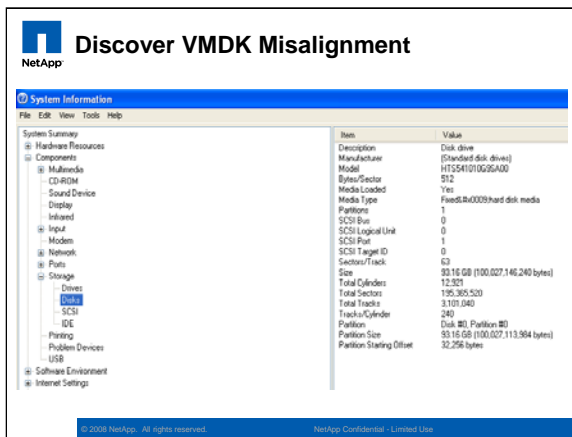
**select disk 1**

- Create the partition:

**create partition primary align=32.**


That is 32K or 32768. For SQL Server. or is it Exchange Microsoft best practice says 64K offset. As long as it is larger than 32 and it is divisible by 4K, the partition offset will not negatively impact performance.





## Discover VMDK Misalignment

This slide shows some examples of tools that can be used to discover misalignment issues.



## Migration to Achieve Partition Alignment

- Migration challenges
  - To convert existing VMs or not?
  - Impact on performance?
  - Impact on FAS Depuplication?
- ReplicatorX can correct misalignment
  - ReplicatorX agent inside of each VM
    - Requires VM reboot
  - Migrate to correctly created VMDK
  - ReplicatorX available as part of a PS engagement

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Migration to Achieve Partition Alignment

Block misalignment exists inside the virtual disk file. So if you copy the file—such as when creating a clone—ESX copies the bad alignment with it. The only thing that fixes misalignment is a volume level migration—either some sort of file level migration on the inside, or you can do ReplicatorX migration from a misaligned VMDK. Customers moving from another vendors arrays to NetApp have an opportunity to fix misalignment when they migrate.

ReplicatorX can correct misalignment using the agent inside each VM, not the fabric level agent. You run the ReplicatorX agent, pre-create a properly aligned volume on the destination and replicate to that. Professional Services is entitled to use ReplicatorX as part of a migration and take it out when they leave. if you have a customer that doesn't want to buy REPX, then sell them the migration and PS can use the REPX for the duration of the migration. It is an option.

Misalignment is one of those things that is not always worth chasing down. If you are getting 10X deduplication performance improvement, and your customer wants to go through the migration process, it will be worthwhile. If they are not experiencing a performance problem, it is an awful lot of work because there is no magic fix for this. Certainly have your customers correct their templates and any virtual machines that will be cloned frequently to avoid future problems.



## Resources and References


- [TR-3428 "Network Appliance and VMware Virtual Infrastructure 3 Storage Best Practices"](#)
- [TR-3593 "Storage Block Alignment with VMware Virtual Infrastructure"](#)
- [Recommendations for Aligning VMFS Partitions](#)
- [kb24492 "VMware LUN Alignment in ESX 3.x"](#)
- [kb8190 "Using partitions on Linux with NetApp LUNs may require alignment for best performance"](#)
- [ONTAP Blocks/Lun Types](#)
- [GNU Parted](#)
  
- NOTE: RDM's are not affected by the alignment issue.

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

## Resources and References

Refer to these resources and references for more details.



## VMDK Alignment Summary

- Understand the VMDK misalignment issue and its potential impact
- Create properly aligned VMDK files on Windows
- Discover VMDK misalignment in a virtual machine
- Migrate to achieve proper VMDK alignment

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMDK Alignment Summary

Now that you have completed Part 2 of VMware on NetApp Solutions, you should have a better understanding of :

- The potential for a VMware virtual disk storage block misalignment that could have a negative impact on storage IO performance;
- How to create properly aligned VMDK files on either Windows or Linux;
- How to discover VMDK misalignment in a particular VMware deployment; and when and how to
- Migrate virtual storage to achieve proper VMDK alignment

As well as when and how to run a VMware on NetApp proof of concept.



## Lab 5 Summary

- Create new virtual machines.
- Properly align the virtual disk partitions before installing the operating system.

© 2008 NetApp. All rights reserved.


NetApp Confidential - Limited Use

In this lab you will practice creating new VMs and properly aligning the virtual disk partitions before installing the operating system.





Next, let's take a quick look at NetApp's replication and disaster recovery solutions for VMware virtualized environments.



## Simple, Rapid, and Reliable DR

**Primary Site**

VM1 VM2 VM3

Storage


**DR Site**

VM1 VM2 VM3

Storage

SnapMirror®

- Practical DR for ALL apps
- Application consistent when combine with SnapManager
- Support for VMware Site Recovery Manager
- Replication of VMs and data
- Replicates only changed data
- Up-to-the-minute



"[NetApp] has really facilitated our move to a virtualized server environment, and that is allowing us to dramatically minimize the risk and duration of any business downtime."  
George White, CIO, Pennsylvania Office of the Attorney General

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 2

## Simple, Rapid, and Reliable DR

One of the most important values of VMware on NetApp is full data center recovery in minutes using SnapMirror. SnapMirror leverages Snapshot and SnapRestore and replicates data between sites. Unlike other solutions, SnapMirror is flexible enough and simple enough to use in various ways across all applications in the environment, as opposed to deploying only on the most critical applications. Additionally, the notion of application consistency is afforded through the use of the application-aware Snapshot copies. Replication is handled through backup scripts or through SnapManager for Virtual Infrastructure. SnapManager can also launch SnapMirror to send a copy of a Snapshot to the disaster recovery site. Scripts or SnapManager can then be used to restore a datastore, or an individual virtual machine in the event of a disaster, hardware failure, or data corruption. As with backup copies, you can set replication policies at a datastore level.

In the VMware environment, SnapMirror is integrated with VMware Site Recovery Manager (SRM), which manages the VMs and ESX servers across the sites. SRM automates fail over of VMware environments. It runs on an external server and monitors all of the ESX servers and virtual machines on a given VMware environment. Storage vendors including NetApp supply plug-ins for storage level replications. The SnapMirror plug-in makes SRM SnapMirror aware. In the event that you have a site disaster, you click a button on the SRM server and it handles the storage-level fail over as well as the ESX- and virtual-machine-level fail over. It will break mirrors and expose secondary copies of a datastores, mount those with the ESX servers on the fail over site and then start the virtual machines on the fail over site.

There are 14 storage vendors that all support SRM, so it is not a differentiator for NetApp storage technologies. Our replication technologies are the differentiators. For example, FlexClone allows live access to replicated virtual machines. You can create a FlexClone of the secondary copy of your data and use the clone to test DR procedures. FlexClone also gives you the option of running test and development on a clone of the DR data copy rather than on a clone of production data.

VMware does not provide the capability to directly copy virtual servers to a remote location, so VMware recommends using storage-based replication for both virtual servers and data. The NetApp replication solution (SnapMirror) is uniquely simple to deploy and administer.

The DR solution depicted in this slide requires neither servers nor storage systems be identical. ReplicatorX can extend DR capabilities to non-NetApp storage.

Regarding customer quote: Pennsylvania Attorney General used VMware to consolidate 120 servers; used FAS3000 for storage. The case study with details is available on netapp.com at <http://www.netapp.com/library/cs/penn-office-attorney.pdf?xCountry=US&xLanguage=EN>





## VMware Challenges with Disaster Recovery

- VMware doesn't have built-in DR or replication
- Storage requirements
  - Most solutions require > 2X space for replication
    - i.e. SRDF by EMC, EqualLogic, etc...
- No ability to 'test' DR implementation
  - Testing may require interruption to production environment or replication
- No "Clean" process for DR

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

3

### Challenges with Disaster Recovery

VMware ESX has no built in disaster recovery replication engine. VMware advises customers to look to storage hardware vendors or to third party software vendors for disaster recovery solutions. There is no simple process for disaster recovery.



## Why NetApp for Backup and Recovery?

- NetApp SnapVault and SnapMirror
- Industry's most flexible replication technology
  - Synchronous
  - Asynchronous
  - Semi-synchronous
  - TCP/IP
  - Fibre Channel Protocol
- Support for replicating both RDMs & VMDKs,
- FlexClone allows a DR location live access to replicated VMs for test, dev, training, etc...

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

4

## Why NetApp for Backup and Recovery?

NetApp has the industry's most robust, most flexible and easy to deploy replication topologies with SnapMirror and SnapVault. Your customers can at any time during the deployment change from a synchronous to an asynch or from replicating over fibre to TCP/IP or back and forth these are the options that are in your tool bag. NetApp supports backup and replication of both VMDKs and RDMs.



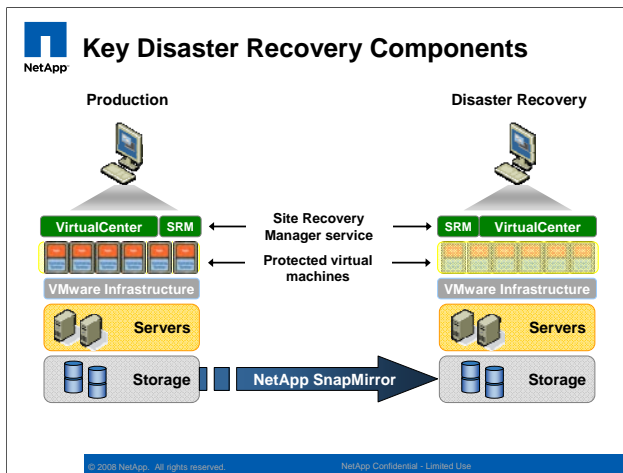
## Storage Savings Over Competitive Solutions

- Storage savings need to include production and backup media
  - Disk and/or tape
- NetApp data deduplication reduces primary storage by 50%
- NetApp data deduplication reduces DR storage by 50%
- Reduce time and bandwidth required for replication by deduplicating source data

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

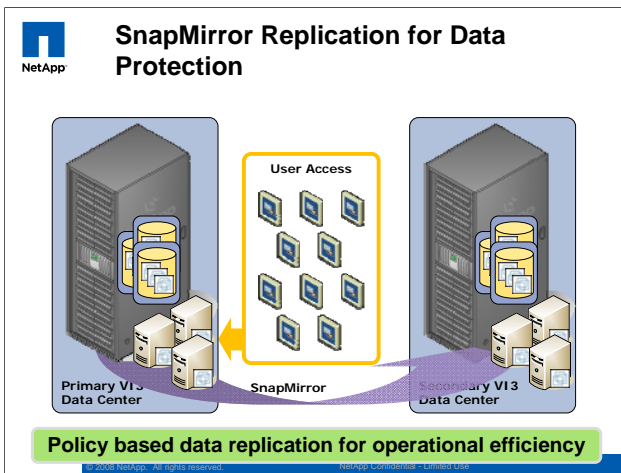
### Storage Savings Over Competitive Solutions

It is important to point out here as well that FAS deduplication can achieve 50% storage savings on both the primary and secondary sites. NetApp storage systems can provide two copies of your data—the production copy and a mirrored DR copy—all for the cost of your original data requirements.



## Key Disaster Recovery Components

This slide depicts the NetApp disk-based disaster recovery architecture for VMware ESX servers with NetApp secondary storage.



## SnapMirror Replication for Data Protection

Replication can be policy based. In this example, a number of datastores are being replicated to a secondary site. In the event of a disaster, you can simply point the ESX server or servers to the secondary datastores and restart virtual machine services.

**NetApp** ... and Disaster Recovery made simple ...

The screenshot shows a window titled 'Secondary Recovery V13 and vCenter Recovery' with a 'Recovery Completed' status. The main area contains a list of recovery steps, including:
 

- 1. What does virtual machine replication from the primary site.
  - 1.1. Run script about: Refresh from primary site
- 2. Create a new method for virtual machines.
  - 2.1. Run script about: All files, all partners
  - 2.2. Run script about: Cache Connection v13.1
  - 2.3. Run script about: Refresh from primary site
- 3. Add Snap Recovery Relationship.
  - 3.1. Run script about: Cache v13.1
  - 3.2. Run script about: Refresh from primary site
- 4. Discover resources for all virtual machines in this profile.
  - 4.1. Run script about: Refresh from primary site
  - 4.2. Refresh virtual machine recovery
- 5. Recover virtual machines.
  - 5.1. Run script about: Refresh from primary site
  - 5.2. Refresh virtual machine v13.1
  - 5.3. Run script about: Refresh from primary site
  - 5.4. Refresh virtual machine v13.1
  - 5.5. Refresh virtual machine v13.1
  - 5.6. Refresh virtual machine v13.1
  - 5.7. Refresh virtual machine v13.1

 A 'Recovery Completed' message at the bottom states: 'This recovery has successfully completed! To review the results of this recovery after closing this window, see the "Health" tool tab under the Recovery Profile.' To the right of the screenshot is an icon of a server rack labeled 'Secondary V13 DataCenter' with a yellow circle around it.

**Applications are back online in minutes - fast and easy!**

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

**... and Disaster Recovery made simple ...**

This fail over process today is manual, but will be automated by SRM. NetApp SnapMirror will be supported when VMware releases their site recovery manager 1.0 in upcoming months

This is an early screen shot of SRM recovery manager. SRM automates the fail over process in real time. It will discover all of the datastores, and all of the virtual machines, and discover and manage the SnapMirror relationships. It will do all of the storage device level work in order to get these secondary copies online. And then, once the copies are online, it will tell the ESX servers to mount and start all of the virtual machines.



## Resources

- [Network Appliance And VMware Virtual Infrastructure 3 Storage Best Practices](#)
- [ESX Server 3 Configuration Guide](#)

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use



## Lab 6: Summary:


- Simulate failure and recovery of a production storage environment.

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

In this lab you simulate a failure of an entire production storage environment. For the purposes of this exercise, your production and DR systems will be the same, which would not be the case in a real world scenario.





## VMware on NetApp Data Deduplication

First, we'll discuss NetApp's unique data deduplication value-add capabilities and how they benefit VMware deployments.



## Learning Objectives

- Deduplication on SAN and NAS based datastores

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

**When you have completed this module, you will have a basic understanding of:**

- How to get the most out of NetApp deduplication on SAN- and NAS-based datastores;
- VMware cloning technologies;
- When to use NetApp's cloning technologies with VMware;
- VMware VDI benefits and deployment challenges;
- When to use NetApp deduplication and FlexClone to facilitate rapid deployment of VMware virtual desktops; and
- How best to grow and shrink VMFS and NFS datastores.



## Unmatched Storage Utilization

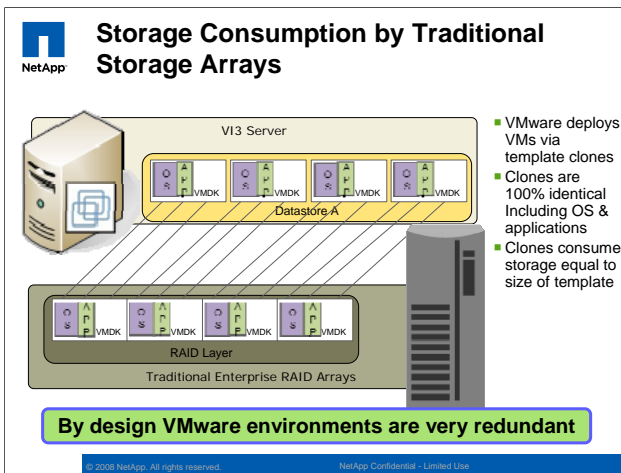
- Data deduplication reduces storage by 50% or more
  - Formerly known as A-SIS
  - Some customers have seen savings in the 80% range, but 50% is more commonly seen in aged production environments
- NetApp is the only storage vendor offering data deduplication for both production and DR data
- NetApp VMware storage solutions free up more storage than other vendors
  - ‘Storage solutions’ include production & DR disk and backup media

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Unmatched Storage Utilization

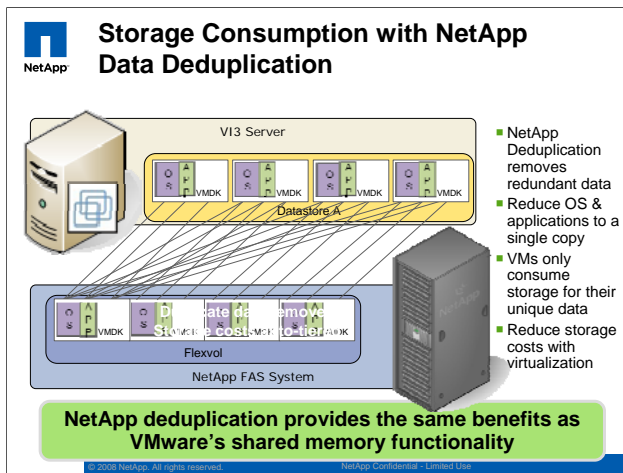
NetApp’s storage utilization story for virtualization is very strong. VMware data is highly redundant. When you create multiple virtual machines you are writing out the same OS and application binaries over and over again. NetApp Deduplication is able to recover as high as 50-60%, and even as high as 80% of the space used by the original virtual machines.

NetApp has a space savings calculator you can run on your own. It works, and will sometimes show potential savings over 80%; but, you should be very comfortable telling customers universally that they will see at least 50% space savings. It is better to under promise and over deliver—tell them 50% and have them hopefully get that or better rather than them be disappointed to only get 75% space savings when you promised 80%. 50% should be an exciting enough story..



## Storage Consumption by Traditional Storage Arrays

Deduplication works so well in a VMware deployment because you typically deploy new machines by cloning an existing virtual machine. As you deploy new machines, you get exact replicas of the original machine hundreds, or even thousands of times. Initially, each clone consumes the same amount of storage as the original machine.



## Storage Consumption with NetApp Data Deduplication

NetApp deduplication takes all of these clones and shrinks them down to the size of a single instance of the original virtual machine. Why store 100 identical copies of Windows when you can just reference the same data 100 times? Rather, save expensive tier 1 enterprise storage for high value, data.

It is often an eye opener to point out to a potential customer the parallels between NetApp features and VMware features. VMware has a shared memory functionality whereby common page files are shared between multiple running virtual machines. You can draw a parallel between storage sharing and memory sharing. If your prospect understands the concept of shared memory in ESX, then should immediately see the benefit of NetApp's storage sharing model. .



## Increased Storage Utilization

- Data deduplication on production data is a key differentiator for VMware on NetApp
- Deduplication policies are set at the datastore level and can be disabled if needed
- For optimal deduplication results, group similar OSs into datastores

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

6

## Increased Storage Utilization

NetApp is the only storage company offering deduplication of production data, with the possible exception of Data Domain in low IO environments only. NetApp deduplication can run in almost any production environment.

NetApp's Deduplication is done at the volume or datastore level. You can also turn it off, or run it on a schedule. For optimum deduplication results, similar OSs should be grouped into a datastore. We discussed in the Backup, Recovery, and Replication module of this course grouping virtual disks in datastores according to backup schedules and replication schedules. Deduplication is another reason to separate storage content into separate datastores. For example, if you are running Linux VMs and Windows VMs and want the best possible deduplication results, you should create a datastore for each.

These storage layout recommendations are different than and usually more complex than legacy customers may be used to. You need to help them do a benefit vs. complexity analysis. Often the more granular the datastores, the more effective the deduplication will be. Sometimes storage layout will be a deployment consideration rather than something that has to be decided during the sales cycle, unless you are explicitly depending on some of these results for sizing or performance.

**NetApp** **Data Deduplication with SAN**

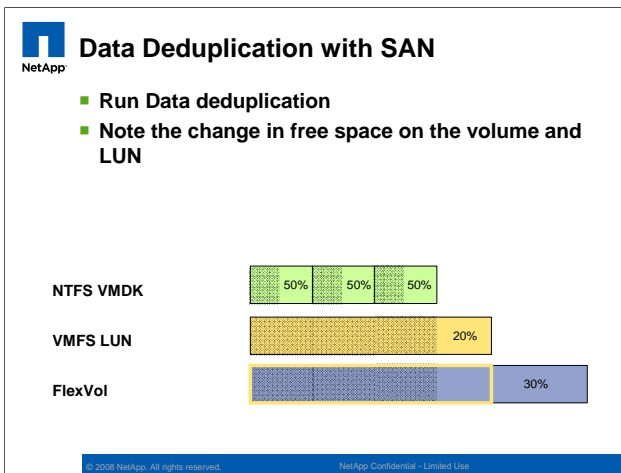
- Create a FlexVol
- Create a VMware LUN – note the reserved space
- Deploy 3 Virtual Machines
- Note the data written in the VMDKs, LUN, and FlexVol
- Note the free space in the VMDKs, LUN, and FlexVol

Storage Component	Utilization
NTFS VMDK	50%
VMFS LUN	20%
FlexVol	30%

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Data Deduplication with SAN

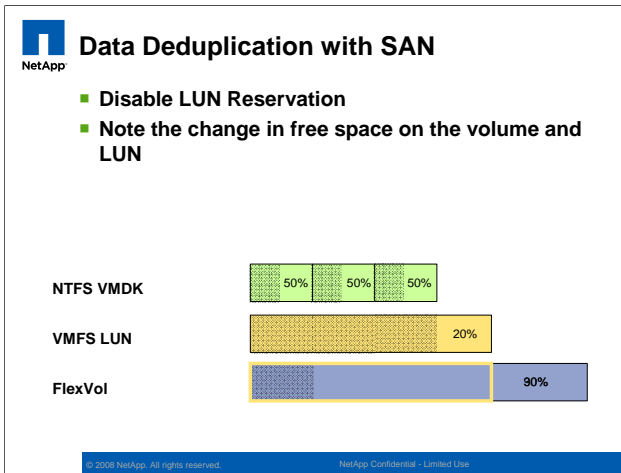
The next several slides depict how deduplication works when you are deploying VMware on NetApp storage. You start with a SAN-based VMFS environment over fiber channel or iSCSI. First you create a Flexible Volume and a VMware type LUN inside that volume. When you create the LUN, it reserves space within the volume. Then you create some virtual machines and their virtual disks inside the LUN. Look at the space reserved by the NetApp storage system versus written data. Note that in this example, each of the virtual disks is 50% full, the VMFS datastore is 80% full, and the NetApp storage system volume is 70% full.



## Data Deduplication with SAN

Now run the CIFS deduplication process on the NetApp storage system volume. See what happens? The duplicate space is compressed down to the size of just one copy. In some cases, the deduplication results could end up taking the space of less than the original virtual machine if that machine contained redundant data. In the example shown in this slide, we deduplicated three copies of the VM down to the size of one. Deduplication didn't change the amount of free space reserved by the LUN reservation. It remains the same. This is a traditional, thick provisioned LUN. We saved all of this space at the volume level, but have realized none of it because the same amount of storage space is being reserved. You can't use it for anything else. And, none of these savings are available to the VMware administrator.





## Data Deduplication with SAN

Now enable LUN thin provisioning by removing the space reservation within the volume. The volume goes from 70% full to only 10% full because now you only have to pay for this amount of space. But, from the VMware side, we still have not realized any space savings. There are two ways to take all of this space that you saved and deliver to the VMware administrator and to the end user. You can either expand the existing LUN or you can create a new LUN and provision it as a brand new datastore.

**NetApp** **Data Deduplication with SAN**

- To add space to the LUN, grow the LUN and add a VMFS extent
- Note the change in free space on the volume and LUN
- Add a new Virtual Machine

The diagram illustrates storage utilization for three different storage configurations:

Storage Configuration	Utilization
NTFS VMDK	50% (repeated three times)
VMFS LUN	20% (original) + 40% (new extent) = 60%
FlexVol	90%

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Data Deduplication with SAN

In this example, you can run the LUN size operation and change the size of the LUN. Now when you rescan, you see that the physical LUN is larger. But you can't just expand a VMFS file system like you can NTFS. You can only add to it as an extent, which is a logical grouping of resources. In this case, the extent is another area in the same LUN. An extent could also include multiple LUNs. You can take two LUNs and make a VMFS extent and the VMFS volume would span both LUNs. Now, your effective free space in the VMFS volume has changed. You can add another virtual machine.



## Data Deduplication with SAN

- Deduplication savings require work before savings are realized
- LUN thin provisioning is required to free up volume capacity
- VMFS extent is required on LUN to add space
- Create a new datastore in a new LUN as an alternative to using extents

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

11

## Data Deduplication with SAN

To summarize, in a SAN environment, running the deduplication scan doesn't free usable space. The only way to free up space is to enable LUN thin provisioning. That frees up space from the storage side, but not from the VMware side. And, using extents to grow VMware datastores is not best practice—except in an emergency. Instead, you should provision that space as a new LUN in a new datastore.

**NetApp** **Data Deduplication with NAS**

- Create a FlexVol
- Connect the volume as a datastore
- Deploy 3 Virtual Machines
- Note the data written in the VMDKs, datastore, and FlexVol
- Note the free space in the VMDKs, datastore, and FlexVol

The diagram illustrates the storage usage in a deduplicated environment. It shows three rows of storage components:

- NTFS VMDK:** Three separate blocks, each representing a virtual machine's disk. Each block is 50% full, indicating that each VM has written unique data.
- NAS Datastore view:** A single block representing the entire datastore. It is 50% full, showing that deduplication has occurred, as the shared data from the three VMs is only stored once.
- FlexVol:** A single block representing the NetApp storage system volume. It is also 50% full, demonstrating that the deduplication is reflected in the physical storage usage.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use


## Data Deduplication with NAS

In a NAS environment deduplication is simpler. Just as before, you create a FlexVol and then mount the volume as a datastore. In this example, we have created three virtual machines and each is half full. At this point, the datastore and the NetApp storage system volume are also half full.



## Data Deduplication with NAS

- Run data deduplication
- Note the free space in the VMDKs, datastore, and FlexVol
- Note How the storage savings are seen from the datastore – no need to grow the datastore
- Add a new Virtual machine



© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Data Deduplication with NAS

Assume now that you run deduplication to remove duplicate data in order to generate free space. If the free space within the volume goes up, to 90% for example, the logical representation stays exactly the same and nothing changes about the virtual disks and the volume itself, but, free space within the datastore will change. The net size of the datastore will be the same, but the free space will have increased. Then, you can create more virtual machines with no need to create an extent.



## Data Deduplication with NAS

- **Deduplication savings require no additional work**
- **With NAS, storage savings are immediately realized in VMware datastores**

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

14

## Data Deduplication with NAS

So, in a NAS environment, deduplication is no additional work. Savings are realized immediately, with no need to create new datastores and no interaction between the storage administrator and the VMware administrator.



## Lab Summary

- Lab 7: Run deduplication on an NFS datastore and observe the effect on storage consumption.
- Lab 8: Enable deduplication on a VMFS datastore and observe the effect on storage consumption.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Lab Summary

In Lab 7 you run deduplication on an NFS datastore and observe the effect on storage consumption. In Lab 8, you enable deduplication on a VMFS datastore and observe the effect on storage.







Next, we'll cover the best ways to use of VMware and NetApp cloning technologies in VMware infrastructures.



## Learning Objectives

- VMware cloning technologies
- When to use NetApp FlexClone with VMware
- VMware Virtual Desktop Infrastructure (VDI) benefits and deployment challenges
- When to use NetApp deduplication and FlexClone to facilitate rapid deployment of VMware Virtual Desktops

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

**When you have completed this module, you will have a basic understanding of:**

- VMware cloning technologies;
- When to use NetApp's cloning technologies with VMware;
- VMware VDI benefits and deployment challenges;
- When to use NetApp deduplication and FlexClone to facilitate rapid deployment of VMware virtual desktops; and



## VMware Cloning Technologies

- VMware allows the cloning of individual Virtual Machines
- VMware clones are copies of files
  - .vmx, .vmdk, etc
- Clones take time relative to the size of the files being cloned
  - Average VMs take 10-30 minutes
  - Not uncommon for large VMs to take ~ an hour
  - Cloning operations are serialized
- A VMware clone of a thin-provisioned VMDK will become a thick-provisioned VMDK

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

3

## VMware Cloning Technologies

Cloning of virtual machines within a datastore is a feature that is available in ESX. It works very well to generate a new virtual machine by cloning an existing VM. If you have a gold master image of a server, you can use it to provision new servers for building out. You can use the ESX clone feature to create a copy of the configuration file and a copy of the virtual disk file, resulting in two identical virtual machines. You can then install applications on the clone. Customers are pretty familiar with it because it is easy to do. It's just a couple clicks within VirtualCenter. But, they do have to wait while the data copies. If they are cloning a number of virtual machines, each containing a large virtual disk file, they have to wait for all those copies to be completed.



## NetApp Clones with VMware

- Clones are available of LUNs or volumes
  - With physical servers this makes a lot of sense
  - With VMs this value does not directly apply
- Clones of LUNs or volumes are VMware datastore clones
  - Datastores serve many VMs to many V13 servers
- Where are datastore cloning use cases?
  - Mass deploy many VMs
    - Requirement of VMware Virtual Desktops
  - Clone a set or group of VMs which function together
    - Test dev, training, etc...

© 2008 NetApp. All rights reserved.

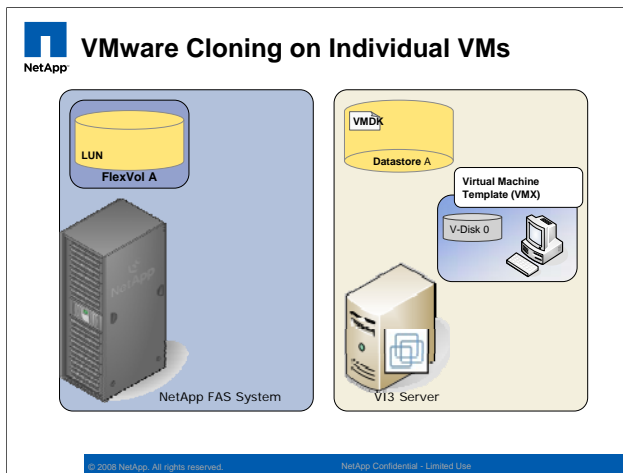
NetApp Confidential - Limited Use

4

## NetApp FlexClone with VMware

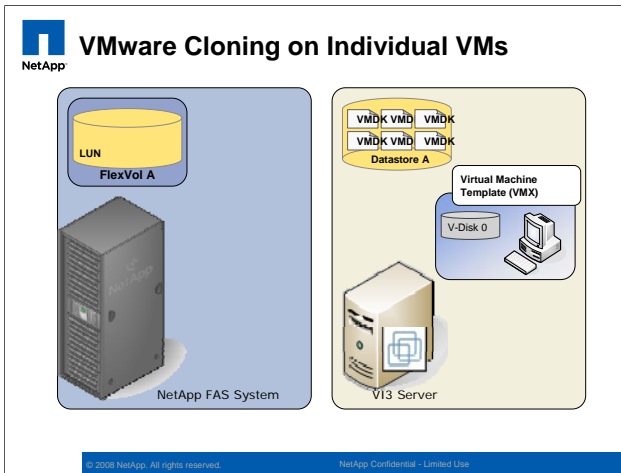
NetApp clone technology also can be used to provision new servers. If you've got just a traditional VMFS file system in a Fibre Channel environment, however, the NetApp clone feature doesn't offer a significant advantage over the native ESX clone feature. If they're using NFS, NetApp clone does offer the benefit of performing the clone procedure from the storage side, thus reducing the load on the ESX server. Also, if using RDMs, you can clone the RDM itself using LUN clone and then split the LUN clone, which also removes the load from the ESX server.

So far, all of these cloning solutions are not saving any storage space. The real value of NetApp clone in a virtual infrastructure is realized when you use it to create temporary virtual machines, especially a large number of VMware Virtual Desktops, or if you are using VMware to provision a test dev center, a demo center, or a training center and you needed 30 servers for testing or for demo. In a traditional VMware environment, 30 copies of a server always take 30 times the storage needed for the original machine, and you have to wait while that data copies 30 times. So, obviously, it can be expensive to provision large numbers of virtual machines in a traditional VMware environment.

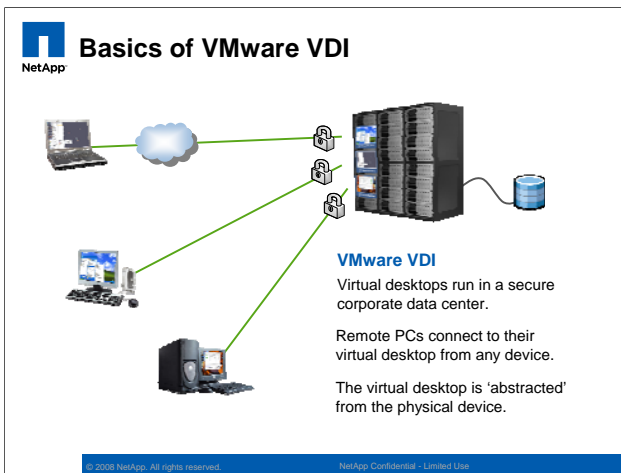


## VMware Cloning on Individual VMs

By contrast, you can use NetApp's clone or LUN clone to quickly provision a large number of virtual disks on NetApp storage, and then attach new virtual machines to the cloned drives. Because of NetApp's cloning technology, the storage space consumed by the cloned virtual disks is only a fraction of the space that non-NetApp storage would use. So when you have customers who need lots of virtual desktops or are constantly creating and recreating temporary machines, NetApp clone and/or LUN Clone technology will provide significant space savings while reducing dramatically the time needed to complete the cloning process.



## VMware Cloning on Individual VMs



## Basics of VMware VDI

Let's consider a scenario for using FlexClone to provision a VMware Virtual Desktop Infrastructure deployment.

VMware Virtual Desktop Infrastructure (VDI) is an integrated desktop virtualization solution through which virtual desktop images (typically Windows XP or Vista) run on servers in the data center. Users can connect to the virtual desktops from any client device in any location. VMware VDI sits in between the client and the server hardware. Regardless of the client, users get a complete, well-managed virtual corporate desktop that behaves just like a normal PC configured with corporate applications and with access to corporate data.



## VDI Benefits: Centralization

- Close to IT Operations
  - Desktops are located in the Data Center – close to support staff.
  - Management tools can access desktop VMs over high-speed local networks for patching and maintenance.
- Application Performance
  - Close proximity to servers maintains application performance, regardless of where the client is located
- Security
  - All data resides on a secure network inside the corporate firewall.
  - Remote users only 'view' data - it never gets transferred to insecure devices which might be lost, stolen, or hijacked.



© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

## VDI Benefits: Centralization

Organizations use VMware VDI to streamline desktop management, increase security, and reduce costs. Using VMware VDI, the time it takes to deploy a desktop is typically reduced to minutes, optimizing the value of IT resources and getting end users productive faster. VMware VDM is a flexible and intuitive desktop management server enabling IT administrators to quickly provision and tightly control user access.





## VDI Benefits: Compatibility

- Real Desktop Operating Systems
  - Virtual Machines run Windows XP just like physical hardware so applications work normally without modification.
  - Applications can make system level changes such as registry writes and DLL replacements where necessary
  - Existing corporate desktop configurations can be easily imported.
- Isolation
  - Each Virtual Machine runs separately, so a crashed VM or a poorly behaving application does not affect other users on the same server.
  - VDI is suitable for developers because any type of change can be made to a VM without affecting other users.



© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VDI Benefits: Compatibility

To the end user, the environment is a real desktop computer system. Whether it's a Windows 2000 workstation, an XP desktop, or a Vista workstation, they get a real operating system environment with applications loaded that behave just as they would on a physical machine. The end user's quality of service remains high.



## VDI Benefits: Virtualization

- Homogenous Virtual Hardware
  - All Virtual Machines use the same virtual hardware.
  - One base image can be used for many different VMs – driver management is greatly simplified.
  - Templates can be created in Virtual Center to aid rapid deployment of new Virtual Machines
- VDI is Virtual Infrastructure
  - V13 Servers and Virtual Center are the core components.
  - Virtualized Server and Desktop spaces can leverage the same hardware, architecture, and infrastructure.
  - Automatic Load-Balancing and High Availability through VMware HA, DRS, and VMotion features.



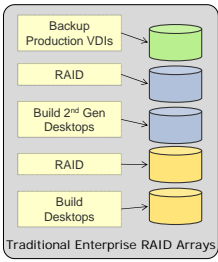
© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VDI Benefits: Virtualization

VMware VDI uses the same core components as VMware Virtual Infrastructure. The bulk of the virtual desktops in a typical deployment can be standardized across the enterprise, although they may be tailored for different departments. You may have one type of desktop that has the enterprise load—the enterprise service packs and patches and tools like antivirus, and Microsoft Office—as well as department-specific application packages, such as a toolset that's just for human resources and that's going to support the human resources department, for example. You may have another desktop configuration that has the corporate standard load, and yet another with tools for sales administrators for the sales administrators group. Each of these cases is a version of the same desktop. When a user is working from a desktop, they don't get to save data to these virtual machines and they don't get to install the applications. Work files and user-interface customization, etc, get redirected to their home directories on shared network storage.

This solution enables central control of patch management, and software deployment from a central copy after they have been tested and validated in lieu of distributing patches and updates to thousands of clients one desktop computer at a time. If one of the users starts to use a lot of system resources causing a server to be overloaded, VMware uses DRS and VMotion can move the virtual desktop dynamically to a server with a lighter load.

**NetApp** **The Challenges Associated with Deploying VDI**

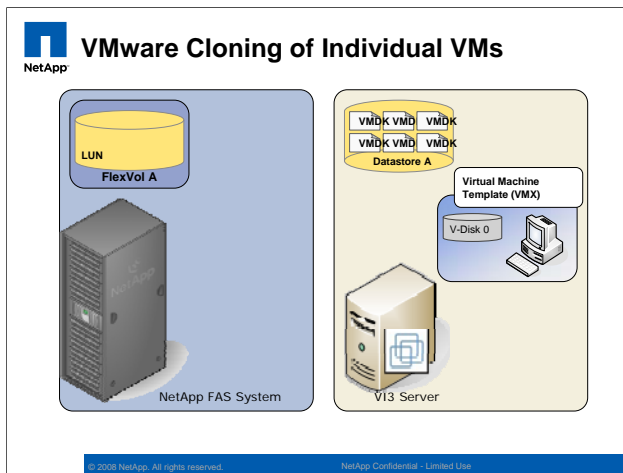


- Storage Acquisition Costs
  - VDI Requires Enterprise Class Storage
    - Data Protection
    - High Availability
    - Backup
  - VDI Provides Central Desktop Updates
    - Requires Additional Storage
- Lengthy Mass Deployment Timeframes
  - VMware Cloning 100s – 1,000s of Virtual Machines Takes Time
    - 6-12 Minutes Per VM or 5-10 VMs deployed per Hour

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## The Challenges Associated with Deploying VDI

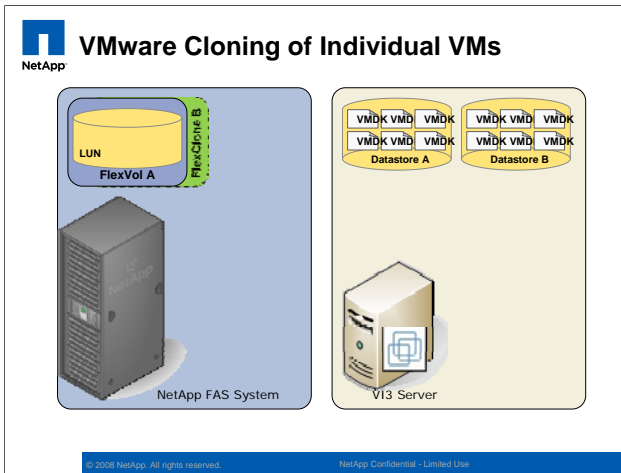
Now let's cover the challenges associated with VMware VDI. When you move from PC-based storage to centralized storage, you need a high level of data protection, and you start to eat a lot of storage. VMware VDI is a great product, and everybody loves it but if you have customers that want 30,000 desktops, 50,000 desktops, 80,000 desktops, the cost of storage and all the operations required for a VDI deployment can become a roadblock.



## VMware Cloning of Individual VMs

This is where NetApp's storage virtualization technologies can play a key role in VDI deployments. A 4 minute video is attached to this module for optional download. The video, which was shown at VMworld, shows the cloning of 100 virtual machines, each of which was 10GB in size. The cloning itself was completed in two minutes without consuming any additional storage. Let's walk through the steps needed to clone a large number of virtual machines, as was shown in the video.

1. First, build a datastore and create a virtual machine, the original virtual machine that will be the prototype for the cloned virtual machines. For Windows desktops, use Sysprep to ensure that when the virtual machines are cloned, they are recognized by the operating system as unique machines.
2. Next, take a snapshot of that datastore and create a FlexClone. You don't want to use the original copy just in case something goes wrong. Then mount the FlexVol in ESX.
3. Now, create VMware clones of the original virtual machine inside the datastore. This slide shows 6 virtual machines. The example shown in the video took one virtual machine and cloned it 24 times. At this point you have a datastore containing multiple clones of the original virtual machine.
4. Run FAS Deduplication on the datastore to reduce the consumed storage space back down to the size of the original machine. Note that the video shows the cloning and deduplication in one step. That feature of Data ONTAP will be available in a future release.

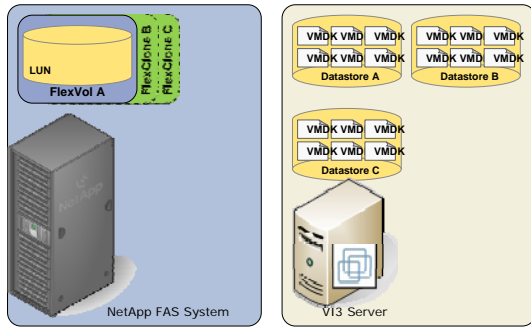


## VMware Cloning of Individual VMs

- Now, use NetApp's FlexClone to create multiple cloned datastores. Every FlexClone contains multiple virtual machines, but consumes no additional storage space.

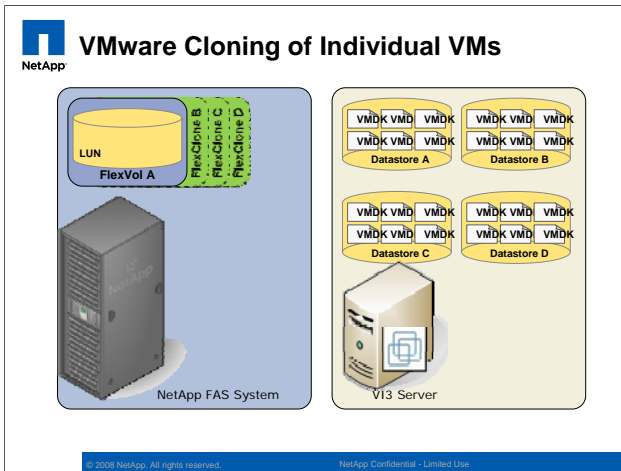


## VMware Cloning of Individual VMs

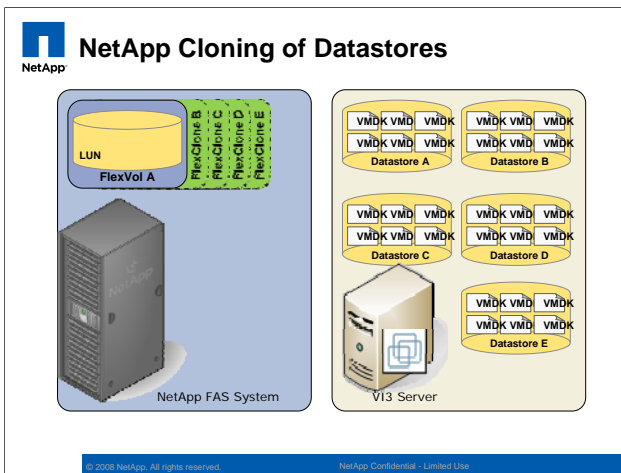


© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMware Cloning of Individual VMs



## VMware Cloning of Individual VMs



## NetApp FlexClone Clones Datastores

- After creating all the clones, all that's left is registering them in the VirtualCenter, and starting them up. In a customer deployment, you can simply have a script that slowly boots the machines up. That's it, that's the process. You went from one to many virtual machines without consuming any additional storage, and you did it very quickly.





## Resources

- [Network Appliance and VMware Virtual Desktop Infrastructure](#)

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use



## Lab 9 Summary

- Use FlexClone to rapidly provision new virtual machines.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Lab Summary

In Lab 9 you use FlexClone and FAS Deduplication to rapidly provision new virtual machines, run deduplication on an NFS datastore, and observe the effect on storage consumption.



Module 9:  
**Growing  
Datastores and  
Virtual Disks**



NetApp Confidential - Limited Use

This module covers growing VMware datastores on NetApp storage.



## Learning Objectives

- Growing and shrinking
  - VMFS and NFS datastores
  - Virtual Disks

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

In this module we will talk a little more about growing datastores and virtual disks.



## Growing Datastores

- VMware allows the expansion of individual Datastores
- Supports growth of both VMFS and NFS datastores

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

3

## Growing Datastores

We have already discussed growing VMware datastores and extents in the context of deduplication in a SAN environment. In fact, you can grow a datastore, whether they are a VMFS datastore or a NFS datastore. Both are online operations. Growing a VMFS datastore is not a best practice, but it can be done in an emergency,



## Growing VMFS Datastores

- Expand datastores by increasing the size of the LUN.
- More LUNs do not increase performance.
- An extent is a partition.
  - The file system appears as one contiguous file system.
- A datastore has a limit of 32 extents.
- Expansion requires an administrative task to be completed on every ESX server in the data center.
- VMFS datastores cannot be shrunk

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

4

## Growing VMFS Datastores

When you expand a VMFS datastore, NetApp best practice is to increase the size of the LUN that contains it rather than distributing the datastore across multiple LUNs. There is no performance advantage to distributing the datastore across multiple LUNs.

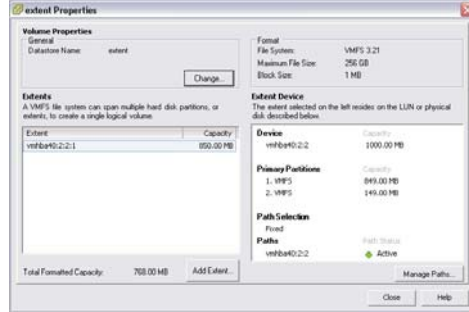
VMFS datastore expansion always requires creation of an extent. An extent is a partition. A data object or file cannot be split between extents. Each file needs to be in a single extent, not spread across multiple extents. You can have a maximum of 32 extents in a VMFS datastore. So, you can only grow a VMFS datastore a maximum of 31 times. When you extend a datastore, you have to manually rescan to detect the new LUN, and to recognize the new size of their datastore.

And, it is impossible to shrink VMFS datastores. If you want to shrink a VMFS datastore you would have to evacuate it, destroy it, and recreate it at the smaller size.



## VMFS Extents

- Extent Properties



## VMFS Extents

This screen shot shows the extent Properties screen. As you can see here, a VMFS file system can span extents, and disk partitions to create a logical volume. This properties screen shows you the relative size of each extent, but you can't get any more information than that to determine the contents of each extent.



## Growing NFS Datastores

- Expand or contract a datastore by increasing or decreasing the size of the FlexVol.
- New datastore size is immediately recognized on every ESX server in the datacenter.
  - VirtualCenter may not report the storage change.
  - VirtualCenter is poll based, not real time.
- Expansion does not have any of the caveats of VMFS.
  - No extents required
  - No limits
  - No server side configurations

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

6

## Growing NFS Datastores

By contrast to VMFS datastores, NFS datastores are nice and easy to expand and contract. You just use the `volsize` command. You can grow or shrink the datastore dynamically. It is immediately recognized by the ESX server. It might not show up immediately within VirtualCenter or the VI client. The client is not a real time view of what the server is seeing. If you go to the service console and execute a `DF` command, you see a real time number that might be different than the number reported by VirtualCenter or the VI client. But eventually VirtualCenter and the client will catch up.

Since VMware isn't applying a file system into the NFS volume, extent limits and volume limits don't apply. There is no limit to the number of times that you can grow or shrink an NFS datastore.





## Storage VMotion

- Perform storage migration while VM is running
- Can be used as an alternative to growing datastores
- Currently not fully supported (Experimental) for NFS datastores

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

A promising new feature in VI 3.5 is Storage VMotion. It is the capability to dynamically move the storage resources of a virtual machine from one datastore to another while the VM continues to run. In VI 3.0 you can move a virtual machine from one datastore to another, but you first have to power down the virtual machine. Then you right click the virtual machine, and select the migrate option. The migrate option then copies all of the files to the other datastore. When the copying is complete, you can start the virtual machine from the other datastore. The new Storage VMotion feature allows you to perform storage migration while the virtual machine is still running. It starts mirroring the files from one datastore to another. When all the files have been copied, the ESX server brings all of the services over without stopping the VM. There is just a pause when the VM switches from one datastore to another.

This Storage VMotion feature could be a solution to the datastore management problem that arises if you need to grow a datastore. Storage VMotion gives you a much better path to move storage resources from one datastore to another. You should expect performance to be pretty bad during the migration, but this feature does give you the option of migrating a virtual machine between datastores, perhaps during a period of low activity, without having to first shut down the virtual machine.

Storage VMotion is not certified between protocols. You can't use this feature to migrate virtual disk resources from VMFS to NFS and vice versa. Storage VMotion support for NFS is currently in experimental status.



## Growing Virtual Disks

- VMDKs can be extended, but only when VM is powered off.
  - vmkfstools -X <new size>
- File system must be extended manually
  - For windows: diskpart
  - For Linux: ext2resize
- Extend root file systems by connecting them to a different VM
  - C:\ or / are root file systems

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Growing Virtual Disks

It is also possible to grow VMDKs, but it is an offline operation. The virtual machine to which the virtual disk is assigned must first be shut down. It is a two step process for extending the virtual disk.

First, resize the virtual disk file, which is equivalent to the physical disk itself. In VMware VI versions 2.5 and 3.0, the command is

**vmkfstools -x**, the size and then the path of the disk.

This disk command is available in the GUI for VI version 3.5. You still have to shut down the virtual machine.


But expanding the size of the virtual disk doesn't automatically change the file system that resides on that disk. So if you have extended anything besides your root disk, you then power on the virtual machine and use the file system tools, such as the **diskpart** command in Windows. This is very similar to what you do with SnapDrive when you dynamically extend a disk. In that case, you physically extend the LUN and then use **diskpart** to extend the partition within the LUN. So, in this case, after you boot up your virtual machine, then you use the file system tool to extend the partition within the expanded virtual disk.

Now that works for data drives; but, since you can't dynamically extend the root drive, you won't be able to dynamically expand a virtual disk on the root drive. If you want to extend the C drive of a virtual machine, you shut down that virtual machine, use the appropriate command to extend the virtual disk and then mount that virtual disk on any other virtual machine. Then run the **diskpart** command from that virtual machine. Then shut down, disconnect the expanded drive, and give it back to the original owner of that disk.



## Growing Virtual Disks

- VM must be powered down / off
- File system still needs to be expanded (i.e. diskpart)



## Growing Virtual Disks

From the VI Client inventory in ESX 3.5 you can use the Capacity section of the Virtual Machine Properties window's Hardware tab, shown in this screenshot, in lieu of the service console's vmkfstools command to change the size of the virtual disk. Otherwise, all of the same restrictions apply that are discussed in Part 4 of this course: You still have to shut down the virtual machine and extend the file system manually; and, you still have to go through extra steps to change the size of the root drive.

Notice that the this window manages properties at the virtual machine level, not at the datastore level.



## Summary

- Growing and shrinking
  - VMFS and NFS datastores
  - Virtual Disks

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Summary

Now that you have completed this module, you should have a basic understanding of how best to grow and shrink VMFS and NFS datastores.



## Lab 10 Summary

- Grow existing VMFS and NFS datastores.

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

### Lab Summary

In Lab 10 you grow existing VMFS and NFS datastores.





Next, we'll discuss sizing and capacity planning for VMware environments on NetApp storage.



## Learning Objectives

- Sizing NetApp storage for VMware environments
- Planning and designing VMware on NetApp connectivity

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Learning Objectives

When you have completed this module, you will have a basic understanding of:

- Sizing storage for VMware environments including available Sizing Tools;
- Planning and Designing VMware Connectivity; and
- VMware on NetApp Performance monitoring using VirtualCenter and the VMware esxstop command;





## Sizing VMware Solutions

- Sizing storage for VMware has 3 aspects
  - Capacity
  - Performance
  - Connectivity
- Ongoing process
  - Initial exercise to build environment
  - Ongoing to keep pace and to predict demand and requirements

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

### Sizing VMware Solutions

When sizing a VMware storage solution, be sure to always address the issues of:

Capacity—how much space do you need to store the virtual machine images, and swap space;

Performance—the raw disk IO throughput needed by the environment; and

Connectivity—IP or fiber channel, and the number of physical links to the ESX servers and to the storage device itself?

While sizing is a critical part of architecting a VMware on NetApp solution, don't forget that it is an ongoing process. You measure capacity, performance, and connectivity needs when you are initially analyzing a customer's environment, but depending on the sales cycle, you may need to measure it again to account for possible growth.



## Which Counters?

- We need to know
  - Read vs. write
  - Random (IOPS) vs. Sequential (MB/s)
  - Queue length vs. idle time
    - Tells us if they're already providing maximum available throughput, and if we should plan for MORE performance than currently being delivered.
- Physical disk

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Which Counters?

At this point, your sizing analysis should be very similar to how you approach any other custom application—read vs. write parameters, random or sequential IO, and queue length versus idle time, all measured at the physical disk. At the end of the day, you are going to plug the data you gather into the custom sizing tool to generate NetApp hardware sizing recommendations.



## Starts with Data Collection

- VMware Capacity Planner
- Other tools including
  - NetApp's VM Insight (Onaro)
  - Perfmon
  - Logman
  - Platespin

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

Often the most difficult part of this analysis is collecting data from a distributed environment. That is the purpose of the VMware Capacity Planner—the data collection tool recommended by VMware. We'll cover Capacity Planner first and then briefly cover other available tools, including NetApp's own VM Insight that came over with the recent Onaro acquisition.



## VMware Capacity Planner

- Paid engagement by VMware or authorized partner
- 1-4 data collectors installed (\$\$ each)
- Collects data on servers with various OS
  - Even OS they don't virtualize!
- Data collected for 1-2 months
- Auto-upload to [optimize.vmware.com](http://optimize.vmware.com)
- Auto-analysis and report generation

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VMware Capacity Planner

Data gathering using VMware Capacity Planner is a paid engagement. Your customer contracts either directly with VMware or with a VMware VAR who then installs Capacity Planner data collectors to gather information on all of the machines in the environment that are going to be virtualized. The collectors in turn automatically upload the data to VMware's [optimize.vmware.com](http://optimize.vmware.com) Web sight. From the collected data, VMware SEs generate capacity reports and recommend virtual server configurations.



## Capacity Planner Data Collector

- Discovery
  - Find servers to analyze
- Inventory
  - Collect manufacturer, model, CPU, BIOS, memory, disk, network, etc.
- Performance
  - Hourly task to collect a default set of counters
- Data sync
  - Auto-upload to optimize.vmware.com
  - Hourly by default

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Capacity Planner Data Collector

Capacity Planner includes four processes:

1. Find the servers in the environment;
2. Take an inventory of the physical resources;
3. Determine and log real time performance statistics; and
4. Upload this information to the VMware data collection Web site.

All of the analysis is done in Palo Alto.



## Capacity Planner Disk Performance Metrics

- Inadequate To Size Storage
- Report lists peak per server
- Can customize collection job to
  - Collect disk-specific counters more frequently
  - Keep in CSV files (One per interval with all servers and counters)
- Customizing collection job overwrites VMware's collection job
  - Use a separate collector or wait until they're done
- Disable auto-upload job
  - VMware doesn't want your data
  - Once uploaded, data is deleted from data collector

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Capacity Planner Disk Performance Metrics

The reports generated from default Capacity Planner data collection will not provide all the sizing information that you need. Typical Capacity Planner output shows only peak disk activity per server, which tells you nothing about what the servers are doing on a day to day basis.

You can set up a custom Capacity Planner data collection job to gather the information that you need, but a data collector can't run your custom job while another data collector job is running. The customer either needs another data collector, at additional cost to the customer, or you will have to wait your turn. And, when you run your custom job, don't forget to turn off the automatic upload of the collected data to VMware's collection Web site. VMware won't know what it is and they won't want it, and what's worse, the job will automatically delete the data files from the data collector before you've had a chance to analyze it.

**NetApp VM Insight**

- Real-time service-level information on:
  - CPU
  - Memory
  - I/O
  - Network bandwidth
  - SAN bandwidth
  - Storage array performance
- Common view of network storage service paths and changes
- Complete datacenter view of both physical and virtual loads impacting the storage environment

Name	State	Volume	Volume ID	Datastore	FCID	Datastore Capacity	Datastore Free Capacity	Virtual Machine Capacity
60	running	33	250.00	33		249.00	65.00	13,312.0
20	running	41	250.00	42		249.00	49.00	12,400.0
95	running	42	250.00	41		249.00	37.00	14,238.0
20	poweredOff	42	250.00	41		249.00	37.00	9,000.0
70	running	44	250.00	44		249.00	51.00	14,336.0
83	running	44	250.00	44		249.00	51.00	21,504.0
92	running	45	250.00	45		249.00	102.00	13,372.0
41	running	45	250.00	45		249.00	102.00	14,238.0
W2H3	poweredOff	45	250.00	45		249.00	102.00	20,480.0

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## NetApp VM Insight

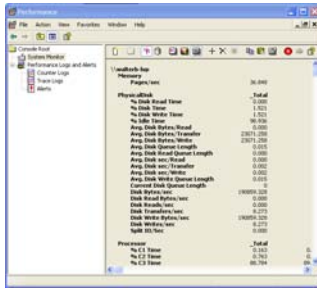
So, you can discover exactly what your customer's IO requirements are with Capacity Planner, but there are other tools that you can use to gather this information. NetApp's VM Insight from the recently acquired Onaro builds on SANscreen's service path awareness and change management technologies. It is the first cross domain application for both VM system administrators and storage administrators that spans both virtual server and traditional server deployments. It is agentless and thus very quick to deploy and provides real-time service-level information on:

- CPU,
- Memory,
- I/O,
- Network bandwidth,
- SAN bandwidth and
- Array performance.

Using this data, you can quickly develop a complete datacenter view of both physical and virtual loads impacting the storage environment and then use this information to architect the optimum VMware on NetApp solution.



## Windows Perfmon




- Collects and graphs counters for local or remote Windows server
- Not easy to collect data for many servers
- Defines the counters and objects used by most other tools.
- Use Perfmon to see what counters are available for other tools

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Windows Perfmon

As an alternative to either Capacity Planner or VM Insight, you can run Perfmon jobs on each of the Windows servers. But for environments with a large number of servers, this could be a labor intensive exercise. As a compromise, you don't have to gather information from every single server in a customer's environment. Rather, you can group servers with similar IO requirements, and then take a sampling of each group using a tool like perfmon. This way you might only have to touch 10-20% of the customer's servers to gather enough data to make informed sizing recommendations.





## Windows Perfmon

- Disk Read Bytes/sec                      Throughput
- Disk Write Bytes/sec                      Throughput
- Disk Reads/sec                              IOPS
- Disk Writes/sec                              IOPS
- Current Disk Queue Length              Bottleneck?
  - Can also look at read vs. write
- Avg. Disk sec/Transfer                      I/O Size
- Avg. Disk Bytes/Transfer                  I/O Size
- % Idle time                                    Busy?
- Split IO/Sec                                  Fragmentation?

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Windows Perfmon

This slide shows a number of Perfmon parameters that would be appropriate for storage-sizing analysis—essentially read and write throughput, IOPS, and actual transfer rates.



## Counter Instances

- Perfmon counters exist for individual drives, all (\*) or total (\_total)
- Individual drives are good for looking at the I/O of a specific app
  - Distinguish database I/O (random) vs. log I/O (sequential)
- Most of the time, \_total is what you want
- Note: All and \_total are not the same

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Counter Instances

When you use Perfmon, you can select the counters that you want to monitor and the instance or instances of that statistic that you want to collect in the Add Counters dialog box. In addition, you can add a collector for a particular drive or drives. The \_Total option aggregates a collector for all instances. Notice that the All and \_Total options are not the same thing. If you only need to monitor the aggregate disk IOPS, for example, and you don't care about IOPS to each individual drive, just choose \_Total in the instances list. By contrast, you can select the All option to monitor the statistic for all instances (drives) individually.






## NetApp Perfstat

- NetApp tool downloadable from NOW Toolchest
- Designed to collect performance data from NetApp storage systems, but can also collect from host
- Linux/Unix
  - Script
  - Multiple local or remote hosts using –h option
  - Requires iostat command from sysstat RPM
- Windows—executable, local host only
- Not the best tool to use
  - Doesn't work remotely with Windows
  - Output not in a format easily imported or analyzed

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## NetApp Perfstat

NetApp's Perfstat is designed to collect performance data from NetApp storage systems, and can collect data from host systems. However, it is not the best tool to use for sizing VMware storage because it doesn't work remotely with Windows and its output format is not easily imported or analyzed.



## Other Third-Party Tools

- PlateSpin
- MRTG
- TekTools

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

### Other Third-Party Tools

Of the other available options, PlateSpin seems to be popular. Customers engage PlateSpin (soon to be acquired by Novell) to do a consulting engagement in which PlateSpin does all of the sizing data discovery and collection and then recommends a VMware architecture. It is likely that PlateSpin's data collection tools can be configured to collect the IO information that you need for sizing NetApp storage, but we are still gaining experience in the field about how best to use these tools.



## Once You Collect Data

- Combine results from CSV files
  - Get aggregate peaks for each counter
  - Also look for simultaneous high values in more than one counter
  - Sometimes useful to graph data and look at peaks and patterns
- Feed aggregate peak values into Custom Application Sizer

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Once You Collect Data

Now you have gathered the necessary capacity, performance, and connectivity data. For example, assume that you run a custom Capacity Planner job for a month for one of your large customers and the capacity planner server generates tens of thousands of text files. You need to aggregate that information into something that is usable. In particular, you could write a perl script add up all of that real-time utilization information to find the average and the peak utilization for all of the servers. Those are the two biggest things that you want to look for.

Like any other application, you are sizing for peak load. So, there are things you can do with scheduling to change peak activity. In a lot of customer environments the peak activity is going to come during backups. If the peak IO activity in any given 24 hour period is always the backup, and you know you are going recommend that the customer backup the virtual machines in a different way than in their current environment, you are probably looking at a lower peak in the virtualized environment.

Once you get all of these values together, feed them into the Custom Application Sizer and that is going to tell you the storage hardware, disks and controllers to include in your VMware on NetApp solution.



## Capacity Planner Output

How big is the data?

- Using CP, a single counter is 100-160 bytes
- #counters x #servers x #intervals x ~130 bytes
- 2 days @5 minute intervals =  $2 \times 24 \times 12 = 576$
- 576 intervals x 6 counters x 140 servers x 130 = 62.9MB !!
- 30 days = 943MB !!!

■ How to send it

- Zip and email (text compresses nicely, 7-10:1)

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Capacity Planner Output

So if you are going to gather Capacity Planner data for a month, how much should you expect to collect? In this example, collecting data from 140 servers for one month generated almost a gigabyte of text. Zipped up, big files can be uploaded to the customer sizer.



## Storage Capacity Planning

- Work from required usable for each VM to physical disks needed
  - Actual disk on physical server is in CP report
  - Customer must say how much they actually *need*
- .vswp = size of memory for each VM
  - Technically, it's the limit minus reservation
  - Add up memory of all VMs
  - Best practice for .vswp to not reside on same datastore as VM
    - Storage Snapshot growth implication

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

## Storage Capacity Planning

For capacity, you need to add up total required space for every virtual machine. Keep in mind that if you are going with VMFS, growing is difficult. Generally customers want to size their virtual machines with some amount of growth in mind.

Additionally, you have to allocate for the .vswp file. The .vswp file is equal to the amount of memory allocated to the virtual machines. So add up the memory allocated to all of the virtual machines and that is the total .vswp requirement for the whole environment. And, in a large environment, it could be a significantly large amount of data. It is best not to store the swap file on the same data store as the VM since that would have significant impact on space needed for Snapshot copies. If you are going to keep 7 days of Snapshot copies on disk, for example, all of that Snapshot space it is going to be significant.





## Storage Capacity Planning (cont)

- VMware snapshot copies?
  - How many? How long kept? How much will they grow?
  - If only for backup scripts, should be short-lived
  - Do they use snapshots with memory?
- Datastore overhead
  - Typically <1GB for metadata overhead
  - How much free space? 20%?
- Once you know what VMware needs, standard NetApp sizing calculations apply
  - Right-size, RAID, WAFL, Snap Reserve, etc.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Storage Capacity Planning (cont)

Other questions you need to ask during storage capacity planning:

Are you going to be taking **VMware snapshot copies**? If so, how many? Are you going to be keeping them around? Are they going to be dumping memory? NetApp Snapshot doesn't use the memory dump, but if they are dumping memory to disk when they take their VMware snapshot copies, you are going to have to account for all of that space. Is any of this going to be captured in a NetApp Snapshot and retained for days and weeks?

**Datastore overhead.** We saw this in the iSCSI datastore lab. We created a 1GB LUN and we ended up with about 150 MB usable. There is about 6 or 7 hundred megabytes of meta data associated with every data store. In most cases this is not an issue. But, if your customers are creating hundreds or thousands of data stores, that number will eventually add up. And, of course, no one wants to run their data stores completely free. What is your target utilization number on deployment? Do you want to make your data stores 60, 70, 80% free? You need to take all of this into account. Then take all of the NetApp stuff, WAFL, right-sizing, Snap Reserve, Snapshot space, RAID-DP, and so on, into consideration before arriving at your final recommendations.



## Planning and Designing Connectivity

- Redundancy, throughput
- VLAN tagging or separate ports
  - VIFs
  - VLANs can allow one NIC or VIF to support multiple requirements, if adequate bandwidth
- Boot from SAN
  - HW iSCSI or FCP
  - Not as useful with ESX as with other OS
  - Does nothing for ESX 3i

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Planning and Designing Connectivity

Part 1 of this course covers connectivity, so we don't need to go into as much depth here. The big connectivity planning and design question is what kind of redundancy do you need? Every customer wants some level of link redundancy within their ESX servers because a link failure on an ESX server will cause all of the virtual machines to go off line with no automatic resolution of that condition. That is a very compelling reason to make sure you have link redundancy on both sides.

And, the second question, assuming you have established redundancy, is what are your throughput requirements? How much IO do you need to be able to drive in and out of the ESX servers? When you add up the numbers on a per ESX server basis, IO requirements are probably not going to be very high. Average utilization on most ESX servers is probably 20-30 megs per seconds. This is largely because the most likely server to be virtualized is a lightly used one. Actual usage in your customer environment might vary. If they are virtualizing Exchange, they are going to have much higher IO requirements than if they are virtualizing a domain controller.

And boot from SAN is not as useful with ESX as it is with tradition network operating systems because an individual ESX server in a data center has very little value. Its only value is that at any given time it is running some of the virtual machines in the environment. But, its individual configuration really means nothing. In fact, it takes about 20 minutes to build an ESX server. You only have to answer a couple of configuration questions. Some customers are interested in boot from SAN from a resiliency perspective. But, boot from SAN doesn't provide any additional resiliency because with ESX you can recreate a server from scratch just as fast as you can restore a boot LUN in the event of a problem. In 3i, a boot from SAN option doesn't even exist.




## The Future of Sizing

- Standardize the collection process and tools
- Streamline upload
- Automate the analysis and sizing process
  - Single web site/page/tool that
    - Imports sizing data
    - Runs scripts to calculate I/O requirements
    - Builds graphs
    - Accepts model, disk, etc. constraints
    - Generates config
- In early stages – manual for now

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## The Future of Sizing

We are working on a VMware specific sizing tool. We are working on interacting with these capacity planning and measuring tools in a better way, making it easier for you guys to gather the information. I don't know what exactly the future of that is going to be. We might end up partnering with one of the other vendors, like Platespin—someone that would be interested in meeting us half way on that stuff—where they would build more disk specific information into their tool by default and that is what we would use. You would hire a professional services engagement from NetApp and they would come in with that tool and measure the customer's environment. Those kind of things are being investigated right now. But, in the meantime, engage your CSEs, get help from us, and we will do all we can to help you.



The last module of the course covers Performance monitoring.



## Learning Objectives

- Performance monitoring using
  - VirtualCenter,
  - VMware VI Client, and
  - esxtop

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

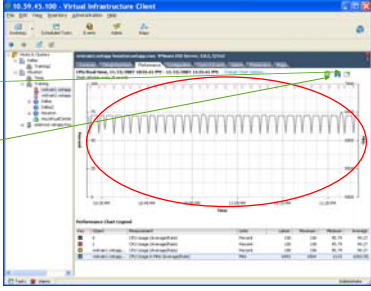
## Learning Objectives

When you have completed this module, you will have a basic understanding of VMware on NetApp Performance monitoring using VirtualCenter and the VMware esxtop command.

**Virtual Infrastructure Client**

NetApp

- ESX Server provides real-time performance reporting from Virtual Infrastructure Client.
  - Click the Chart icon to generate a custom chart.
  - Click the Save icon to export data in CSV format that you can import into Excel.



The screenshot shows the Virtual Infrastructure Client interface. A red oval highlights a performance chart with a grid and data points. A green arrow points from the 'Chart' icon in the text to the chart area. Another green arrow points from the 'Save' icon in the text to a save icon in the interface. Below the chart is a table with columns for 'Performance Chart Legend', 'Host', 'Metric', 'Unit', 'Min', 'Max', 'Avg', and 'StdDev'. The table contains several rows of data for different hosts and metrics.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Virtual Infrastructure Client

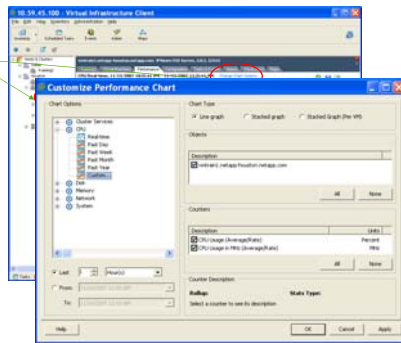
After deployment of a VMware environment, it is still necessary for your customer to be able to track performance so they can keep their constituents happy and can plan for the future. You can get real-time ESX server and virtual machine utilization and performance of the using the VI client, as shown in this screenshot, but this data is not logged. You can however, export the performance data as a printed chart or into a CSV file that can then be imported into Excel.



## Virtual Infrastructure Client

- To select statistics, click the desired entity and click Change Chart Options.
- Available Statistics depend on selected entity

	Cluster	Server	VM	Res Pool
Cluster Services	✓	✓		
CPU	✓	✓	✓	✓
Disk		✓	✓	✓
Memory	✓	✓	✓	✓
Network		✓	✓	
System		✓	✓	



## Virtual Infrastructure Client

You can customize the information included in the chart or CSV file. Cluster services, CPU, disk, memory, network, and system utilization and performance statistics are available broken out by cluster, server, virtual machine, and resource pool.



## VirtualCenter

- VirtualCenter provides historical and real-time performance data collection and reporting.
- Historical data stored in database
  - Built-in MSDE (not recommended for production environments)
  - SQL
  - Oracle
- Fewer metrics available than with real-time performance display

Category		Metrics	
Cluster Services	→	CPU Fairness	Memory Fairness
CPU	→	CPU Usage %	CPU Usage MHz
Disk	→	Disk Usage KBps	
Memory	→	Memory Usage %	
Network	→	Network Usage KBps	
System	→	Uptime (Days)	

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## VirtualCenter

The Performance tab in VirtualCenter gives you both a historical and a real-time picture of server and virtual machine utilization and performance. This data that is being logged into a database.





## VirtualCenter

- VirtualCenter and ESX Server Performance Tabs display real-time performance statistics as shown below.

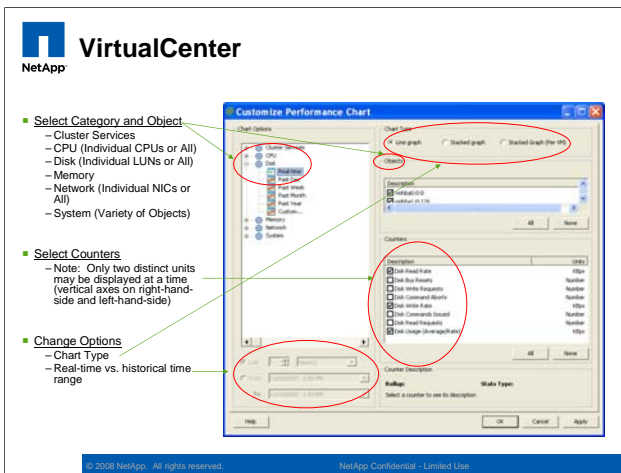
		Real-Time Metrics											
		CPU Fairness						Memory Fairness					
		System Overall						Resource CPU Usage					
Cluster	→												
System	→												
CPU	→	CPU Usage %	CPU Usage MB/s	CPU Used MB	CPU Usage %	CPU Usage MB/s	CPU Reserved Capacity MB/s	Idle %					
Disk	→	Disk Usage MB/s	Disk Read Rate MB/s	Disk Bus Streams	Disk Write Requests	Disk Command Alerts	Disk Write Rate MB/s	Disk Command/Issued	Disk Read Requests	Disk Usage MB/s			
Network	→	Network Usage MB/s	Network Packets Received	Network Data Received	Network Data Rate MB/s	Network Usage MB/s	Network Packets Transmitted	Network Data Transmitted	Network Usage MB/s				
Memory	→	Memory Shared Common KB	Memory Granted KB	Memory Consumed KB	Memory Balloon KB	Memory Shared KB	Memory Usage %	Memory Swap In KB	Memory Active KB	Memory Zaps KB			
	→	Memory Heap KB	Memory State KB	Memory Unreserved KB	Memory Reserved Capacity KB	Memory Overhead KB	Memory Used by VMs KB	Memory Swap Out KB	Memory Heap Free KB	Memory Swap Used KB			

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use


## VirtualCenter

VirtualCenter and ESX Server Performance Tabs display real-time performance statistics as shown in this slide.



## VirtualCenter

To customize a chart, first select the resource categories and specific objects for which you want to see performance and utilization data. You can see statistics on cluster services, CPU, disk, memory, network, and on a variety of system objects. Next, select the metrics that you want to see from the Counters check list. Finally, select a chart type (line graph, stacked graph, or stacked graph per VM), and a time frame—the last X hours or a specific date range.



## esxtop Command

- Provides granular real-time resource utilization data
- Runs on the ESX Server Service Console.
- Viewing modes:
  - Interactive (default)
  - Batch
  - Replay
- Displays in four panels:
  - CPU
  - Memory
  - Storage
  - Network
- Can only be run by user root.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 8

## esxtop

In addition, you can use the esxtop command from the Service Console and look at utilization in real time. It provides granular real-time resource utilization data in four panels: the CPU panel, the memory panel, the storage panel, and the network panel. Can only be run by user root. When run from the Service Console command line, it displays information in interactive mode, but it also can be run as a batch process, and in replay mode to replay resource utilization statistics collected using vm-support.





## esxstop Memory Panel

- Memory utilization statistics by:
  - Server
  - Worlds consuming memory
  - Resource pool
  - Running virtual machine

```

root@sanak01:~#
7:25:11pm up 44 min, 30 wor:det MEM overcommit average: 0,00, 0,00, 0,00
MEM (MB): 15215 total: 272   cap, 202   vml, 278 other, 15465 free
VMEM (MB): 15620 managed: 509 minfree, 1016   read, 14947 unred, high state
CPOSEM (MB): 72   free: 2947 mem_s, 2047 mem_f: 0,00   r/s, 0,00 w/s
NMR (MB): 3615 (3265), 4906 (4900), 4096 (4090), 4021 (4015)
PMRE (MB): 20 shared, 8   commit, 12   saving
CASP (MB): 0   curr, 0   target: 0,00   r/s, 0,00 w/s
MEMCTL (MB): 0   curr, 0   target, 0   max

```

ID	UID	NAME	MEM	MEM%	SIZE	ICED	ZIK IV	ZK IVS	ZK IV	ZK IVN	UNDEL	QMS	QMS%
12	12	Umem-cslauthd	1	2,20	2,00	0,74	0	0	0	0	0	0,00	0,00
15	15	PH4L4HS	5	256,00	317,87	17,92	5	2	4	7	16,91	28,96	72,98

## esxstop Memory Panel

The esxstop memory panel, displayed by pressing the m key in interactive mode, lists memory utilization statistics by server, worlds consuming memory, resource pool, and running virtual machine. The first line of the panel shows the current time, time since last reboot, number of currently running worlds, and memory over commitment averages.



**esxstop Storage Panel**  
**Interactive Commands**

- Key commands to toggle between expanded or unexpanded statistics:
  - e Toggles display of storage adapter statistics
  - a Toggles display of storage channel statistics
  - t Toggles display of storage target statistics
  - l Toggles display of LUN statistics
- Key commands to sort statistics:
  - r Sorts by Reads column
  - w Sorts by Writes column
  - R Sorts by MB read column
  - T Sorts by MB written column
  - N Default sort order—Sorts first by ADAPTR column, then by CID column within each ADAPTR, then by TID column within each CID, then by LID column within each TID, and by WID column within each LID.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 10

## esxstop Storage Panel Interactive Commands

In addition to the commands that display the different statistics panels, there are eight more key commands that change the way data is listed in the Storage Panel:

By default, storage statistics are displayed in an unexpanded mode. If you need to see more granularity:

Pressing the “e” key toggles an expanded set of storage resource utilization statistics broken down by individual channels belonging to a particular storage adapter. You are prompted for the adapter name.

The “a” key command toggles expansion of storage resource utilization statistics by individual targets belonging to a particular storage channel. You are prompted for adapter name and channel ID. You have to expand the adapter first before you can expand the channel.

Pressing “t” expands or hides storage resource utilization statistics aggregated by individual LUN and by storage target. You are asked to supply the adapter name, channel ID, and target ID. Expand channel, and adapter before expanding the target.

The “l” command toggles expansion of utilization statistics by individual worlds utilizing a particular LUN. You have to supply the adapter name, the channel ID, the target ID, and the LUN ID. Expand the LUN target, channel, and adapter before expanding the LUN.

By default, the Storage Panel sorts first by ADAPTR column, then by CID column within each ADAPTR, then by TID column within each CID, then by LID column within each TID, and by WID column within each LID. If you want to see a different order, you can use the lower case R, lower case W, upper case R, and upper case T key commands sort statistics by the Reads, Writes, MB read, or MB written columns, respectively. The N key returns to the default sort order.



## esxtop Storage Panel Statistics

- **ADAPTR**—Name of the storage adapter.
- **CID**—Storage adapter channel ID. Visible only if adapter is expanded.
- **TID**—Storage adapter channel target ID. Visible only if the corresponding adapter and channel are expanded.
- **LID**—Storage adapter channel target LUN ID. Visible only if adapter, channel and target are expanded.
- **WID**—Storage adapter channel target LUN world ID. Visible only if the corresponding adapter, channel, target and LUN are expanded.
- **NCHNS**—Number of channels.

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

13

## esxtop Storage Panel Statistics

The various esxtop panels display a plethora of resource utilization statistical information. This screen and following three screens describe the statistics summarized on the esxtop Storage Panel. Refer to the Resource Management Guide for the ESX Server and VirtualCenter for more details.





## esxtop Storage Panel Statistics (continued)

- **NTGTS**—Number of targets
- **NLUNS**—Number of LUNs
- **NVMS**—Number of worlds
- **SHARES**—Number of shares
- **BLKS**—Block size in bytes—LUNs only
- **AQLEN**—Storage adapter queue depth—maximum number of ESX Server VMKernel active commands supported by adapter
- **LQLEN**—LUN queue depth—maximum number of ESX Server VMKernel active commands LUN can handle

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## esxtop Storage Panel Statistics (continued)



## esxtop Storage Panel Statistics (continued)

- **WQLEN**—World queue depth--maximum number of ESX Server VMKernel active commands a world supports per LUN
- **%USD**—Percentage of queue depth (adapter, LUN or world) used by ESX Server VMKernel active commands
- **LOAD**—Ratio of ESX Server VMKernel active commands plus ESX Server VMKernel queued commands to queue depth (adapter, LUN or world).
- **ACTV**—Number of currently active ESX Server VMKernel commands
- **QUED**—Number of currently queued ESX Server VMKernel commands
- **CMDS/s**—Commands issued per second

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

15

## esxtop Storage Panel Statistics (continued)



## esxtop Storage Panel Statistics (continued)

- **READS/s**—Read commands issued per second.
- **WRITES/s**—Write commands issued per second.
- **MBREAD/s**—Megabytes read per second.
- **MBWRTN/s**—Megabytes written per second.
- **DAVG/cmd**—Average device latency per command, in milliseconds.
- **KAVG/cmd**—Average ESX Server VMKernel latency per command, in milliseconds.
- **GAVG/cmd**—Average virtual machine operating system latency per command, in milliseconds.
- **ABRTS/s**—Commands aborted per second.
- **RESETS/s**—Commands reset per second.

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## esxtop Storage Panel Statistics (continued)



## esxstop Network Panel

- Network utilization by
  - Port
  - Configured network or virtual network device
- Refer to the row for the port to which the physical network adapter or virtual network adapter is connected.

```
root@sanas001 ~
7:26:43pm up 45 min, 39 users, CPU load average: 0.11, 0.05, 0.09
```

PORT ID	UPLINK	USED BY	DTYP	DNAM	PkT/TX/s	MB/TX/s	PkT/RX/s	MB/RX/s	100PTX	200PRX
1677217	Y	01NCP	H	pps	0,00	0,00	0,00	0,00	0,00	0,00
1677218	N	01NCP	S	vSwitch0	4,15	0,01	2,57	0,03	0,00	0,00
1677219	N	0cvsaf0	S	vSwitch0	0,00	0,00	0,00	0,00	0,00	0,00
1677220	N	0cvsaf0	S	vSwitch0	4,15	0,01	0,53	0,00	0,00	0,00
33554433	Y		S	vSwitch1	10270,79	5,02	31536,92	341,39	0,00	0,00
33554434	N	01NCP	S	vSwitch1	0,00	0,00	0,00	0,00	0,00	0,00
33554435	N	1069:RHEL-RIS	S	vSwitch1	10270,89	5,01	31536,25	342,39	100,00	0,00
50331649	Y		S	vSwitch2	0,00	0,00	0,53	0,00	0,00	0,00
50331650	N	01NCP	S	vSwitch2	0,00	0,00	0,00	0,00	0,00	0,00

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use 17

## esxstop Network Panel

The esxstop network panel, displayed by pressing the n key in interactive mode displays network utilization statistics arranged by port per configured network and virtual network device. Find to the row for the port to which the physical network adapter or virtual network adapter is connected for statistics on that adapter. Refer to the Resource Management Guide for the ESX Server and VirtualCenter for more details on the wealth of network utilization statistics found in this panel including packets transmitted per second, packets received per second, and . MegaBits transmitted and received per second.



## esxtop Network Panel Sort Order Keys

- Key commands to sort statistics:

- T Sorts by Mb Tx column.
- R Sorts by Mb Rx column.
- t Sorts by Packets Tx column.
- r Sorts by Packets Rx column.
- N Default sort order—sorts by PORT ID column

© 2008 NetApp. All rights reserved.

NetApp Confidential - Limited Use

18

### esxtop Network Panel Sort Order Keys

By default, the Network Panel sorts first by Port ID column. If you want to see a different order, you can use the upper case T, upper case R, lower case T, and lower case r key commands sort statistics by the Mb transmitted, Mb received, packets transmitted, and packets received columns, respectively. The upper case N key returns to the default Port ID sort order.



## Learning Objectives

- Performance monitoring using VirtualCenter, the VMware VI Client, and esxtop

© 2008 NetApp. All rights reserved. NetApp Confidential - Limited Use

## Learning Objectives

Now that you have completed this module, you should have a basic understanding of VMware on NetApp Performance monitoring using VirtualCenter and the VMware esxtop command.



**NetApp™**  
Go further, faster

NETAPP UNIVERSITY

# VMware on NetApp Solutions Lab Manual

Revision Number: 1.0

Course Number: SALES-ED-ILT-VMNSOLLAB

This page is intentionally left blank.



## ATTENTION

The information contained in this guide is intended for training use only. This guide contains information and activities that, while beneficial for the purposes of training in a closed, non-production environment, can result in downtime or other severe consequences and therefore are not intended as a reference guide. This guide is not a technical reference and should not, under any circumstances, be used in production environments. To obtain reference materials, please refer to the NetApp product documentation located at [www.now.com](http://www.now.com) for product information.

## COPYRIGHT

Copyright © 2008 NetApp. All rights reserved. Printed in the U.S.A. Specifications subject to change without notice.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products or materials described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product or materials does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

## RESTRICTED RIGHTS LEGEND (SUBTITLE FORMATTING)

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## TRADEMARK INFORMATION (SUBTITLE FORMATTING)

NetApp, the NetApp logo, and Go further, faster, FAServer, NearStore, NetCache, WAFL, DataFabric, FilerView, SecureShare, SnapManager, SnapMirror, SnapRestore, SnapVault, Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and SpinStor are registered trademarks of Network Appliance, Inc. in the United States and other countries. Network Appliance, Data ONTAP, ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, gFiler, MultiStore, SecureAdmin, Smart SAN, SnapCache, SnapDrive, SnapMover, Snapshot, vFiler, Web Filer, SpinAV, SpinManager, SpinMirror, and SpinShot are trademarks of NetApp, Inc. in the United States and other countries.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

## TABLE OF CONTENTS

### LABS:

**LAB 1: CONNECTING DATASTORES AND CLONING VMS**

**LAB 2: ALIGNING VIRTUAL DISK FILE SYSTEMS**

**LAB 3: VIRTUAL MACHINE BACKUPS USING NETAPP SNAPSHOT**

**LAB 4: VIRTUAL MACHINE RECOVERY FROM NETAPP SNAPSHOT COPIESL**

**LAB 5: FILE LEVEL RECOVERY FROM NETAPP SNAPSHOT COPIES**

**LAB 6: VIRTUAL MACHINE REPLICATION AND DISASTER RECOVERY USING SNAPMIRROR**

**LAB 7: FAS DEDUPLICATION OF NFS AND VMWARE**

**LAB 8: THIN PROVISIONING & FAS DEDUPLICATION OF VMFS DATASTORES**

**LAB 9: PROVISIONING DATASTORES & VMS WITH FLEXCLONE**

**LAB 10: GROWING DATASTORES**



## LAB 1: Connecting Datastores and Cloning VM's

Last Updated: Wednesday, May 28, 2008

## Summary:

In this lab we will create and configure VMFS data stores using FCP and iSCSI, we will also connect to an existing volume and use it as a NFS data store.

## Participant Prerequisites

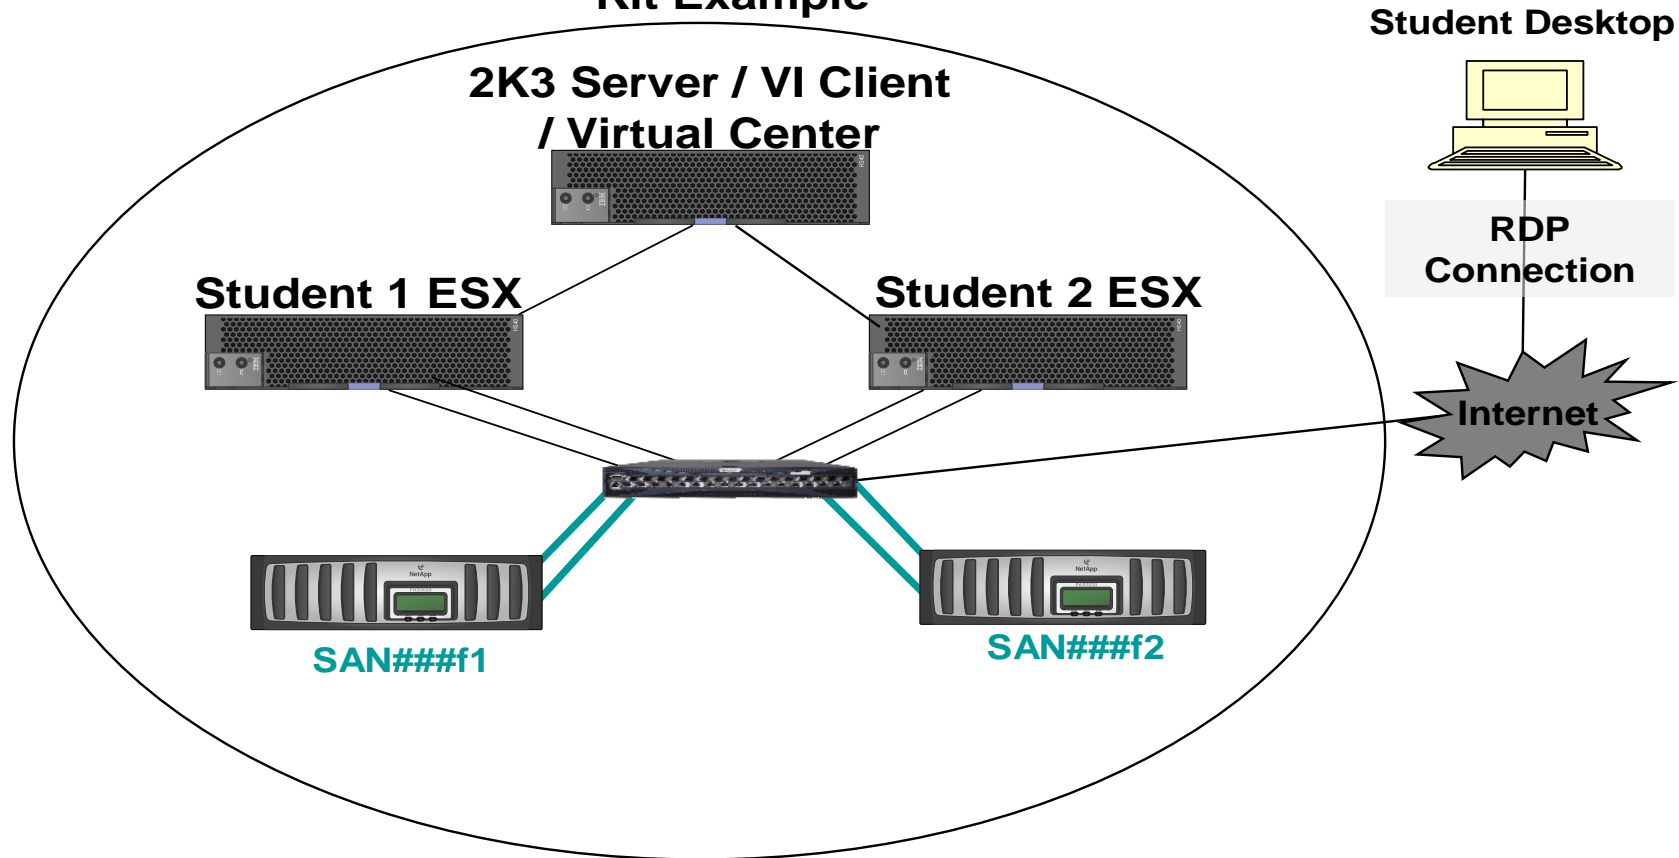
- None

## Lab Environment

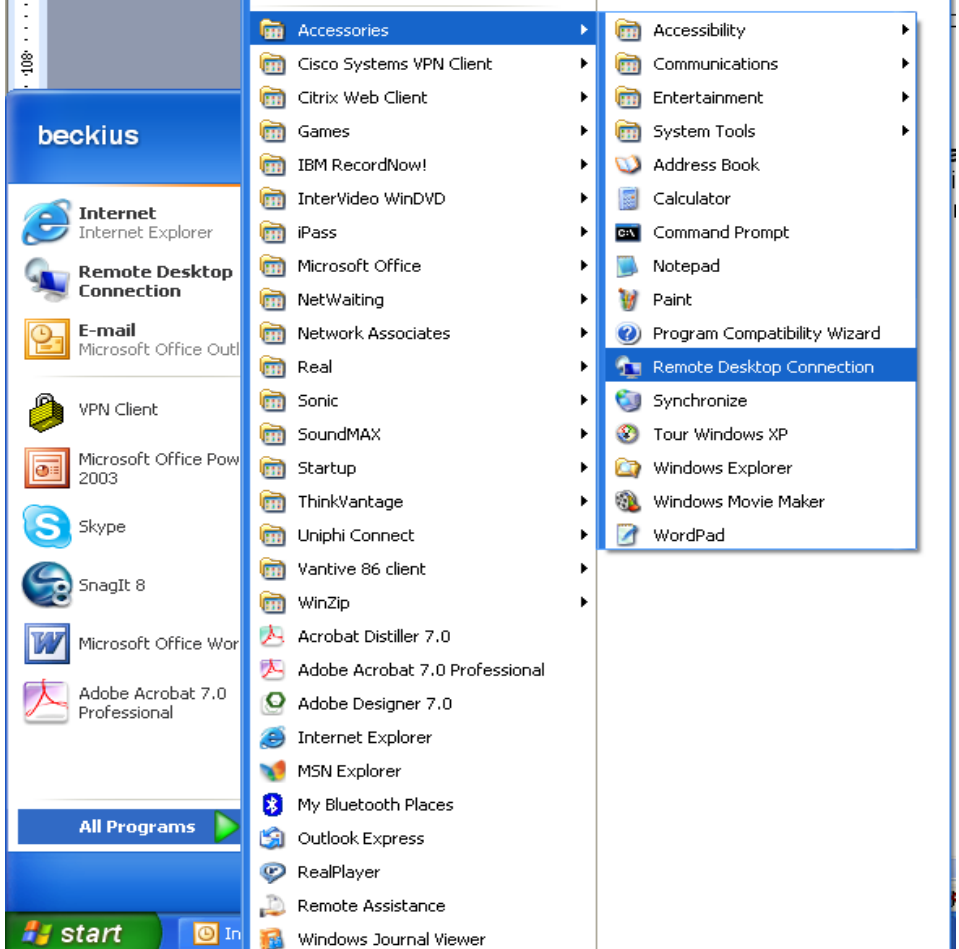
- An ESX server, VMware VI client installed on the local system, a FAS storage controller, and a Virtual Center Server. The FAS storage controller requires both FC and iSCSI connection

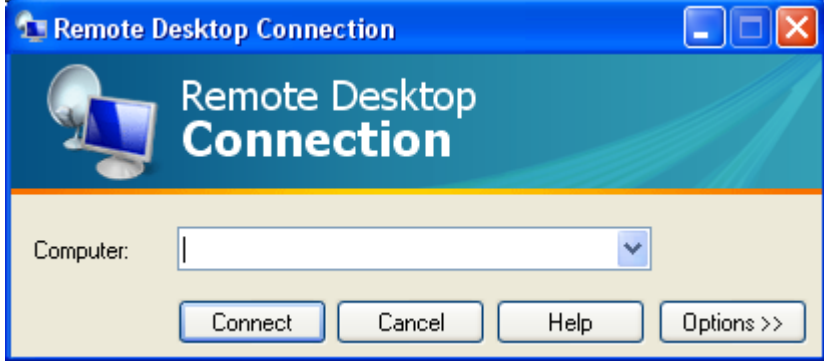
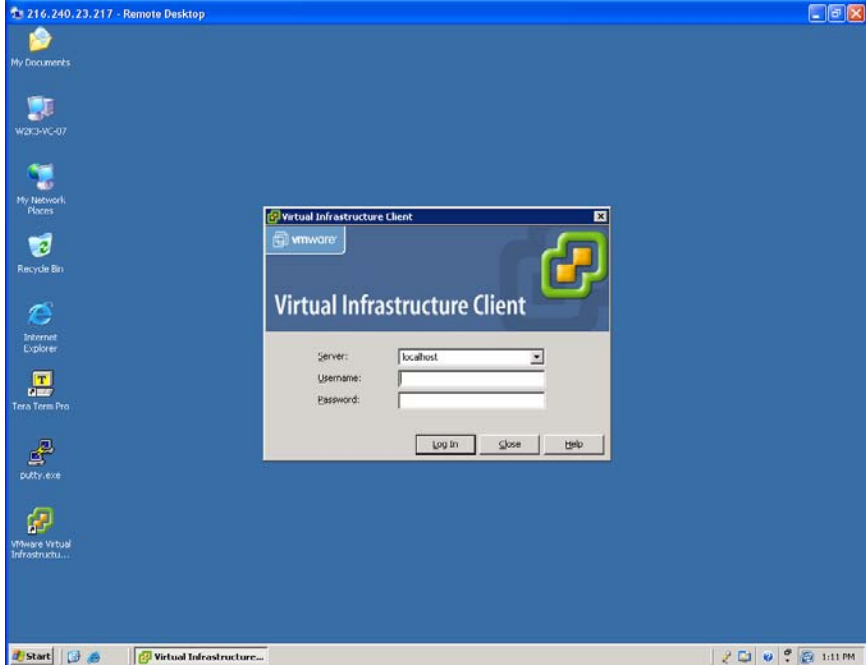
# VMware Lab Environment

## Kit Example



## Part I – Create the VMFS Datastore Using FCP.

	Click Stream	Screen Shots
1	<p>From your windows operating system open the Remote Desktop client.</p> <ul style="list-style-type: none"> <li>This icon should be on the desktop of your classroom system</li> </ul> <p>--OR--</p> <ul style="list-style-type: none"> <li>If using your personal system, this can be found by going to <b>Start &gt; All Programs &gt; Accessories</b>. The icon will either be on the base menu or in the communications sub-heading.</li> </ul>	 <p>The screenshot shows a Windows XP Start menu for a user named 'beckius'. The 'All Programs' button is highlighted, and the 'Accessories' folder is expanded. Within 'Accessories', the 'Remote Desktop Connection' icon is highlighted. Other visible items include Internet Explorer, Remote Desktop Connection, E-mail, VPN Client, Microsoft Office PowerPoint 2003, Skype, Snagit 8, Microsoft Office Word, Adobe Acrobat 7.0 Professional, Adobe Designer 7.0, Internet Explorer, MSN Explorer, My Bluetooth Places, Outlook Express, RealPlayer, Remote Assistance, and Windows Journal Viewer.</p>

<p>2</p>	<p>In the Remote Desktop Connection window enter the <b>IP address</b> of the Windows 2003 system assigned to you by the instructor.</p> <p>Enter the login credentials (provided by instructor)  <b>Login:</b> administrator  <b>Password:</b> (provided by instructor)</p>	
<p>3</p>	<p>Once on the Remote windows server, <b>double click</b> on the <b>VMware Virtual Infrastructure</b> icon to open the VI Client.</p> <p><b>Note:</b> Your VirtualCenter server is running on the Windows 2003 localhost system you RDP'ed into.</p> <p>Log in using <b>localhost</b> as the server using the same username and password as the Windows system to access the VirtualCenter.</p>	

4 Select **your** ESX server in the left hand pane. (Assigned by instructor) under the NetApp datastore

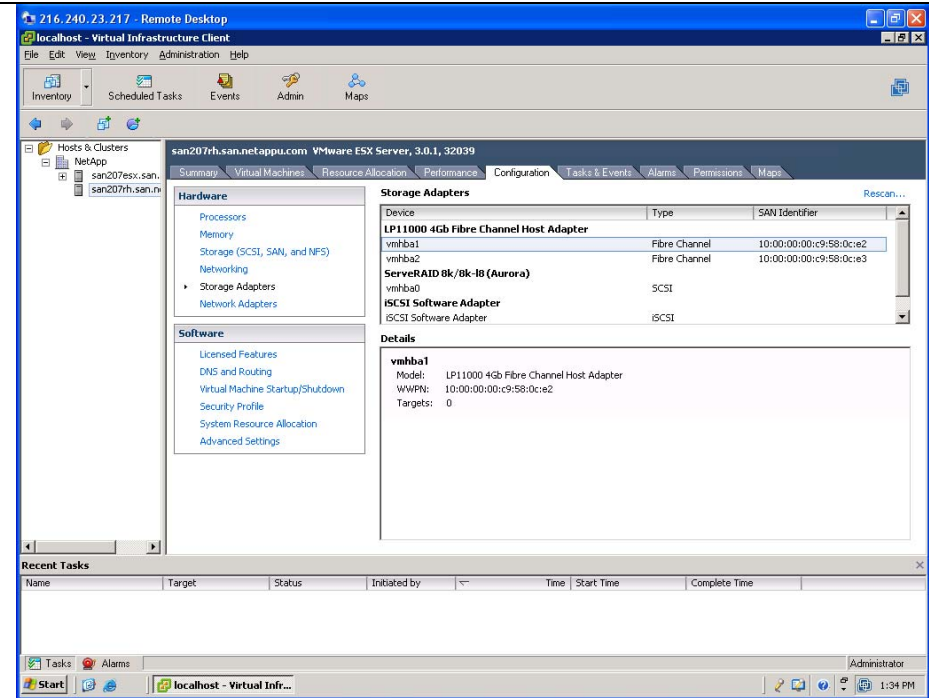
- Click on the **Configuration** tab
- In the Hardware section, click on the **Storage Adapters** link
- Note the fiber channel adapter name(s) (they start with vmhba#). You will only be using one of the adapters for this exercise.

HBA name: \_\_\_\_\_

SAN Identifier(WWPNs): \_\_\_\_\_

HBA Name: \_\_\_\_\_

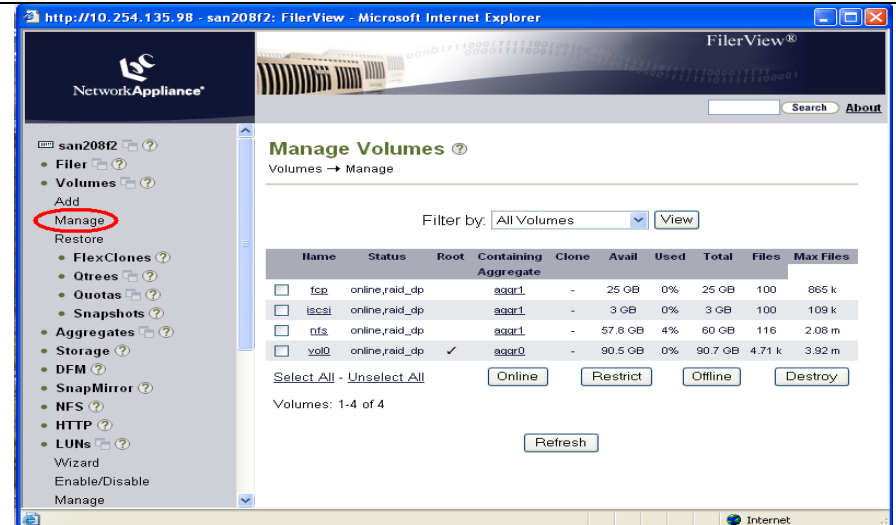
SAN Identifier(WWPNs): \_\_\_\_\_





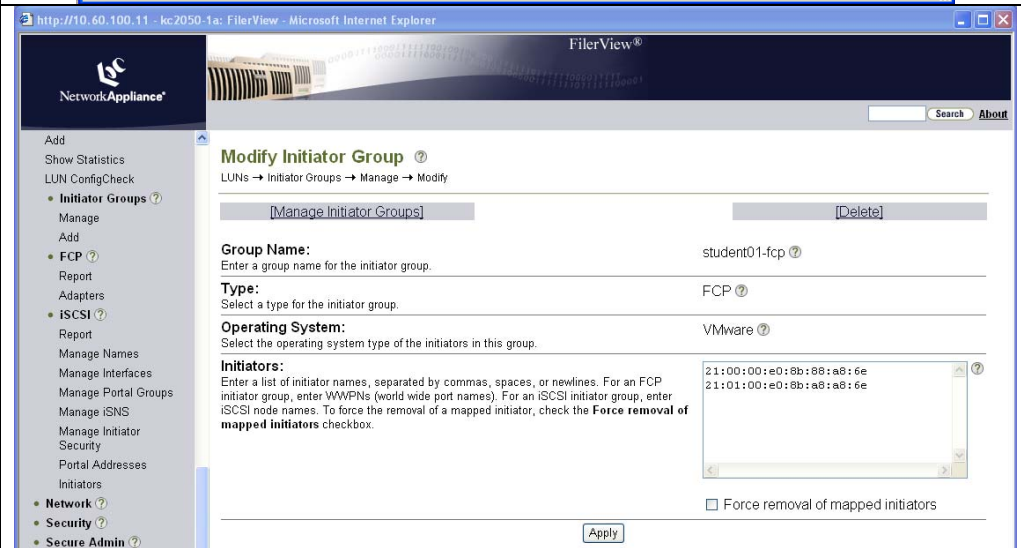
5 Connect to your assigned Storage Controller using FilerView.

- use the storage controller IP (provided by instructor) in a browser window [http://xxx.xxx.xxx.xxx/na\\_admin/](http://xxx.xxx.xxx.xxx/na_admin/)
- Within the FilerView window choose **Volumes > Manage** in the list to the right, note the volumes that are precreated.
- They should be: fcp, iscsi, nfs, and vol0



6 Now you will use one of the initiators you discovered earlier to create an FCP igroup.

- Verify that FCP service is started by going to the CLI interface for the controller and issuing a **fcp start**.
- From FilerView choose **LUNs** in the left column then select **Initiator Groups > Add** in the sub menu.
- Use the group name of **stu#-fcp**
- Verify the Type is **FCP**, and the Operating System is **VMware**
- Enter the WWPN (SAN Identifier) of **vmhba1** from your assigned ESX server into the **initiators** box and choose apply.



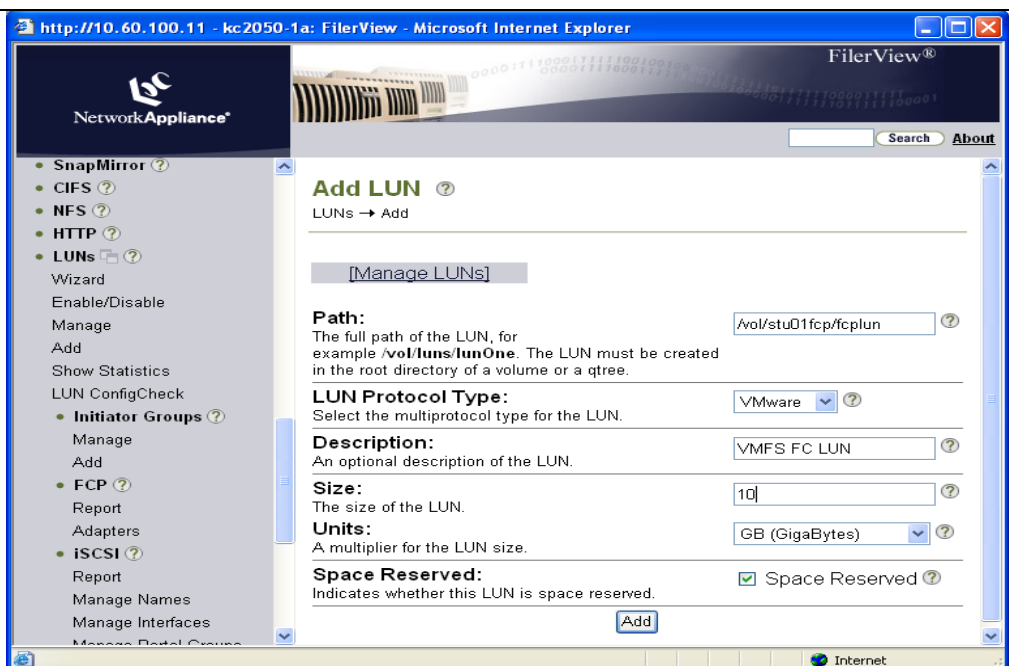
7 Use FilerView to create a 10 GB LUN in the fcp volume on your assigned storage controller . From FilerView choose: **LUNs > Add**

- LUN path: **/vol/fcp/fcplun**
- Set the LUN Protocol type to **VMware** (Since it will be a VMFS LUN).
- Enter a description of your choice, make it 10GB in size and choose Add

When your LUN has been successfully created, choose **Manage** under the LUNs heading on the left.

- Click on the **No Maps** link following your LUN to map the igroup you created previously to the LUN.
- Choose the **Add Groups to Map** header
- Select the igroup you created, and choose **Add**

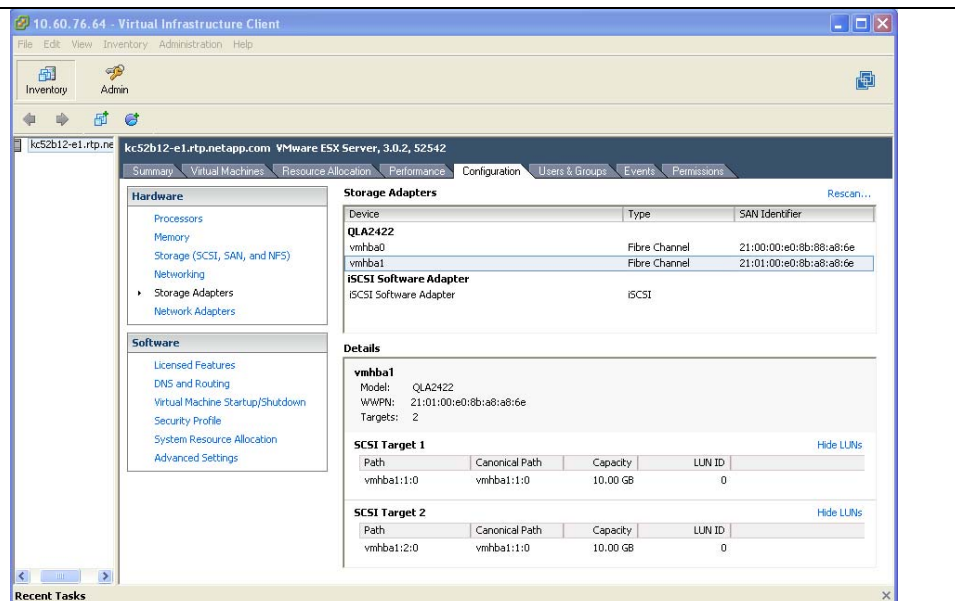
**Note:** This should be the only LUN mapped to that adapter on your ESX server so you may use LUN ID 0.



8 Use the **VI client** to rescan for the new LUN on your ESX server.

- Return to the VI client and select your ESX server.
- Click on the **configuration tab**.
- In the Hardware section, click on the **Storage Adapters** link.
- Click on the **rescan** link and wait for the ESX server to rescan its adapters. You should now be able to see the new LUN on vmhba1.

**Note:** If the new LUN is not visible, try to run the rescan a second time. Depending on the pathing configuration, you may see the LUN displayed multiple times.



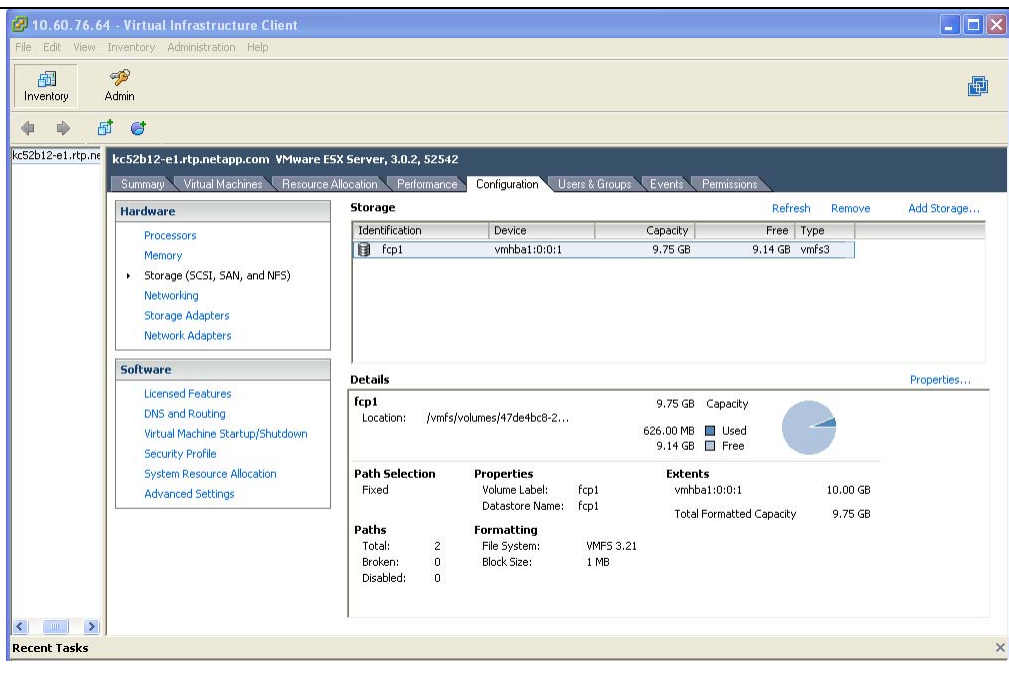
9 To create a VMFS datastore, start by clicking on **Storage** in the Hardware section of the VI client.

- Click on the **Add Storage** link which will start the storage wizard.
- On the select storage type page, select Disk/LUN and click **next**.
- On the select Disk/LUN page, choose the 10 GB LUN you just created on vmhba1 and click **next**.

**NOTE:** There may be more than one LUN listed)

- The current disk layout page will appear observe the data layout and click **next**.
- In the Disk/LUN properties page, enter **stu#fcp** (where # is your student number assigned by instructor) under datastore name and click **next**.
- On the Disk/LUN formatting page leave these values at their defaults and click **next**.

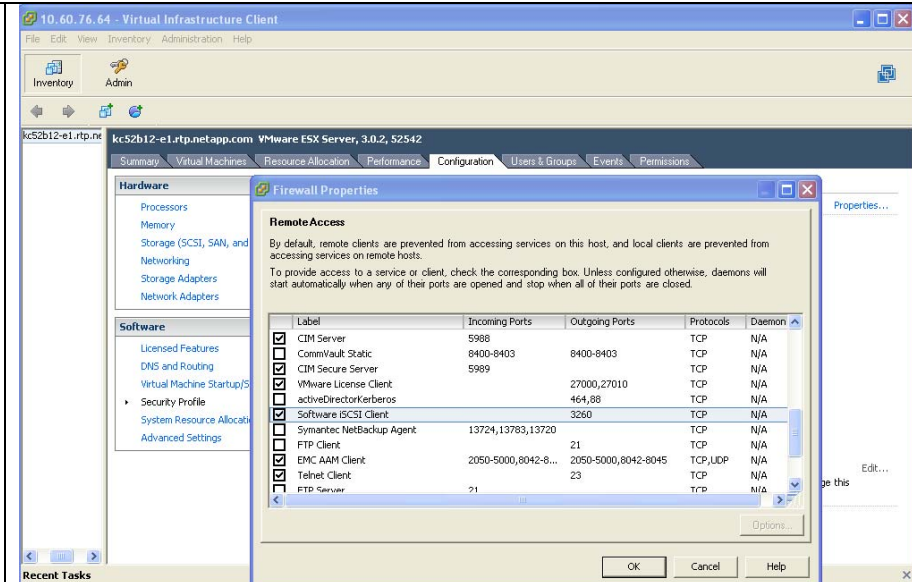
Click **finish** to complete the add storage wizard.



## Part II – Create VMFS Datastore Using iSCSI Software Initiator

The iSCSI protocol is not enabled by default in ESX 3.0X, use the procedures below to enable the initiator and open the ESX firewall for iSCSI traffic.

- 1 In the VI client select the configuration tab and then click on the **security profile** link under the software heading.
  - Click on **Properties** to open the Firewall Properties page.
  - Find the Software iSCSI Client and **check the box** to enable the protocol.
  - Click OK to close the Firewall properties window.

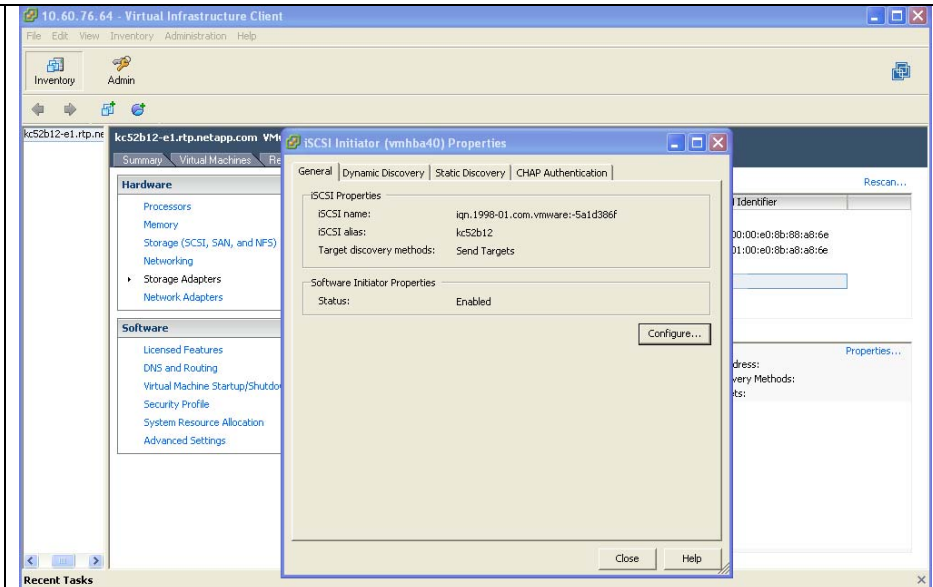


2 To enable the iSCSI initiator click on the **storage adapters** link in the hardware field.

- Select the **iSCSI Software Adapter** and click on the properties link.
- In the General tab, click on **configure** to open the general properties window and check the box to enable the initiator.
- Click **OK** to close the properties window.

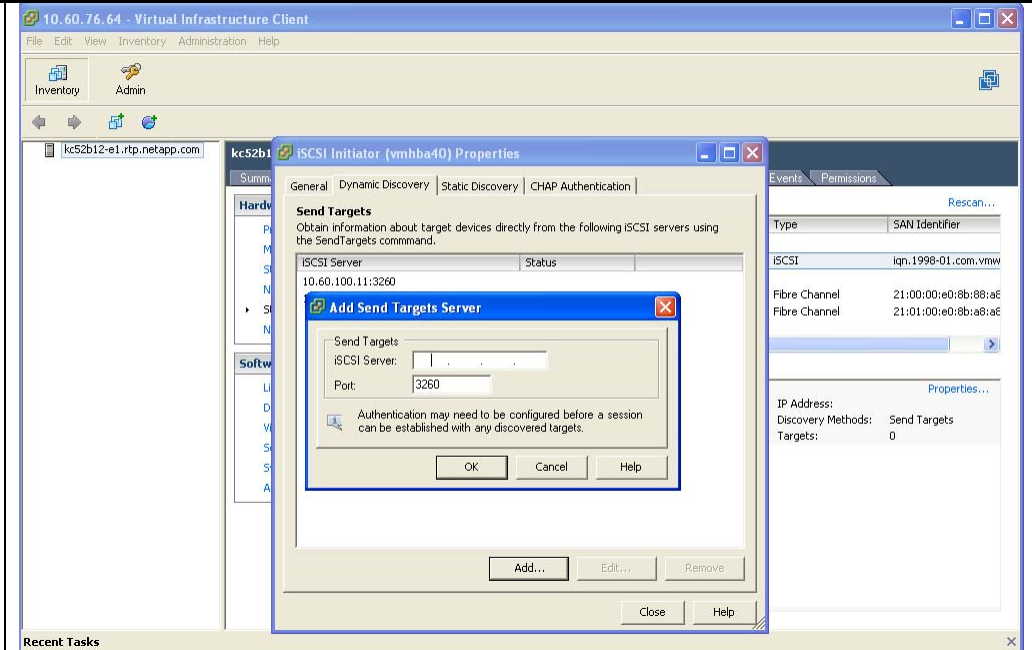
**Note:** Record the iSCSI Name/iqn of the ESX server for later use.

iSCSI name: \_\_\_\_\_

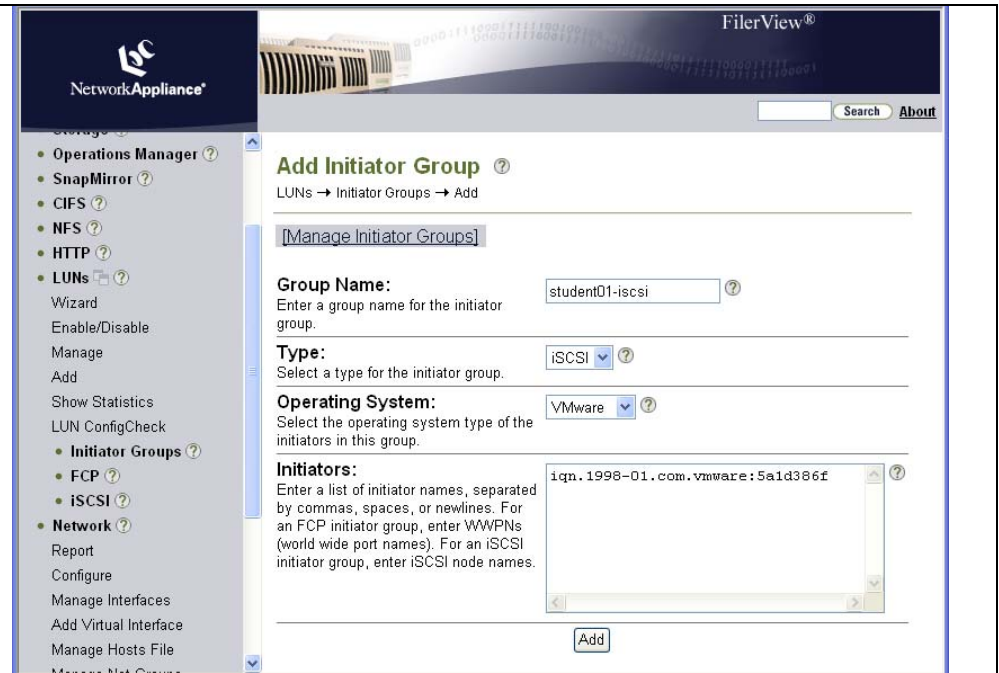


3 Use the VI client to establish an iSCSI session between your ESX servers and your storage controller.

- Open the configuration tab. In the Hardware section, click on the **Storage Adapters** link under hardware.
- Select the iSCSI Software Adapter and click on **Properties** to bring up the initiator properties window.
- Click on the Dynamic Discovery tab then click on the **Add** button.
- Enter the **Storage IP** address (provided by Instructor) for your storage controller then click on OK.  
**Note:** this may take a moment to discover



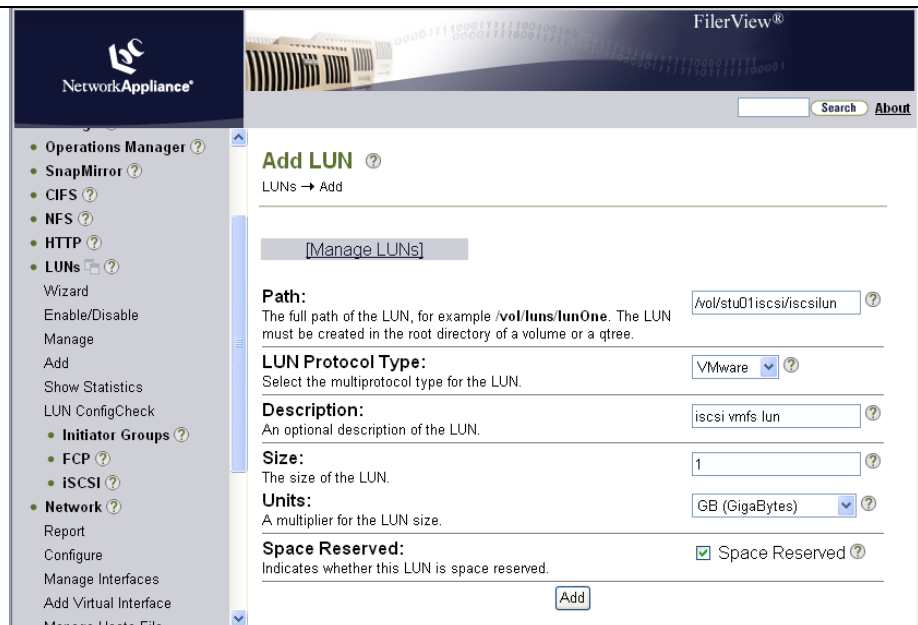
- 4 Creating a VMFS datastore using iSCSI is much like the process when using FCP, but you must establish an iSCSI session before you can scan for the LUN.
- Connect to your storage controller using FilerView, under the initiator groups heading choose **add**.
  - Create an iSCSI igroup named **stu#-iscsi**, OS type of **VMware** containing the iqn's of your ESX server from the previous step





5 Under the LUN heading click add and create a 1 GB LUN in the iSCSI volume (**/vol/iscsi/iscsilun**).

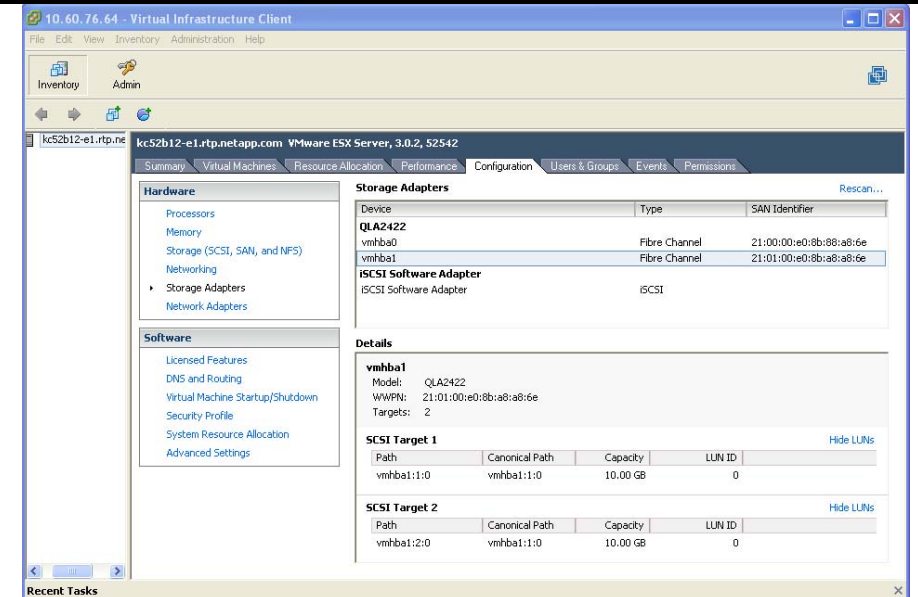
- This will be a VMFS LUN so set the LUN Protocol Type to **VMware**. Click **Add** to create the LUN
  - Choose **Manage LUNs** to map the LUN to the igroup you created in step 1.
- Note:** This should be the only LUN mapped to that adapter so you may use LUN ID 0.



6 Use the **VI client** to rescan for the new LUN on your ESX server.

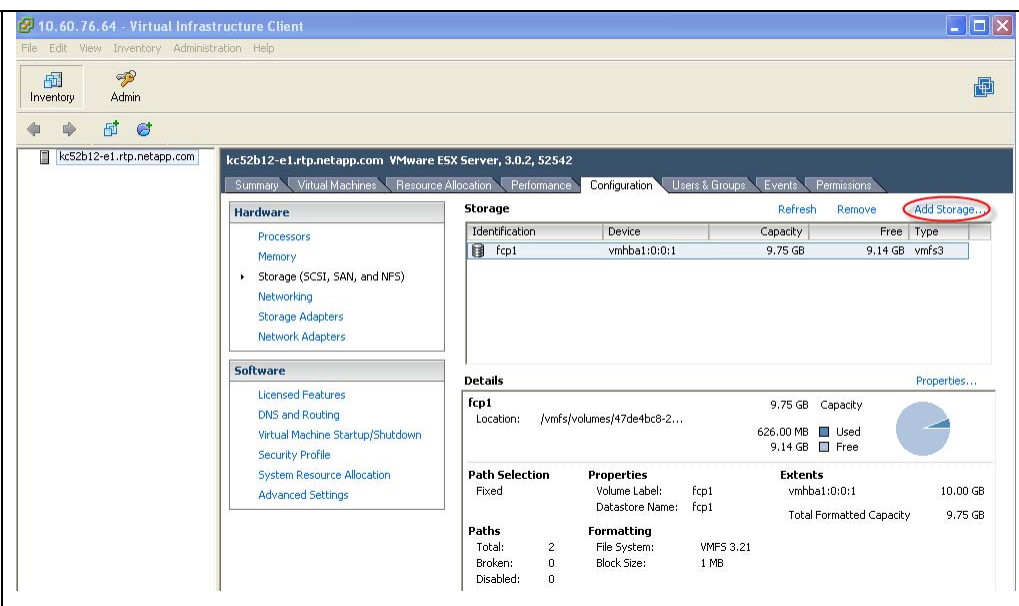
- Return to the VI client and select your ESX server.
- Click on the **configuration tab**.
- In the Hardware section, click on the **Storage Adapters** link.
- Click on the **rescan** link and wait for the ESX server to rescan its adapters. You should now be able to see the new LUN on vmhba40.

**Note:** If the new LUN is not visible, try to run the rescan a second time. Depending on the pathing configuration, you may see the LUN displayed multiple times.

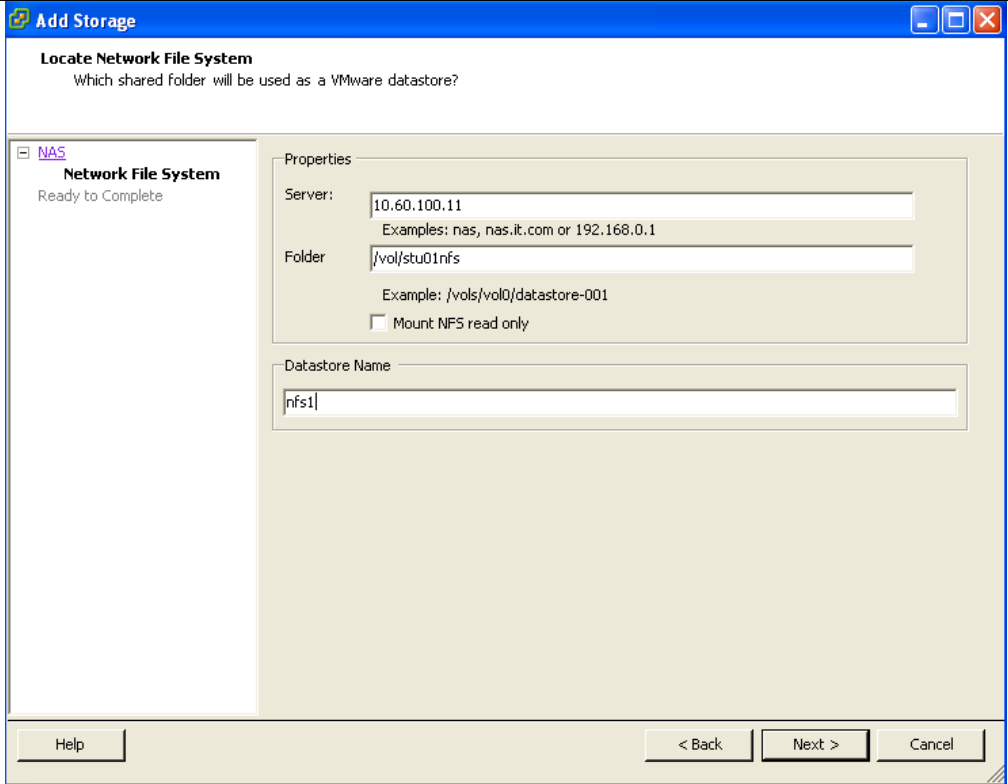




- 7 To create a VMFS datastore on the iSCSI LUN, select the Storage link in the Hardware section of the VI client.
- Start by clicking on the **Add Storage** link. In the wizard on the select storage type page, choose **Disk/LUN** and click next.
  - On the select Disk/LUN page, choose the 1 GB iSCSI LUN and click **next**
  - The current disk layout page will appear next, observe the data layout and click **next**.
  - In the Disk/LUN properties page, enter **stu#iscsi** as the name for the new datastore click **next**.
  - The Disk/LUN formatting page has options for maximum file size and capacity, leave these values at their defaults and click **next**.
  - Click **finish** to complete the add storage wizard. The new datastore will be visible on the storage page.

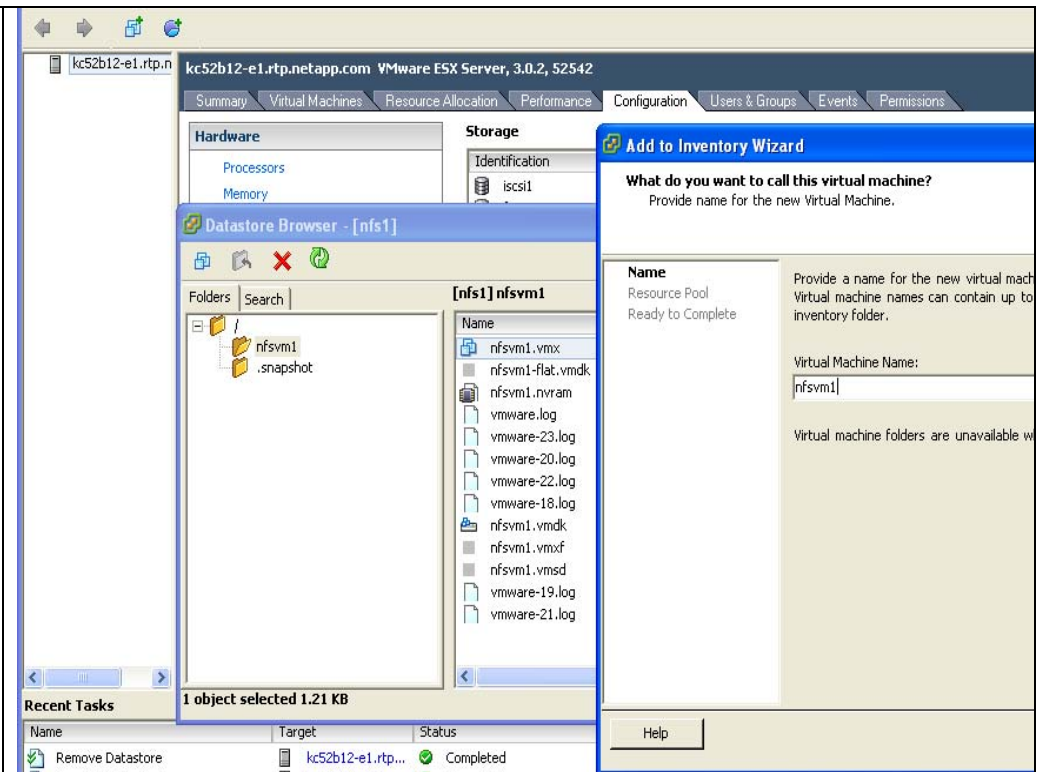


## Part III – Create an NFS Datastore.

1	<p>The NFS protocol is not enabled by default in ESX 3.0X, use the procedures below to open the ESX firewall for NFS traffic.</p> <ul style="list-style-type: none"><li>• In the VI client, select the configuration tab and then click on the <b>security profile</b> link.</li><li>• Click on <b>Properties</b> to open the Firewall Properties page.</li><li>• Find the NFS Client and <b>check the box</b> to enable the protocol. Click OK to close the Firewall properties window.</li></ul>	
2	<p><b>NOTE: This step may already have been completed. If so move to step 3</b></p> <p>A NFS volume has been pre-created for you. <b>(/vol/nfs)</b></p> <ul style="list-style-type: none"><li>• Use FilerView to confirm that your NFS volume has been exported to root on your ESX server. Look under the <b>NFS</b> heading then <b>Manage Exports</b>. Note the path of the NFS volume: _____</li></ul> <hr/> <ul style="list-style-type: none"><li>• To connect your NFS datastore open the VI client and click on the <b>configuration tab</b>.</li><li>• Click on the <b>Storage</b> link in the hardware section.</li><li>• Click on the Add Storage Link to open the Add Storage wizard, then on the Select Storage Type page, choose <b>Network File System</b> and click <b>Next</b>.</li><li>• On the locate Network File System page you will need to provide the <b>IP address</b> for your storage controller and the <b>path</b> to your NFS volume noted above.</li><li>• Leave the Mount NFS Read Only option <b>unchecked</b>.</li><li>• Name the datastore <b>stu#nfs</b>. (# is the number assigned to you by the instructor)</li><li>• Click <b>Finish</b> and verify that your NFS datastore is connected.</li></ul>	

3 Now we want to discover the pre-created Windows VM in your NFS datastore

- In the VI client click on the **configuration tab** then click on storage.
- Right click on the **stu#nfs** datastore and select **Browse Datastore**
- Open the **nfsvm1** directory, right click on the **nfsvm1.vmx** file and select **add to inventory**.
- In the window that opens, you will be asked to name the VM, use **stu#nfsvm1**.
- In the next window choose your datacenter click on **next**
- Select your ESX server and click **next**.
- There should only be one resource pool, select it, click **next**, then **finish**.
- Close the Datastore Browser window.

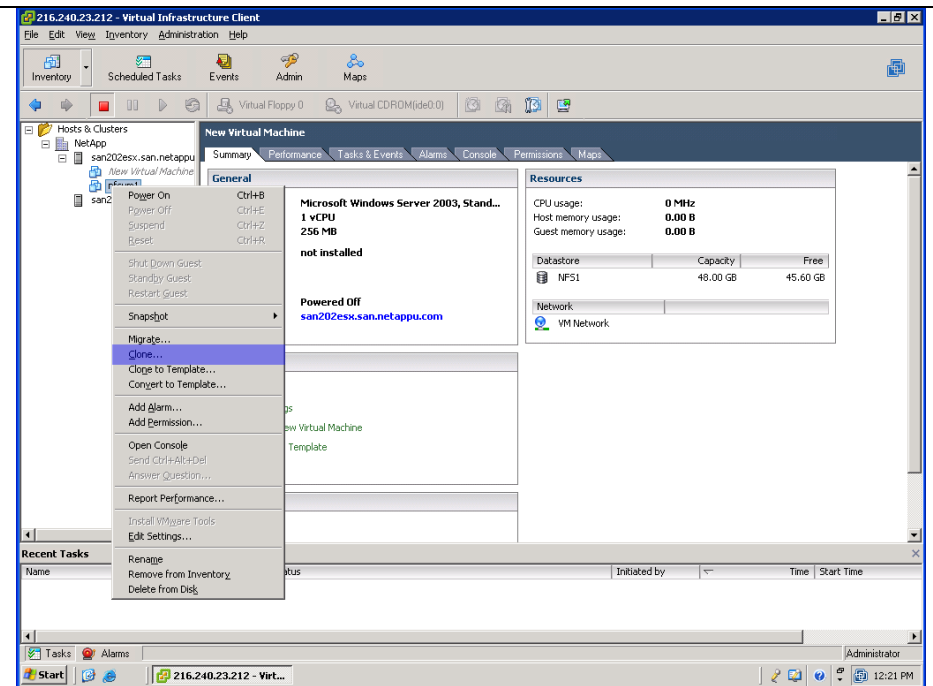


4 Next you will verify that no snapshots exist on your VM, and then create a clone of the VM just discovered.

**NOTE:** This step must be done from your Virtual Center server running on your Windows system, (the instructor will provide the IP address)

- To verify no snapshots exist on your VM, **right click** the VM and choose **Snapshot > Manage Snapshots> Delete Snapshots**
- Now that you have removed all snapshots, **right click** on the virtual machine **stu#nfsvm1** and choose **clone** off the menu.
- In the Clone VM wizard you will need to name this new VM, for the purpose of this exercise call it **stu#fcpvm1**.
- For the virtual machine inventory location, choose the NetApp datacenter and click **next**.
- Select one of your ESX servers and click **next**.
- Select the **stu#fcp** datastore and click **next**.
- Select do not customize and click **next**.
- Click **finish**, it will take several minutes for the clone to complete.

You will now see the newly created VM listed under your Resource Pool



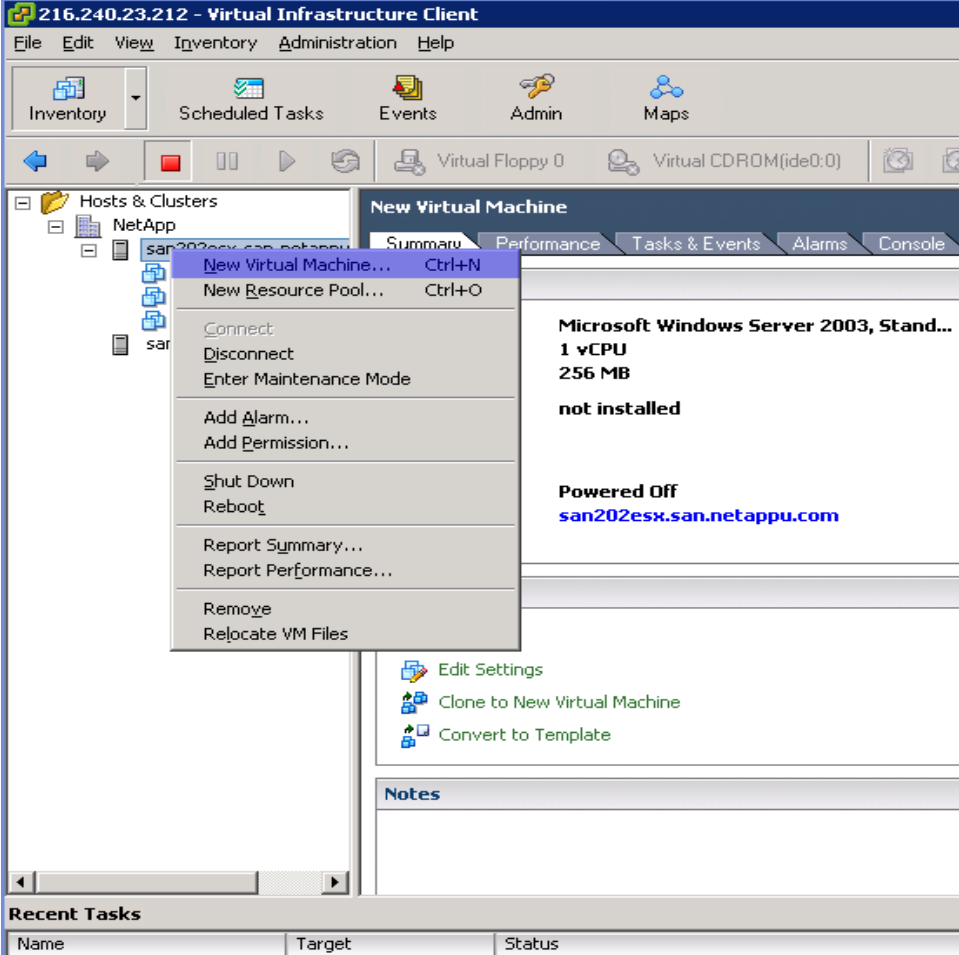


## LAB 2: Aligning Virtual Disk File Systems

### **Summary:**

In this lab you will be creating new Virtual Machines and properly aligning the virtual disk partitions within. This will be done before installing the operating systems to allow the system to run at an optimal level.

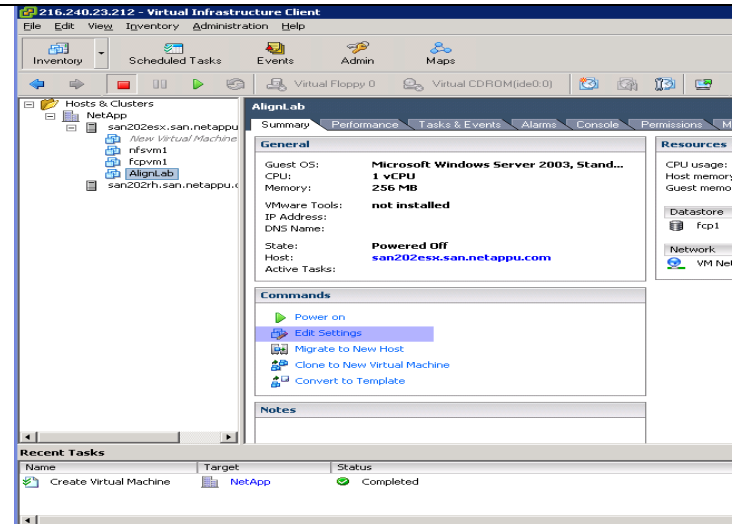
## Part I – Creating New VM's.

	Click Stream	Screen Shots
1	<p>Use your VI client and connect to the VirtualCenter server (IP provided by instructor).</p> <ul style="list-style-type: none"> <li>• Right click your ESX server and select <b>New Virtual Machine</b>.</li> <li>• In the new virtual machine window, select <b>typical</b> and click <b>next</b>.</li> <li>• Name your new Virtual Machine <b>stu#alignvm1</b>, select the datacenter (if needed) and click <b>next</b>.</li> <li>• Select the <b>stu#fcp</b> datastore and click <b>next</b>.</li> <li>• Keep the default selection of Microsoft Windows and click <b>next</b>.</li> <li>• Take the <b>default values</b> for # of processors, Virtual Machine memory &amp; Network Connections clicking next for each.</li> <li>• Change the Virtual Disk size to <b>1 GB</b> and click <b>next</b>.</li> <li>• Click <b>finish</b> and wait for the new VM to be created.</li> </ul>	 <p>The screenshot shows the VMware Virtual Infrastructure Client (VIX) interface. The main window displays the 'New Virtual Machine' wizard. The wizard is currently in the 'Summary' step, showing the following configuration:</p> <ul style="list-style-type: none"> <li>Operating System: Microsoft Windows Server 2003, Standard Edition</li> <li>Processors: 1 vCPU</li> <li>Memory: 256 MB</li> <li>Status: not installed</li> <li>Power State: Powered Off</li> <li>URL: san202esx.san.netapp.com</li> </ul> <p>The interface also includes a menu bar (File, Edit, View, Inventory, Administration, Help), a toolbar with various icons, and a left-hand navigation tree showing the hierarchy of hosts and clusters. At the bottom, there is a 'Recent Tasks' table with columns for Name, Target, and Status.</p>

## Part II – Connecting to the Window PE Boot Disk

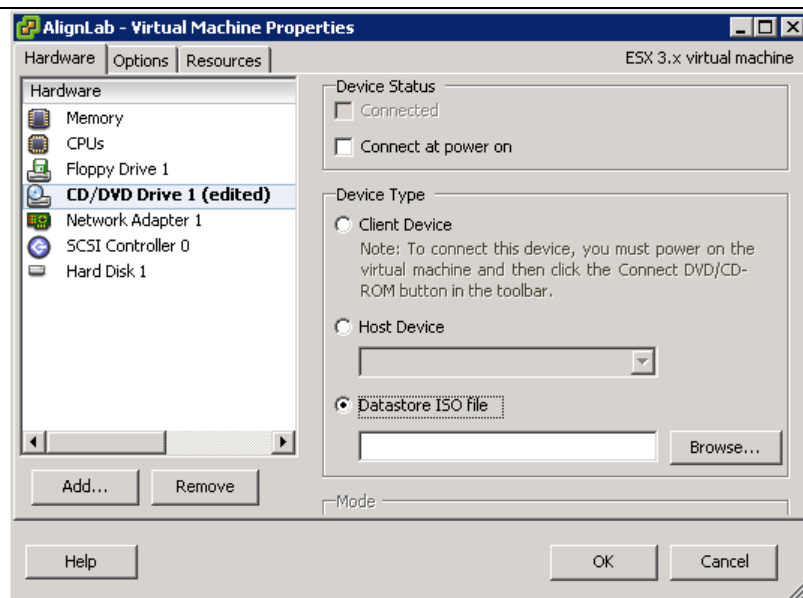
In this section we will be starting the newly created VM and aligning the partition to optimize the operating system performance.

- 1 In the VI client, select the stu#alignvm1 VM then click on **edit settings**.



2 In the virtual machine properties window, select **CD/DVD Drive 1**, then select the **Datastore ISO File** radio button.

- Click on **browse** and double click the **stu#nfs** datastore
- Select the **winpe.iso** file, then click **OK**.
- On the main properties window Check the '**connect at power on**' box then click **OK**.



3 Power on the VM, then click on the console tab. The system may take a few minutes to load.

- Once the Windows Preinstall Environment has loaded, run the '**diskpart**' command in the open command window.
- Enter '**select disk 0**'
- Enter '**create partition primary align=32**'
- type '**Exit**' to leave diskpart
- **shutdown** the virtual machine.

(You can move between VM's and VI client using ctrl-alt)

**Note:** At this point you can install the operating system as normal, using the previously created partition. We will not be installing the operating system in this workshop.





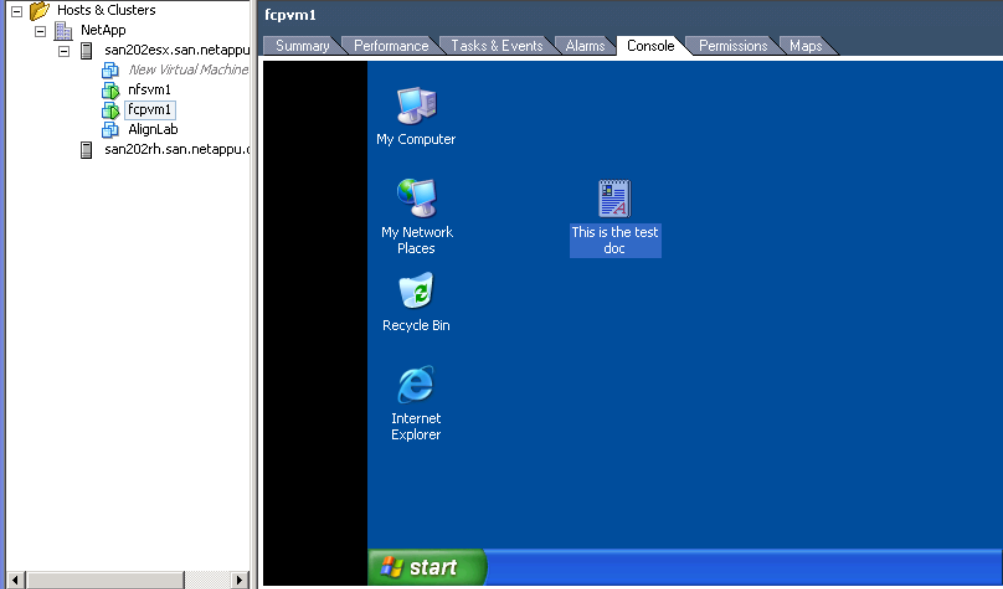


## LAB 3: Virtual Machine Backups Using NetApp Snapshots

### **Summary:**

In this lab you will take NetApp snapshots of your VMFS and NFS datastores. This will give you a crash consistent point in time copy of the VMs contained within the datastores. You will see how quick and easy it is to complete these tasks when coupled with NetApp's snapshot technology.

## Part I – Preparing VM's for Backup.

	Click Stream	Screen Shots
1	<p>You will create files on the VMs to simulate data change, then delete these files to simulate data loss.</p> <ul style="list-style-type: none"><li>• Open the VI client and connect to your VirtualCenter server.</li><li>• Power on the VMs <b>stu#nfsvm1</b> and <b>stu#fcsvm1</b></li><li>• On the console of <b>each</b> VM, log in and <b>create a file</b> in an easily accessible location such as the root of c:\ or the desktop.</li></ul>	

## Part II – Create a Snapshot Backup of the VM's

You will freeze the file system of your virtual machine by taking a VMware snapshot. While the file system is frozen you will take a NetApp snapshot of the virtual machine to create a recovery point. You can then delete the VMware snapshot as it is no longer needed.

**Note:** This exercise includes executing commands on the service console of your ESX server. There is a copy of putty on your VirtualCenter server (windows server). Use that to connect to your ESX server via SSH. The username, password and IP of your ESX server has been provided by the instructor.

You will use one snapshot creation technique for FCP, and another for NFS. This is done only for your information, in the future you can choose which technique you like best.


1 Connect to the service console on your ESX server using the putty application on the Windows server desktop.

- Enumerate the VMs on your server with the command '**vmware-cmd -l**'.

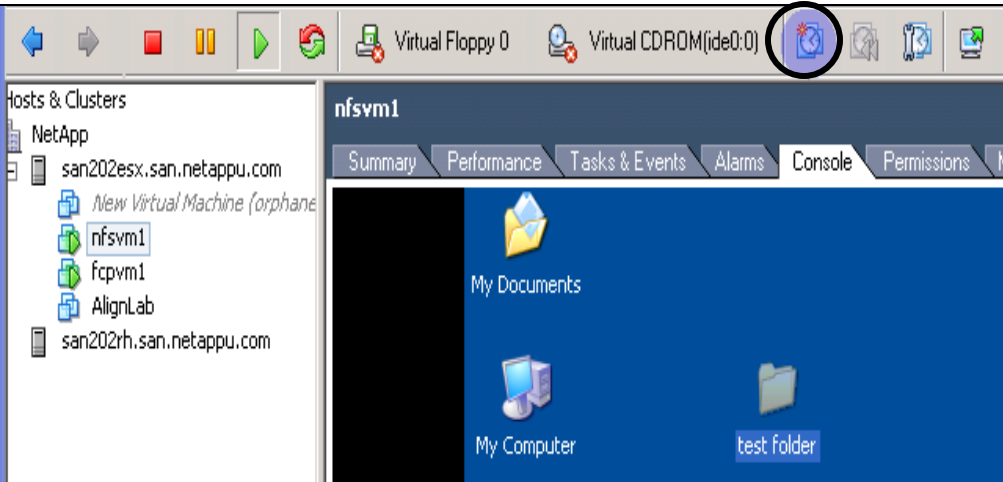
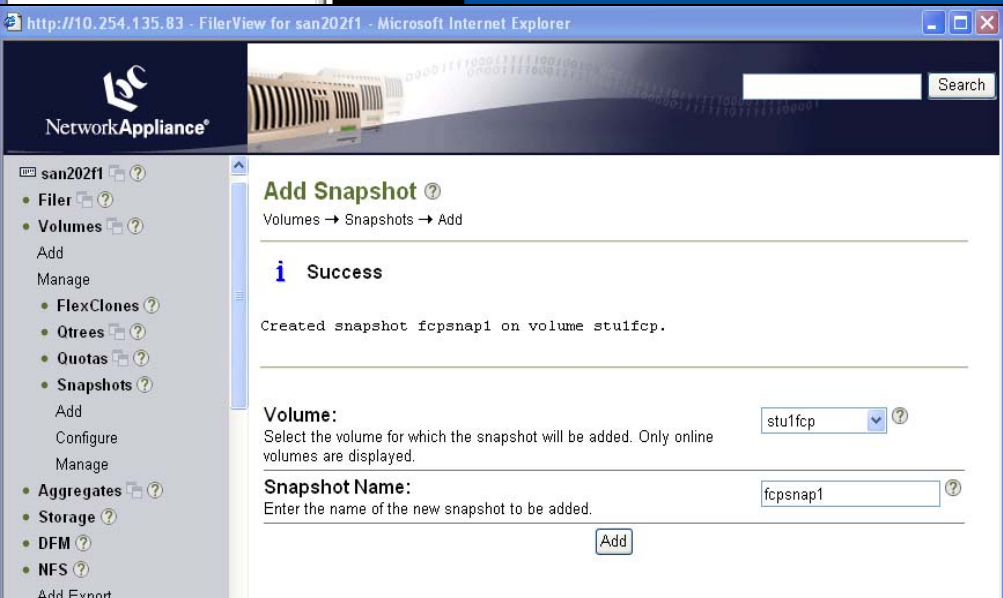
**Note:** This output will show the absolute path to the vmx file of each of your virtual machines.

- Create a VMware snapshot of your fcp VM using the command '**vmware-cmd <vmxfile> createsnapshot backup quiesce**' where <vmxfile> is the absolute path listed in the previous step.

**Note:** This will create a snapshot called 'backup' on the fcp VM.



```
root@san202esx~  
[root@san202esx root]# vmware-cmd -l  
/vms/volumes/1c76ea27-e73eaf62/nfsvm1_1/nfsvm1.vmx  
/vms/volumes/47ee7ce9-49ad733b-05ef-001a641170c8/fcpvm1/fcpvm1.vmx  
/vms/volumes/47ee7ce9-49ad733b-05ef-001a641170c8/alignLab/alignLab.vmx
```

<p>2</p>	<p>In Virtual Center GUI you will select the NFS VM to create the VMware snapshot.</p> <ul style="list-style-type: none"> <li>• Click on the take a <b>snapshot icon</b> from the main window</li> <li>• Give the snapshot the name <b>backup</b> and <b>remove</b> the check from the box to “snapshot the virtual machines memory” click <b>OK</b></li> <li>• To verify that the backup has occurred click on the snapshot manager icon.</li> </ul>	 <p>The screenshot shows the VMware vSphere Client interface. The top toolbar contains several icons, with the 'Snapshot Manager' icon (a blue circle with a white 'S') circled in red. The main window displays the 'nfsvm1' virtual machine, and the 'Hosts &amp; Clusters' pane on the left shows the VM's location within the NetApp environment.</p>
<p>3</p>	<p>Log into FilerView on your storage controller (<i>http://storage controller IP/na_admin</i>) and create a snapshot on the volumes that contain your FCP and NFS datastores.</p> <ul style="list-style-type: none"> <li>• Click on <b>Volumes &gt; Snapshots</b></li> <li>• Under Snapshots choose <b>Add</b></li> <li>• Select the fcp volume and name the snapshot <b>fcpsnap1</b></li> <li>• Select the nfs volume and name the snapshot <b>nfssnap1</b></li> </ul>	 <p>The screenshot shows the Network Appliance FilerView web interface in a Microsoft Internet Explorer browser. The page title is 'http://10.254.135.83 - FilerView for san202f1'. The main content area displays a 'Success' message: 'Created snapshot fcpsnap1 on volume stu1fcp.' Below the message, there are input fields for 'Volume:' (set to 'stu1fcp') and 'Snapshot Name:' (set to 'fcpsnap1'), with an 'Add' button at the bottom.</p>

4 Now that you have taken a NetApp snapshot of the datastore you can remove the FCP VMware snapshot from the VM.

- Return to the putty connection for the ESX server.
- The command to remove the VMware snapshot is **'vmware-cmd <vmxfile> removesnapshots'** where <vmxfile> is the absolute path listed in a previous step of the FCP VM.
- 

```

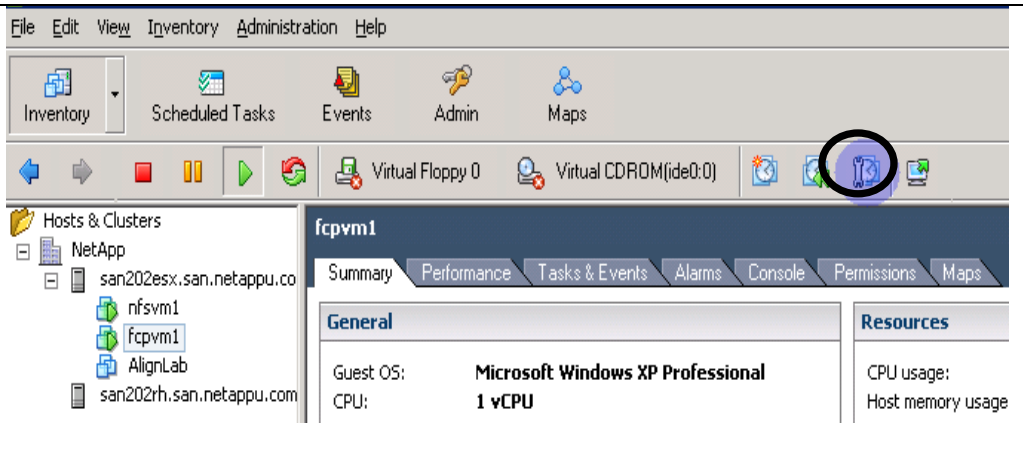
root@san202esx:~
[root@san202esx root]# vmware-cmd -l
/vmfs/volumes/1c76ea27-e73eaf62/nfsvm1_1/nfsvm1.vmx
/vmfs/volumes/47ee7ce9-49ad733b-05ef-001a641170c8/fcpvm1/fcpvm1.vmx
/vmfs/volumes/47ee7ce9-49ad733b-05ef-001a641170c8/AlignLab/AlignLab.vmx
[root@san202esx root]# vmware-cmd /vmfs/volumes/1c76ea27-e73eaf62/nfsvm1_1/nfsvm1.vmx removesnapshots

```

For the NFS VM instead of using the service console command line, you can use the VI client to delete the VMware snapshot.

- Select **nfsvm1**
- Click on the **snapshot manager** icon in the VI client.
- Then choose **delete all** on the snapshot window
- Answer **yes** to the confirm delete window
- Then click close

All NFS VMware snapshots have been removed for that system.







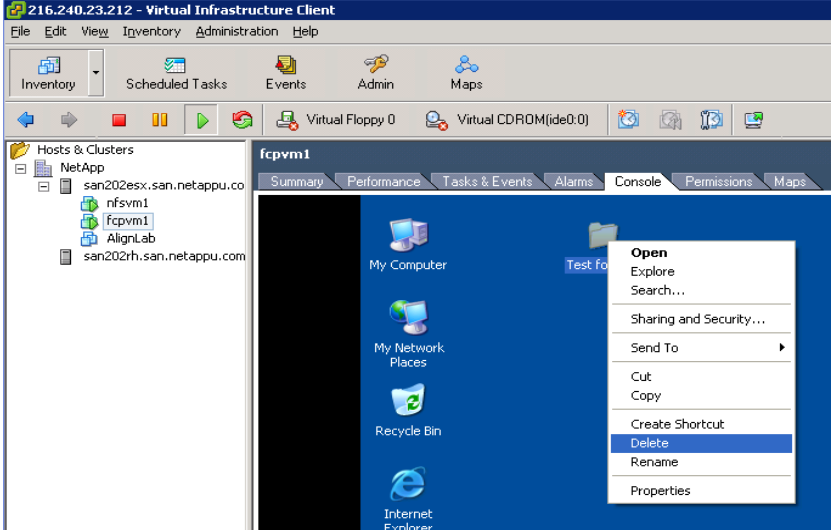
## LAB 4: Virtual Machine Recovery from NetApp Snapshots

Last Updated: Wednesday, May 28, 2008

### **Summary:**

In this lab you will recover the VMs to their previous state after simulating a failure by removing the files that were created in the previous lab and restoring the virtual machines from a NetApp snapshot copy.

## Part I – Prepare VMs for Recovery

	Click Stream	Screen Shots
1	<p>Open the VI client and connect to your VirtualCenter server.</p> <ul style="list-style-type: none"><li>• Select <b>stu#fcsvm1</b> from the menu on the left</li><li>• Click on the <b>console tab</b>, log in and delete the file created in the previous lab.</li><li>• Right click on the Recycle Bin and choose empty</li><li>• Confirm that the file does not exist on the VM</li><li>• Repeat the above steps for <b>stu#nfsvm1</b></li><li>• From within the windows console <b>shut down and power off</b> both of the VMs</li></ul>	 <p>The screenshot displays the VMware Virtual Infrastructure Client interface. The title bar reads '216.240.23.212 - Virtual Infrastructure Client'. The main window shows the 'Hosts &amp; Clusters' tree on the left with 'fcsvm1' selected. The 'Console' tab is active, showing a Windows desktop environment. A context menu is open over the 'Recycle Bin' icon, with the 'Delete' option highlighted. The desktop includes icons for 'My Computer', 'My Network Places', 'Recycle Bin', and 'Internet Explorer'. The taskbar shows 'Virtual Floppy 0' and 'Virtual CDRROM(ide0:0)'.</p>





2	<p>In this exercise we will use LUN clone on the storage controller to connect to the snapshot version of the VMFS datastore, we will then copy the VMs virtual disk file from the clone to the production datastore.</p> <ul style="list-style-type: none"> <li>Using the putty utility on the desktop of you windows server, connect to the service console of your ESX server using the ip address provided</li> <li>Go to the directory containing the files owned by stu#fcsvm1 using the command <b>'cd /vmfs/volumes/stu#fcp/stu#fcsvm1'</b></li> <li>List the contents of this directory using <b>'ls -al'</b></li> <li>Rename the current VMDK file used by the VM fcpvm1, the command to do so will be <b>'mv /vmfs/volumes/stu#fcp/stu#fcsvm1/fcpvm1-flat.vmdk /vmfs/volumes/stu#fcp/stu#fcsvm1/fcpvm1-flat.vmdk.old'</b></li> <li>Run the <b>ls -al</b> command to verify that the file was changed</li> </ul>
3	<p>Using the putty utility connect to the console of your storage controller and create a LUN clone of the VMFS datastore using the snapshot taken in the previous exercise.</p> <ul style="list-style-type: none"> <li>In the storage controller putty connection enter the command <b>snap list</b> to display all system snapshots Note what snapshots exist.</li> <li>Enter the command <b>lun show all</b> to see all LUNs on the sytem. Note what LUNs exist.</li> <li>Now use the following command to create a LUN clone from your previously created snapshot. The command will be <b>'lun clone create &lt;clone LUN</b></li> </ul>

```

root@san202esx:/vmfs/volumes/fcp1/fcpvm1
[root@san202esx /]# cd /vmfs/volumes/fcp1/fcpvm1
[root@san202esx fcpvm1]# ls -al
total 2099008
drwxr-xr-x 1 root root 2100 Mar 30 05:52 .
drwxrwxrwt 1 root root 1260 Mar 29 12:58 ..
-rw----- 1 root root 2147483648 Mar 30 05:52 fcpvm1-flat.vmdk
-rw----- 1 root root 8664 Mar 30 05:52 fcpvm1.nvram
-rw----- 1 root root 336 Mar 30 05:39 fcpvm1.vmdk
-rw----- 1 root root 783 Mar 30 05:24 fcpvm1.vmsd
-rwxr-xr-x 1 root root 1617 Mar 30 04:58 fcpvm1.vmx
-rw----- 1 root root 250 Mar 30 04:58 fcpvm1.vmxfs
-rw-r--r-- 1 root root 23188 Mar 30 04:44 vmware-10.log
-rw-r--r-- 1 root root 23410 Mar 29 12:27 vmware-5.log
-rw-r--r-- 1 root root 20612 Mar 29 12:27 vmware-6.log
-rw-r--r-- 1 root root 23305 Mar 29 12:27 vmware-7.log
-rw-r--r-- 1 root root 23000 Mar 29 12:27 vmware-8.log
-rw-r--r-- 1 root root 25788 Mar 29 15:59 vmware-9.log
-rw-r--r-- 1 root root 37326 Mar 30 05:52 vmware.log
[root@san202esx fcpvm1]# mv /vmfs/volumes/fcp1/fcpvm1/fcpvm1-flat.vmdk /vmfs/volumes/
fcp1/fcpvm1/fcpvm1-flat.vmdk.old
[root@san202esx fcpvm1]# ls -al
total 2099008
drwxr-xr-x 1 root root 2100 Mar 30 06:24 .
drwxrwxrwt 1 root root 1260 Mar 29 12:58 ..
-rw----- 1 root root 2147483648 Mar 30 05:52 fcpvm1-flat.vmdk.old
-rw----- 1 root root 8664 Mar 30 05:52 fcpvm1.nvram
-rw----- 1 root root 336 Mar 30 05:39 fcpvm1.vmdk
-rw----- 1 root root 783 Mar 30 05:24 fcpvm1.vmsd
-rwxr-xr-x 1 root root 1617 Mar 30 04:58 fcpvm1.vmx
-rw----- 1 root root 250 Mar 30 04:58 fcpvm1.vmxfs
-rw-r--r-- 1 root root 23188 Mar 30 04:44 vmware-10.log
-rw-r--r-- 1 root root 23410 Mar 29 12:27 vmware-5.log
-rw-r--r-- 1 root root 20612 Mar 29 12:27 vmware-6.log
-rw-r--r-- 1 root root 23305 Mar 29 12:27 vmware-7.log
-rw-r--r-- 1 root root 23000 Mar 29 12:27 vmware-8.log
-rw-r--r-- 1 root root 25788 Mar 29 15:59 vmware-9.log
-rw-r--r-- 1 root root 37326 Mar 30 05:52 vmware.log
[root@san202esx fcpvm1]#

```

```

10.254.135.83 - PuTTY
login as: root
root@10.254.135.83's password:

Data ONTAP (san202f1.SAN.NetappU.com)
login: root
san202f1>
FCP service is running.
san202f1> lun clone create /vol/stu1fcp/fcplunclone -b /vol/stu1fcp/fcplun fcpsnap1

```

**path> -b <original LUN path> <snapshot>** where <clone LUN path> is your desired path to the clone <original LUN path> is the path to the current production LUN and <snapshot> is the name of the snapshot you created in the previous lab

**Example:** lun clone create /vol/fcp/fcplunclone -b /vol/fcp/fcplun fcpsnap1

4 Use FilerView to online the LUN

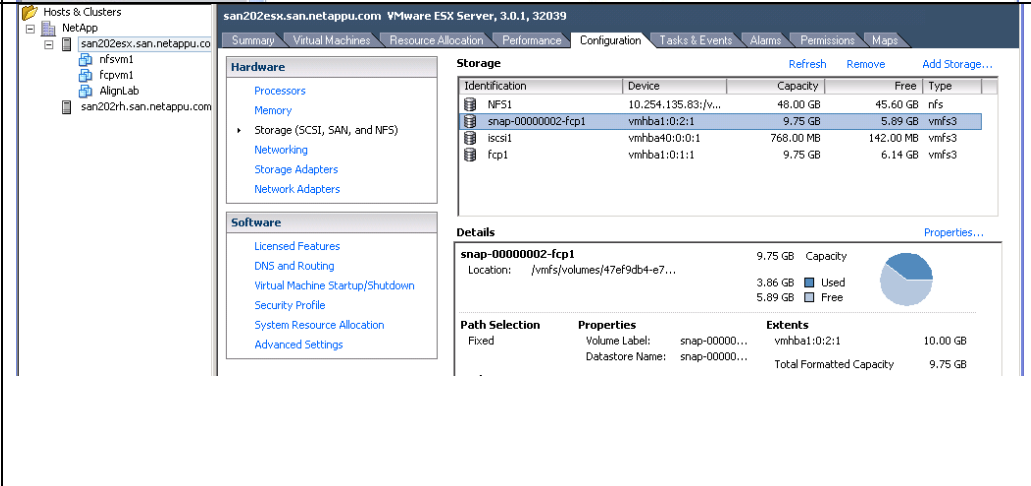
- Map it to your FCP igroup
- In the LUN ID field, **increment** the LUN id (use 2) so it is not the same as the production LUN



5 In the VI client under the configuration tab of your ESX server choose storage adapters and then click rescan to discover the new LUN and datastore.

- You may need to run the process twice, once to find the LUN and a second time to discover the datastore.
- To verify that the datastore has been discovered go to the storage heading under hardware. The cloned datastore will automatically be renamed to something different than the production datastore (Example: snap-00000001-fcp1).

**Note** this datastore name as it will be important later.



6	<p>Return to the service console on your ESX server and change directory to the cloned version of the FCP datastore.</p> <ul style="list-style-type: none"> <li>The command will be <b>'cd /vmfs/volumes/&lt;datastore name from previous step&gt;/stu#fcsvm1'</b>.</li> <li>List the contents of this directory using <b>'ls -al'</b></li> <li>Note that all of the same files as the production datastore are present.</li> <li>From the cloned datastore copy the VMs virtual disk file to the production datastore. The command is <b>cp fcsvm1-flat.vmdk /vmfs/volumes/stu#fcp/stu#fcsvm1/</b></li> </ul> <p>Change directories to the original datastore to verify the copy has occurred</p> <p>In the VI client select stu#fcsvm1 and power it on.</p> <ul style="list-style-type: none"> <li>Click on the console tab of the VM, and check that the file deleted earlier in the lab has been restored.</li> </ul>	<pre>[root@san202esx snap-00000002-fcp1]# cd /vmfs/volumes/snap-00000002-fcp1/fcsvm1 [root@san202esx fcsvm1]# ls -al total 2361216 drwxr-xr-x 1 root root 2380 Mar 30 04:53 . drwxrwxrwt 1 root root 1260 Mar 29 12:58 .. -rw----- 1 root root 268435456 Mar 30 04:50 fcsvm1-ed5328ea.vswp -rw----- 1 root root 2147483648 Mar 30 04:53 fcsvm1-flat.vmdk -rw----- 1 root root 8664 Mar 30 04:50 fcsvm1.nvram -rw----- 1 root root 18597 Mar 30 04:53 fcsvm1-Snapshot5.vmsn -rw----- 1 root root 336 Mar 30 04:53 fcsvm1.vmdk -rw----- 1 root root 783 Mar 30 04:53 fcsvm1.vmsd -rwxr-xr-x 1 root root 1617 Mar 30 04:58 fcsvm1.vmx -rw----- 1 root root 250 Mar 30 04:58 fcsvm1.vmx.f -rw-r--r-- 1 root root 23188 Mar 30 04:44 vmware-10.log -rw-r--r-- 1 root root 23410 Mar 29 12:27 vmware-5.log -rw-r--r-- 1 root root 20612 Mar 29 12:27 vmware-6.log -rw-r--r-- 1 root root 23305 Mar 29 12:27 vmware-7.log -rw-r--r-- 1 root root 23000 Mar 29 12:27 vmware-8.log -rw-r--r-- 1 root root 25788 Mar 29 15:59 vmware-9.log -rw-r--r-- 1 root root 29212 Mar 30 04:58 vmware.log [root@san202esx fcsvm1]# cp fcsvm1-flat.vmdk /vmfs/volumes/fcp1/fcsvm1/</pre>

7


\*\*\*\* **Optional Step:** Only complete this step if you are **not** doing the file level recovery lab next. Ask instructor for Verification \*\*\*\*\*

**Remove** the LUN clone either using the storage controller command line (**lun destroy -f <cloned LUN path>**), or FilerView under manage LUNs.

- Use the VI client to rescan the storage adapters on your ESX server
- Confirm that the clone datastore is no longer connected to the ESX server by clicking the storage link under the hardware heading

Delete the renamed virtual disk files from the stu#fcp datastore.

- To do this double click on the stu#fcp datastore then browse the stu#fcpvm1 folder
- Right click the fcpvm1-flat.vmdk.old and choose delete from disk
- Now select yes to confirm, the file has been removed



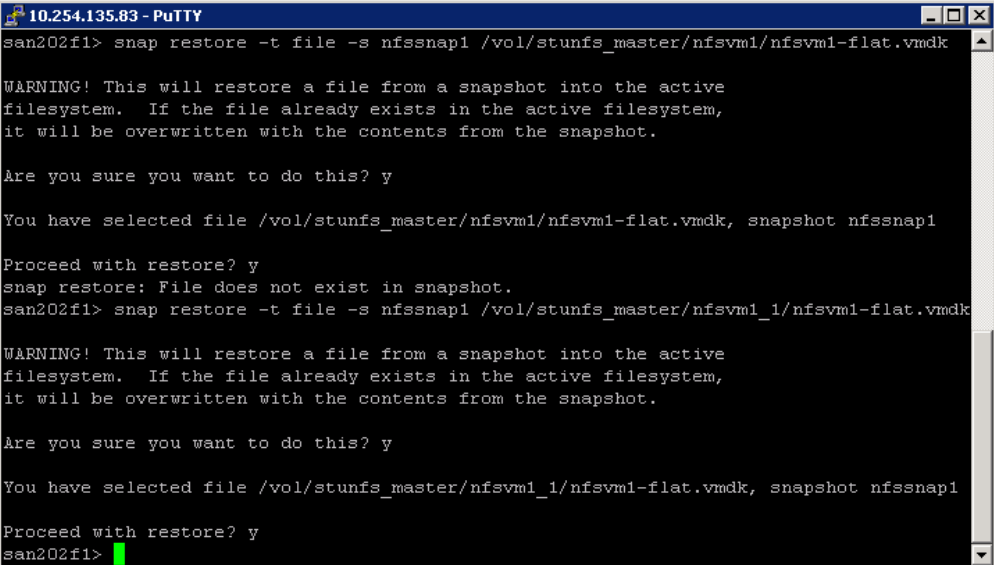
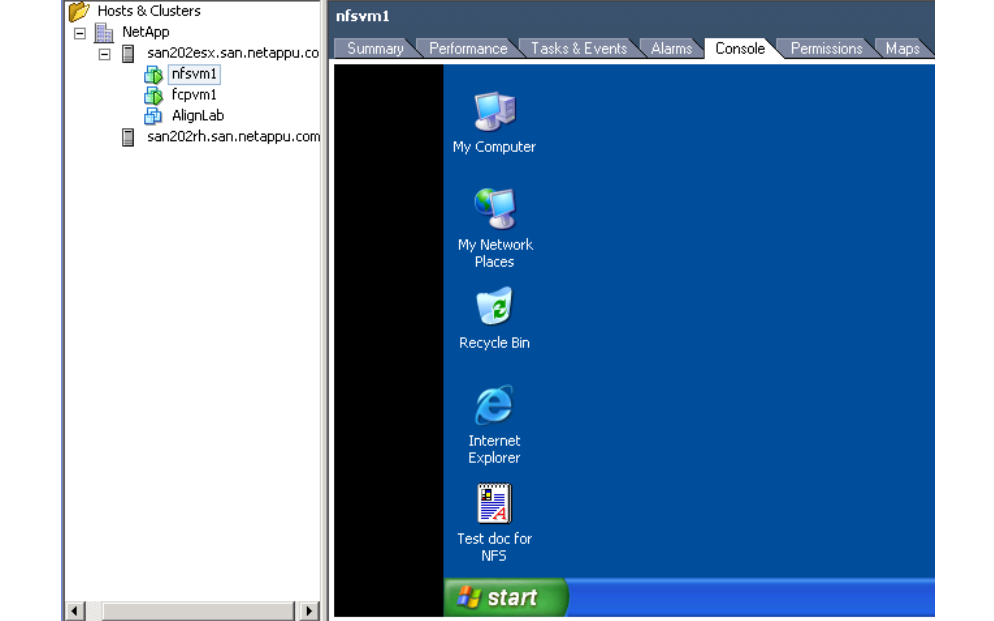
## Part III – Recover a VM from a NFS datastore

	<p>You will now recover the state of the NFS VM to the point prior to the file deletion. You could use a FlexClone and similar procedures as part 2 of this lab, but in this case we will use Single File SnapRestore (SFSR) to restore the virtual disks to demonstrate the efficiency of the NetApp system.</p>
<p>1</p>	<p>Connect to the service console on your ESX server using the putty tool and the ESX server IP address.</p> <ul style="list-style-type: none"> <li>Go to the directory containing the files owned by the VM <code>stu#nfsvm1</code> using the command <code>'cd /vmfs/volumes/stu#nfs/nfsvm1'</code></li> <li>List the contents of this directory using <code>'ls -al'</code></li> <li>Rename the current VMDK file used by the VM <code>stu#nfsvm1</code>, the command to do so will be <code>'mv /vmfs/volumes/stu#nfs/nfsvm1/nfsvm1-flat.vmdk /vmfs/volumes/stu#nfs/nfsvm1/nfsvm1-flat.vmdk.old'</code></li> <li>Run the <code>ls -al</code> command to verify that the file was change</li> </ul>

```

root@san202esx:/vmfs/volumes/nfs1/nfsvm1_1
-rwxr-xr-x  1 root    root      84 Mar 19 16:16 .lck-7b07170000000000
-rw-r----- 1 root    root    268435456 Mar 19 16:16 nfsvm1-2d325418.vswp
-rw-r----- 1 root    root    2147483648 Mar 30 05:51 nfsvm1-flat.vmdk
-rw-r----- 1 root    root     8664 Mar 30 05:51 nfsvm1.nvram
-rw-r----- 1 root    root     336 Mar 30 05:06 nfsvm1.vmdk
-rw-r----- 1 root    root     793 Mar 30 05:06 nfsvm1.vmsd
-rwxr----- 1 root    root    1744 Mar 30 10:41 nfsvm1.vmx
-rw-r----- 1 root    root     250 Mar 30 10:41 nfsvm1.vmaxf
-rw-r--r--  1 root    root    26690 Mar 29 16:36 vmware-10.log
-rw-r--r--  1 root    root    22037 Mar 29 16:37 vmware-11.log
-rw-r--r--  1 root    root    22862 Mar 30 04:44 vmware-12.log
-rw-r--r--  1 root    root    25491 Mar 30 04:51 vmware-13.log
-rw-r--r--  1 root    root    25773 Mar 30 04:57 vmware-14.log
-rw-r--r--  1 root    root    22365 Mar 30 04:58 vmware-15.log
-rw-r--r--  1 root    root    25614 Mar 30 05:51 vmware.log
[root@san202esx nfsvm1_1]# mv /vmfs/volumes/nfs1/nfsvm1_1/nfsvm1-flat.vmdk /vmfs/vol
mes/nfs1/nfsvm1_1/nfsvm1-flat.vmdk.old
[root@san202esx nfsvm1_1]# ls -al
total 2364172
drwxr-xr-x  1 root    root      4096 Mar 30 10:53 .
drwxr-xr-x  1 root    root      4096 Mar 29 13:25 ..
-rwxr-xr-x  1 root    root      84 Mar 19 16:16 .lck-6f07170000000000
-rwxr-xr-x  1 root    root      84 Mar 19 16:16 .lck-7b07170000000000
-rw-r----- 1 root    root    2147483648 Mar 30 05:51 nfsvm1-flat.vmdk.old
-rw-r----- 1 root    root    268435456 Mar 19 16:16 nfsvm1-2d325418.vswp
-rw-r----- 1 root    root     8664 Mar 30 05:51 nfsvm1.nvram
-rw-r----- 1 root    root     336 Mar 30 05:06 nfsvm1.vmdk
-rw-r----- 1 root    root     793 Mar 30 05:06 nfsvm1.vmsd
-rwxr----- 1 root    root    1744 Mar 30 10:41 nfsvm1.vmx
-rw-r----- 1 root    root     250 Mar 30 10:41 nfsvm1.vmaxf
-rw-r--r--  1 root    root    26690 Mar 29 16:36 vmware-10.log
-rw-r--r--  1 root    root    22037 Mar 29 16:37 vmware-11.log
-rw-r--r--  1 root    root    22862 Mar 30 04:44 vmware-12.log
-rw-r--r--  1 root    root    25491 Mar 30 04:51 vmware-13.log
-rw-r--r--  1 root    root    25773 Mar 30 04:57 vmware-14.log
-rw-r--r--  1 root    root    22365 Mar 30 04:58 vmware-15.log
-rw-r--r--  1 root    root    25614 Mar 30 05:51 vmware.log
[root@san202esx nfsvm1_1]#

```

<p>2</p>	<p>Connect to the console of your storage controller and restore the virtual disk file using Single File SnapRestore.</p> <ul style="list-style-type: none"> <li>At the command prompt enter '<b>snap restore -t file -s &lt;snapshot&gt; /vol/&lt;nfsvol&gt;/nfsvm1/nfsvm1-flat.vmdk</b>' where &lt;snapshot&gt; is the name of the snapshot you created in the previous lab and &lt;nfsvol&gt; is the name of your NFS volume on the NetApp system</li> <li>The warning will appear that a file will be overwritten if it exists. Answer <b>Yes</b></li> <li>It will again ask to continue with the restore. Answer <b>Yes</b></li> </ul> <p><b>Note:</b> The single file SnapRestore operation is not instantaneous. Return to the ESX service console and run '<b>ls -al</b>' in your VMs directory to check the progress of the SnapRestore.</p>	 <pre> 10.254.135.83 - PuTTY san202f1&gt; snap restore -t file -s nfssnap1 /vol/stunfs_master/nfsvm1/nfsvm1-flat.vmdk WARNING! This will restore a file from a snapshot into the active filesystem.  If the file already exists in the active filesystem, it will be overwritten with the contents from the snapshot.  Are you sure you want to do this? y  You have selected file /vol/stunfs_master/nfsvm1/nfsvm1-flat.vmdk, snapshot nfssnap1  Proceed with restore? y snap restore: File does not exist in snapshot. san202f1&gt; snap restore -t file -s nfssnap1 /vol/stunfs_master/nfsvm1_1/nfsvm1-flat.vmdk WARNING! This will restore a file from a snapshot into the active filesystem.  If the file already exists in the active filesystem, it will be overwritten with the contents from the snapshot.  Are you sure you want to do this? y  You have selected file /vol/stunfs_master/nfsvm1_1/nfsvm1-flat.vmdk, snapshot nfssnap1  Proceed with restore? y san202f1&gt; </pre>
<p>3</p>	<p>When the SnapRestore is complete (the VMDK stops growing in size) Open the VI client for your ESX server.</p> <ul style="list-style-type: none"> <li>Power on <b>stu#nfsvm1</b></li> <li>Log into the VM, check that the file deleted earlier in the lab has been restored.</li> </ul> <p>Clean up by deleting the renamed virtual disk files either by:</p> <ul style="list-style-type: none"> <li>Using the command '<b>rm /vmfs/volumes/stu#nfs/nfsvm1/nfsvm1-flat.vmdk.old</b>' on the ESX server</li> </ul> <p>OR</p> <ul style="list-style-type: none"> <li>Double clicking on the <b>stu#nfs</b> datastore in the VI client and right clicking on the file and choosing remove.</li> </ul>	







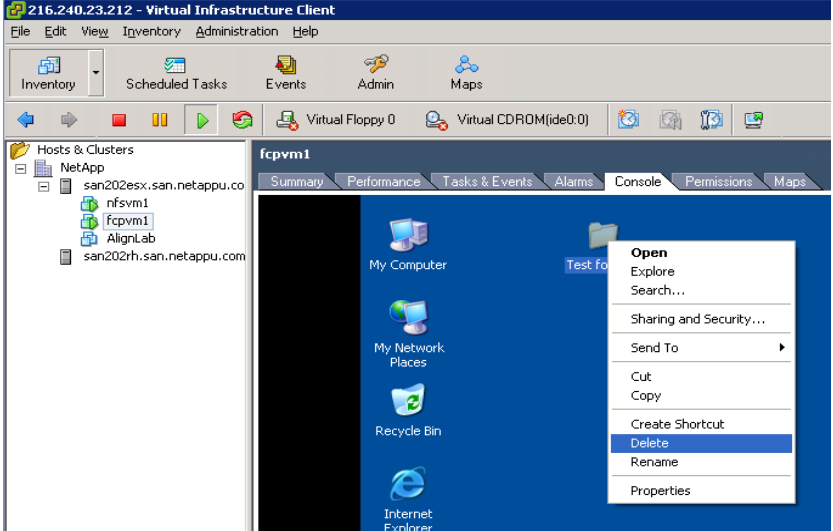
## LAB 5: File Level Recovery from NetApp Snapshots

Last Updated: Wednesday, May 28, 2008

### **Summary:**

In this lab you will recover the file we created on the VMs in the previous lab. In this exercise, you will connect to the snapshot and copy just the desired file back to the VM, rather than revert the entire VM. You will see how granular a restore can be.

## Part I – Prepare VMs for Recovery

	Click Stream	Screen Shots
1	<p>Open the VI client and connect to your VirtualCenter server.</p> <ul style="list-style-type: none"> <li>• Select <b>stu#fcpvm1</b> from the menu on the left</li> <li>• Click on the <b>console tab</b>, log in and delete the file created in the previous lab.</li> <li>• Confirm that the file does not exist on the VM</li> </ul> <p><b>Note:</b> You may also want to remove it from the recycling bin so it can not be recovered there.</p> <ul style="list-style-type: none"> <li>• Repeat the above steps for <b>stu#nfsvm1</b></li> <li>• From within the windows console shut down and power off both of the VMs</li> </ul>	

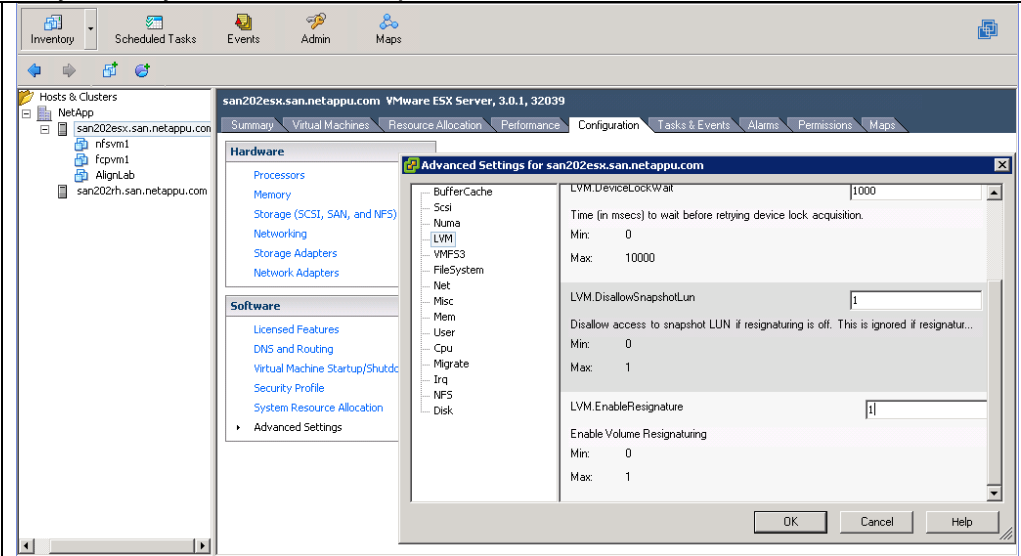
## Part II – Recover a file from a VMFS datastore

This exercise will include executing commands on the service console of your ESX server. There should be a copy of putty on your VirtualCenter server. Use this or a client on your workstation to connect to your ESX server via SSH. The username and password of your ESX server has been provided by your instructor. You will also need to be able to run commands on the console of your storage controller ensure that SSH is enabled. Using FilerView got to **Secure Admin > SSH > Enable/Disable** and verify that SSH v2 is enabled. If you have completed these steps in a previous exercise, you may move on to step 1.

1 **\*\*You may have completed this step in a previous exercise. If so, skip to the next step.**

You will be connecting to the NetApp snapshot copy of a VMFS datastore, before you can do so we need to enable the Volume Resignature option. Use the VI client to connect to your VirtualCenter server.

- Select your ESX server, click on the **configuration tab** then click on **advanced settings** under the software heading
- When the advance settings window opens Select **LVM**
- Set the value of **LVM.EnableResignature** to 1 then click **OK**



You will use LUN clone on the Storage Controller to connect to the snapshot version of the VMFS datastore, you will then mount the virtual disk in the clone datastore to the production virtual machine.

**NOTE:** For the purposes of this exercise you will be mounting the recovery disk on the production VM, in a customer's environment you would use a nonproduction VM or a recovery VM dedicated to this purpose.

2 \*\*\*\*\***Optional Step:** If you just completed the VM Recovery from snapshots lab skip to step 5. \*\*\*\*\*

Using the putty utility connect to the console of your Storage Controller and create a LUN clone of the VMFS datastore using the snapshot taken in the previous exercise.

- The command will be '**lun clone create <clone LUN path> -b <original LUN path> <snapshot>**' where <clone LUN path> is your desired path to the clone <original LUN path> is the path to the current production LUN and <snapshot> is the name of the snapshot you created in the previous lab

**Example:** lun clone create /vol/fcp/fcplunclone -b /vol/fcp/fcplun fcpsnap1

```
10.254.135.83 - PuTTY
login as: root
root@10.254.135.83's password:

Data ONTAP (san202f1.SAN.NetappU.com)
login: root
san202f1>
FCP service is running.
san202f1> lun clone create /vol/stu1fcp/fcplunclone -b /vol/stu1fcp/fcplun fcpsnap1
```

3 Use FilerView to online the LUN

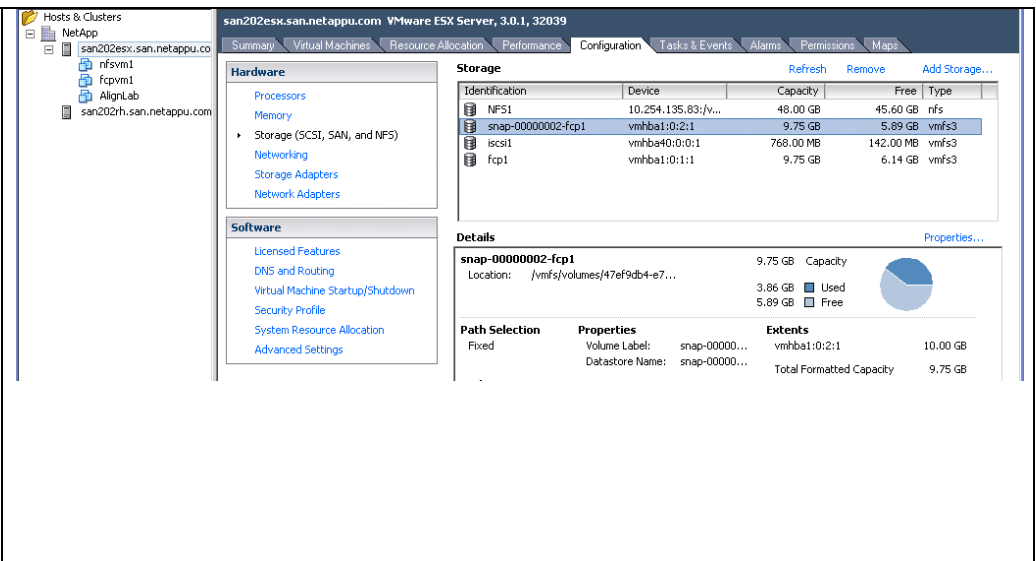
- Map it to your FCP igroup
- In the LUN ID field, **increment** the LUN id (use 2) so it is not the same as the production LUN



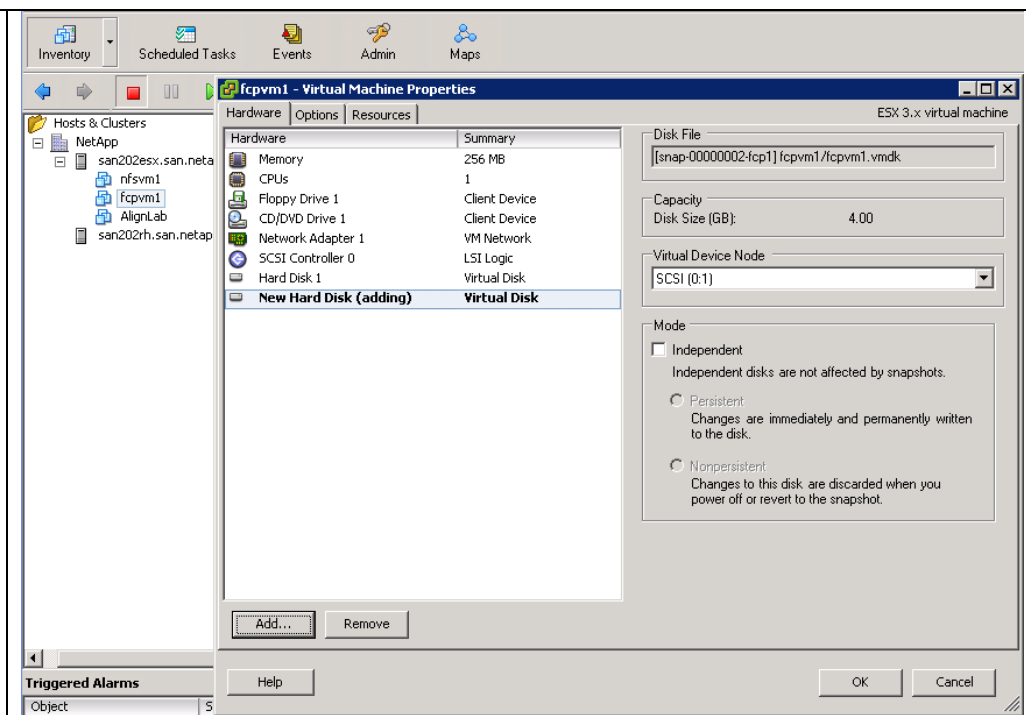
4 In the VI client under the configuration tab of your ESX server choose storage adapters and then click rescan to discover the new LUN and datastore.

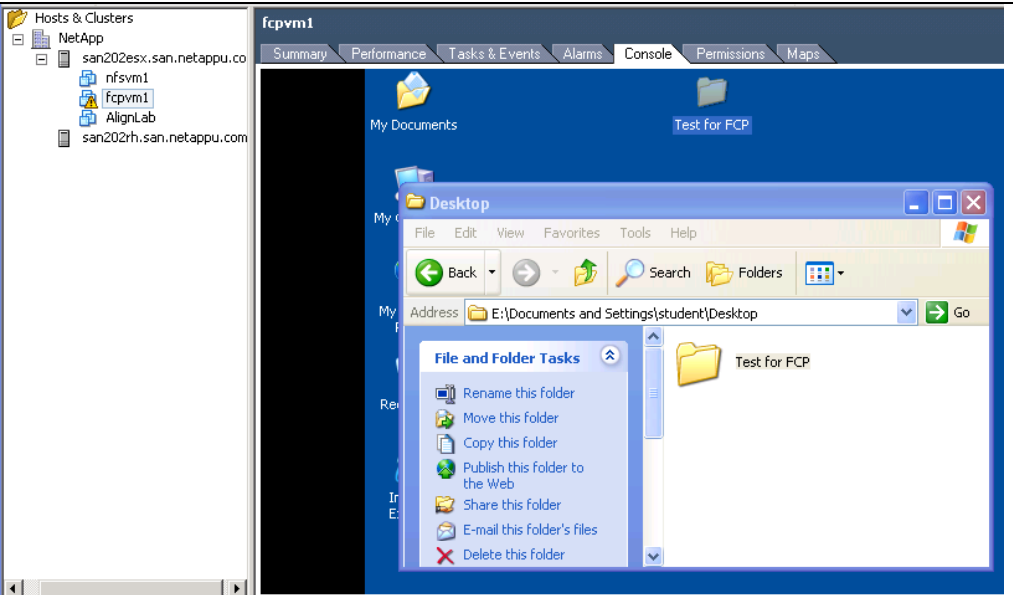

- You may need to run the process **twice**, once to find the LUN and a second time to discover the datastore.
- To verify that the datastore has been discovered go to the storage heading under hardware. The cloned datastore will automatically be renamed to something different than the production datastore (Example: snap-00000001-fcp1).

**Note** this datastore name as it will be important later.

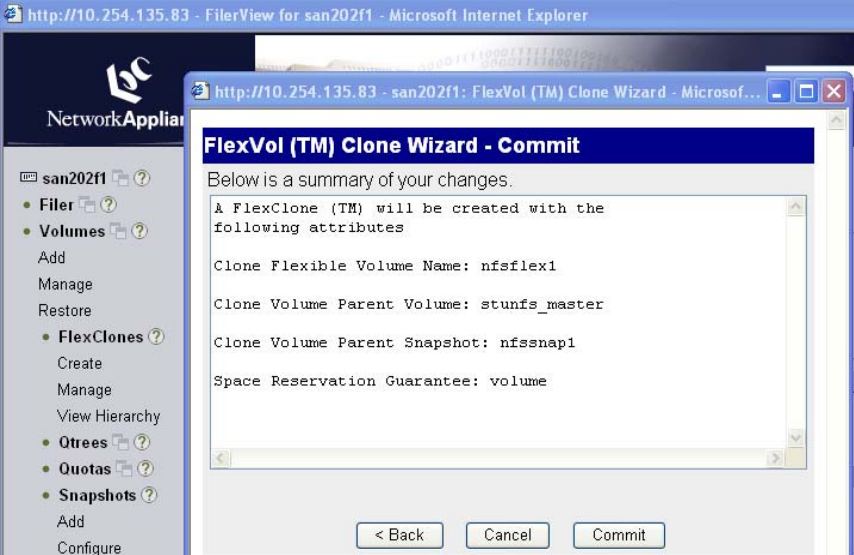


- 5 Now you will mount the clone virtual disk onto the production virtual machine.
- Open the VI client, select the **stu#fcpvm1** and click on Edit Settings under the summary tab.
  - In the Virtual Machine Properties window that will open, click on the **Add** button.
  - In the Add Hardware wizard, select **hard disk** and click **next**
  - Select the “Use an existing virtual disk” radio button and click **next**.
  - Click on **browse**, select the clone datastore created previously. (Example: snap-00000001-fcp1)
  - Open the fcpvm1 folder, select the **stu#fcpvm1.vmdk** file and click **OK**.
  - Keep the default values in the specific advanced options screens that follow clicking **next**.
  - Click on **finish** then **OK** to close the properties window

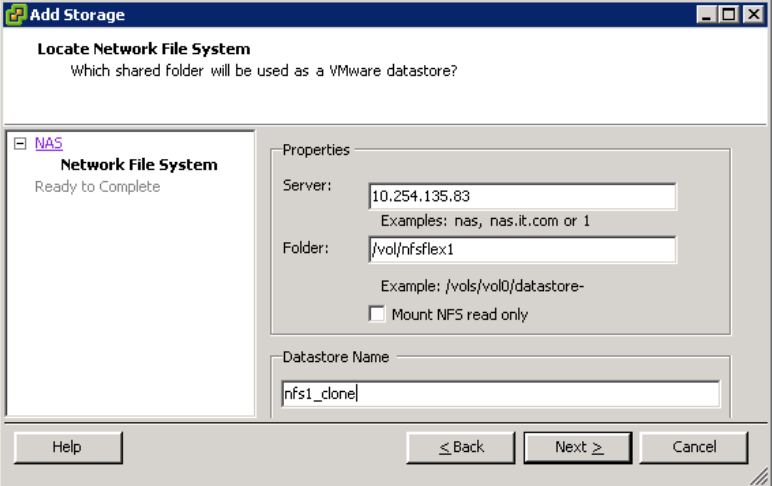


<p>6</p>	<p>Power on the <b>stu#fcsvm1</b> in the VI client</p> <ul style="list-style-type: none"> <li>• Log into the VM, browse the new drive connected to the machine (most likely drive e:)</li> <li>• Find the file created you created in the previous lab and copy it back to its original location.</li> <li>• Shut down the <b>stu#fcsvm1</b></li> <li>• In the VI client, select the <b>stu#fcsvm1</b> VM and click on <b>Edit Settings</b> under the summary tab.</li> <li>• Select the virtual disk you added earlier (it should be Hard Disk 2) and click remove.</li> <li>• Click OK to close the window</li> </ul>	
<p>7</p>	<p><b>**** Mandatory Step</b> This step must be completed or other labs will be affected. Please see your instructor if you have questions. <b>****</b></p> <p>Clean up by removing the LUN clone.</p> <ul style="list-style-type: none"> <li>• On the storage controller console run the command <b>'lun destroy -f &lt;clone LUN path&gt;'</b> using the same value for &lt;clone LUN path&gt; as when it was created.</li> </ul> <p>Use the VI client to rescan the storage on your ESX server. Confirm that the clone datastore is no longer connected to the ESX server</p>	

## Part III – Recover a file from a NFS datastore

	<p>In this exercise you will recover a individual file from a virtual disk which resides in the snapshot of the NFS datastore. You will create a FlexClone using the snapshot created in the previous labs, and then export it to your ESX server so you can connect the virtual disk to your VM.</p>
1	<p>Open FilerView to create a FlexClone of your NFS datastore volume.</p> <ul style="list-style-type: none"><li>• Under the volumes heading choose <b>FlexClones</b> and then click <b>Create</b></li><li>• When the wizard starts, choose <b>next</b></li><li>• Under Clone Name enter <b>nfsflex1</b>. The parent volume is your assigned nfs volume. Choose <b>next</b></li><li>• For the Parent volume Snapshot select the snapshot you created in the previous lab <b>nfssnap1</b></li></ul> <p><b>Note:</b> A FlexClone volume is automatically exported to the same hosts as their parent volumes, so there is no need to create an NFS export for your FlexClone volume.</p> 



<p>2 In the VI client select your ESX server and click the <b>storage</b> link on the configuration tab.</p> <ul style="list-style-type: none"> <li>Click on the <b>Add Storage</b> Link to open the Add Storage wizard.</li> <li>On the Select Storage Type page, choose <b>Network File System</b> and click <b>next</b></li> <li>On the Locate Network File System page enter the <b>IP address</b> for your NetApp system and the path to your newly created FlexCloned NFS volume. Leave the Mount NFS Read Only option unchecked, and name the datastore <b>stu#nfs_clone</b> click <b>next</b>.</li> </ul> <p>Click <b>Finish</b> and the datastore will be visible in storage</p>	
<p>3 Now you will mount the cloned virtual disk onto the production virtual machine.</p> <ul style="list-style-type: none"> <li>Open the VI client, select stu#nfsvm1 and click on <b>Edit Settings</b> under the summary tab</li> <li>In the Virtual Machine Properties window click on the <b>Add</b> button</li> <li>In the Add Hardware wizard, select hard disk and click <b>next</b></li> <li>Select “Use an existing virtual disk” radio button and click <b>next</b></li> <li>Click on <b>browse</b>, select the stu#nfs_clone datastore</li> <li>Select the nfsvm1 folder, select the nfsvm1.vmdk file and click OK.</li> <li>Keep the default values in the specific advanced options screen and click <b>next</b> then <b>finish</b>.</li> </ul>	
<p>4 Power on stu#nfsvm1, log in and browse the new drive connected to the machine.</p> <ul style="list-style-type: none"> <li>Find the file you created in the previous lab and copy it back to its original location.</li> <li>Shut down <b>stu#nfsvm1</b></li> </ul> <p>In the VI client, select nfsvm1 and click on Edit Settings.</p> <ul style="list-style-type: none"> <li>Select the virtual disk you added earlier (it should be Hard Disk 2) and click <b>remove</b>.</li> <li>Click OK.</li> </ul>	

5	<p>Clean up by removing the clone datastore and deleting the FlexClone.</p> <ul style="list-style-type: none"><li>• In the VI client, click on the <b>configuration tab</b>, then click on <b>storage</b></li><li>• Right click on the stu#nfs1_clone datastore and select <b>remove</b></li><li>• In FilerView go to <b>Volumes &gt; Manage</b> select your cloned volume click on <b>offline</b>, select it again and choose <b>destroy</b></li></ul>



## LAB 6: Virtual Machine Replication and Disaster Recovery Using SnapMirror

Last Updated: Wednesday, May 28, 2008

### **Summary:**

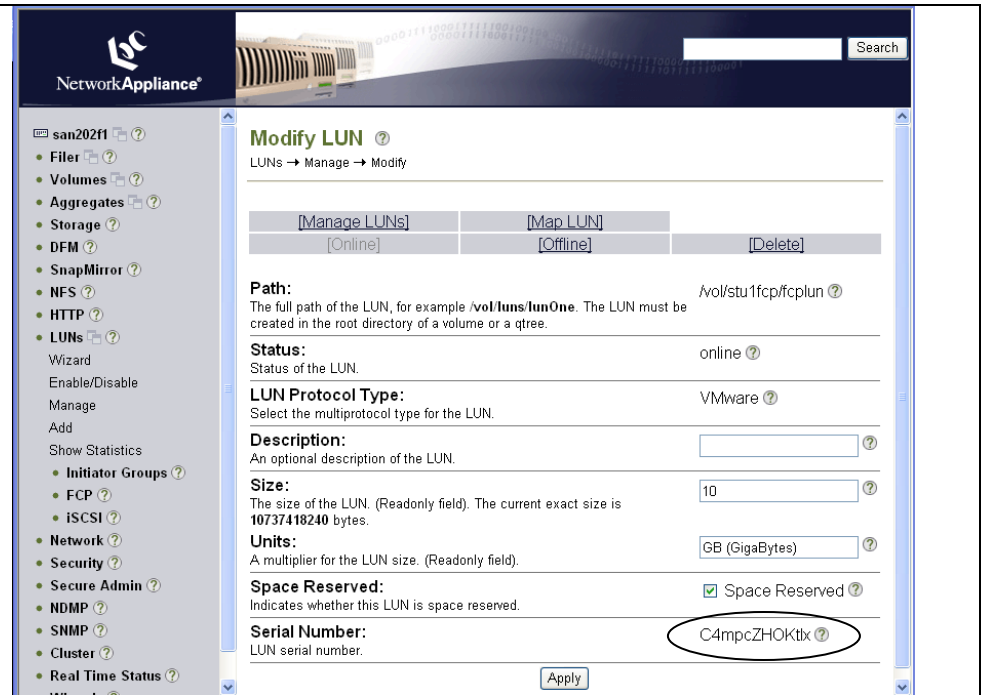
In this lab you will simulate a failure of the entire production storage environment. For the purposes of this exercise, your production and DR systems will be the same. This would not be the case in a real world scenario, but the procedures used below would not change. You will replicate Datastores using NetApp's SnapMirror technology between multiple volumes and then simulate a failure of your main system. After the failover to your Disaster Recovery site you will re-establish the link and transfer the changed data back to the production system.

## Part I – Replicate Datastores with SnapMirror

	Click Stream	Screen Shots
1	<p>In this exercise, you will replicate your two Datastores to secondary volumes</p> <p>You will open FilerView to create two volumes of equal size to your FCP and NFS volumes.</p> <ul style="list-style-type: none"> <li>In FilerView under the <b>Volumes</b> heading choose <b>Add</b> to start the Volume Wizard</li> <li>Create a Flexible Volume named <b>fcp_mir</b>, in <b>aggr1</b> and <b>25 GB</b> in size</li> <li>Create a Flexible Volume named <b>nfs_mir</b>, in <b>aggr1</b> and <b>60 GB</b> in size</li> </ul> <p>Under the <b>Manage Volumes</b> heading select each new volume and choose <b>restrict</b></p> <p>Create a SnapMirror relationship to replicate your FCP and NFS volumes and initialize the mirrors.</p> <ul style="list-style-type: none"> <li>In FilerView select the SnapMirror heading in the left pane and choose <b>add</b></li> <li>For the destination volume choose <b>fcp_mir</b></li> <li>In the source section enter your <b>storage controller name</b> and the volume name of <b>fcp</b></li> <li>Use the preset schedule and choose <b>add</b></li> <li>Complete the same steps to create the NFS mirror</li> <li>When you click on the Manage heading under SnapMirror you will see both relationships choose <b>advanced</b> operations and <b>Initialize</b> both</li> </ul> <p><b>Note:</b> This may take a few minutes to complete</p>	

2 Use FilerView to view the serial number of the LUN on the production FCP volume. To do this choose manage under LUNs then select the name of the LUN. The serial number will be at the bottom.

**Record this serial number for later use.**\_\_\_\_\_



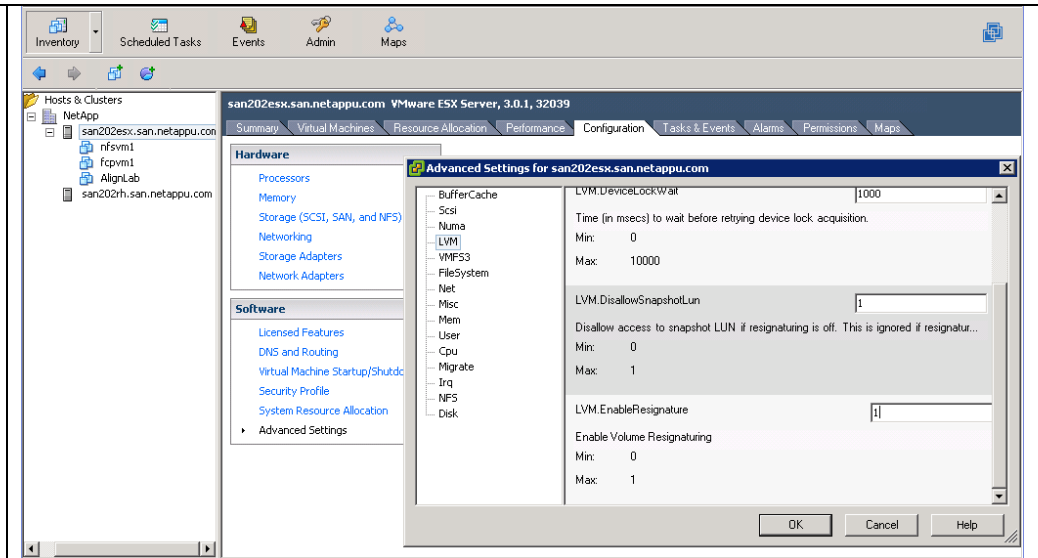
## Part II – Perform a DR failover of a VMFS datastore

This exercise will include executing commands on the service console of your ESX server. There should be a copy of putty on your VirtualCenter server. Use this or a client on your workstation to connect to your ESX server via SSH. The username and password of your ESX server has been provided by your instructor. You will also need to be able to run commands on the console of your storage controller ensure that SSH is enabled. Using FilerView go to **Secure Admin > SSH > Enable/Disable** and verify that SSH v2 is enabled.

1 **\*\*Skip this step if you completed it in a previous lab\*\***

You will be connecting to the NetApp snapshot copy of a VMFS datastore, before you can do so we need to enable the Volume Resignature option. Use the VI client to connect to your VirtualCenter server.

- Select your ESX server, click on the **configuration tab** then click on **advanced settings** under the software heading
- When the advance settings window opens Select **LVM**
- Set the value of **LVM.EnableResignature** to 1 then click **OK**



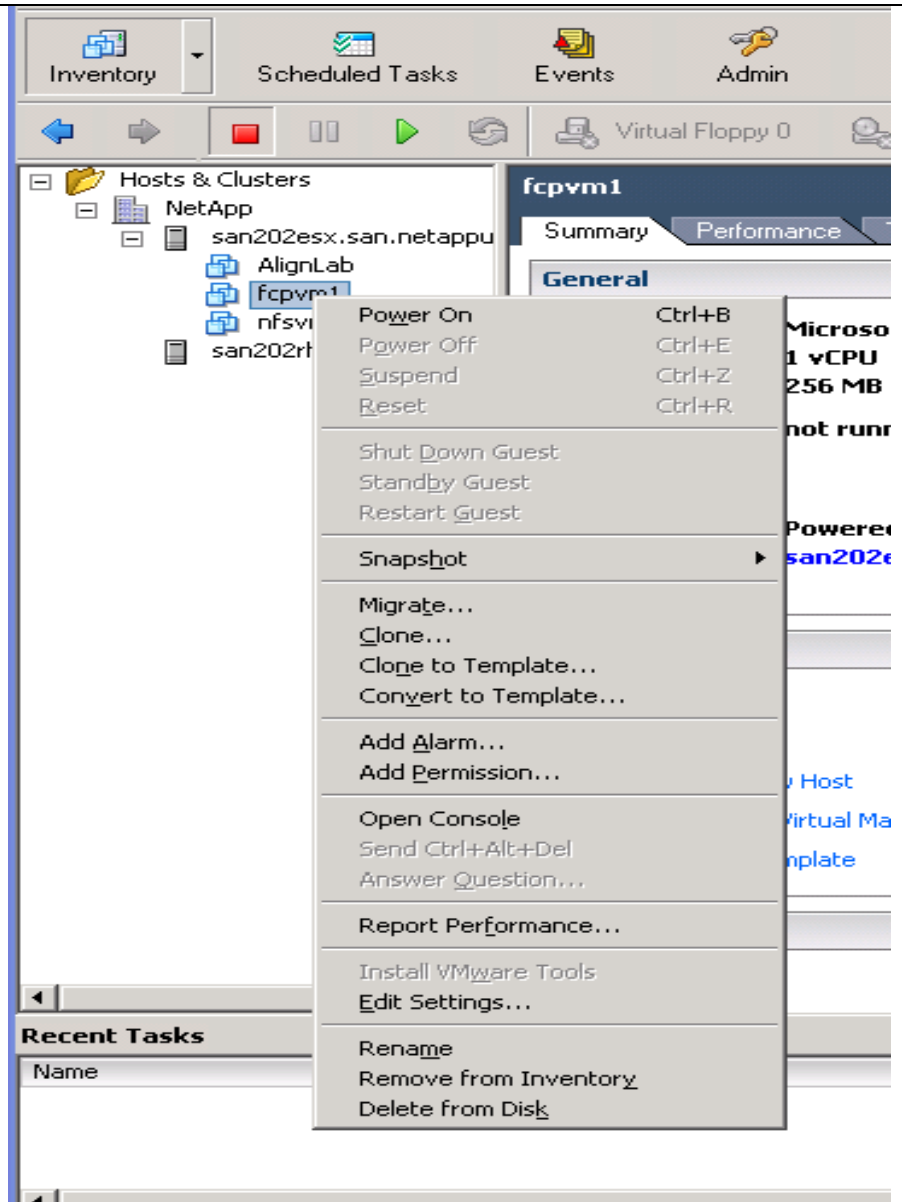
2 You will be simulating a failure of a production datastore by disconnecting the production LUN from the ESX servers.

- In the VI client, right click on **stu#fcpvm1** and select **Remove from Inventory**.
- A confirmation box will pop up, click on **yes**.
- Repeat remove from inventory steps for the **stu#alignvm1**.

In FilerView click on **LUNs > Manage** then select the **fcplun** that you created previously.

- Select the **offline** option to disconnect the LUN from the ESX server.

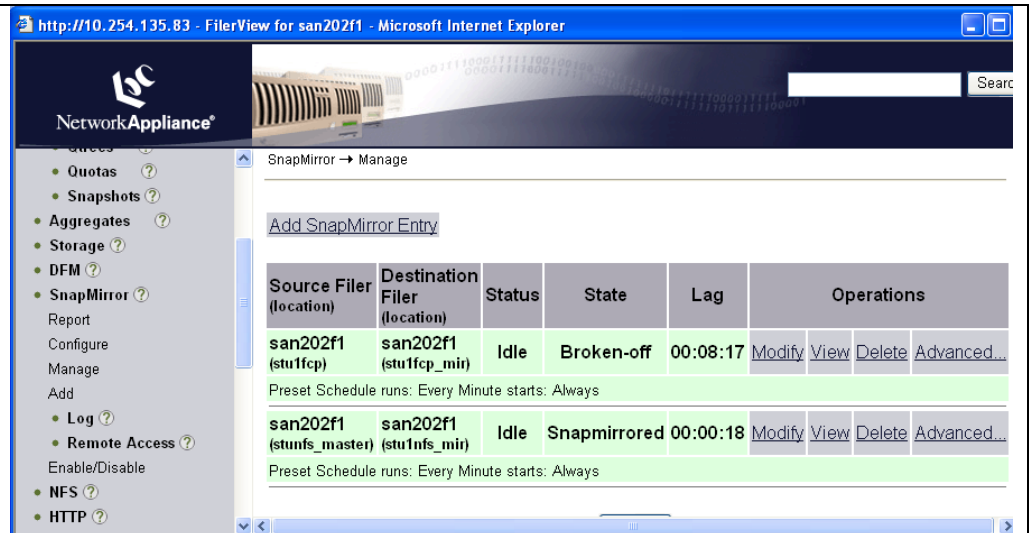
In the VI client rescan the storage on your ESX server. Confirm that the FCP LUN is no longer connected to the ESX server by clicking of the HBA's listed under storage adapters and seeing no drives available.



3 Next you will break the snapmirror relationship to make the destination writeable. Then connect your ESX server to the LUN in the mirrored volume. Then you will discover your datastore and VMs and reactivate them.

Use FilerView to break the snapmirror relationship of your FCP volume that you created earlier in this lab.

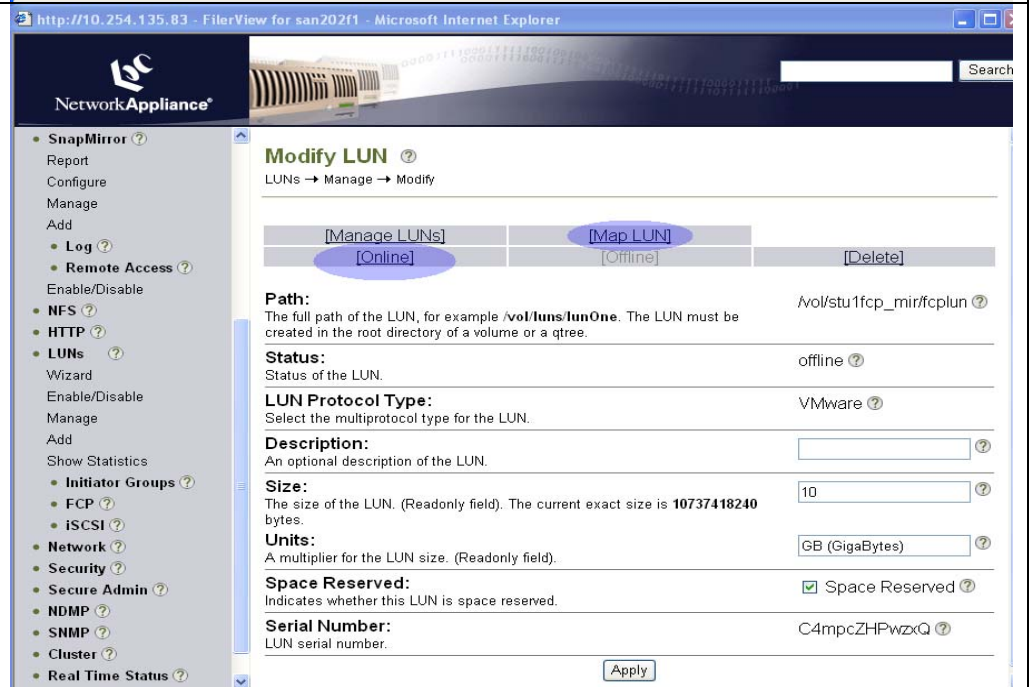
- Under **SnapMirror > Manage** select your mirror and go to **advanced options**
- Click on **Quiesce**, and then **break** the relationship



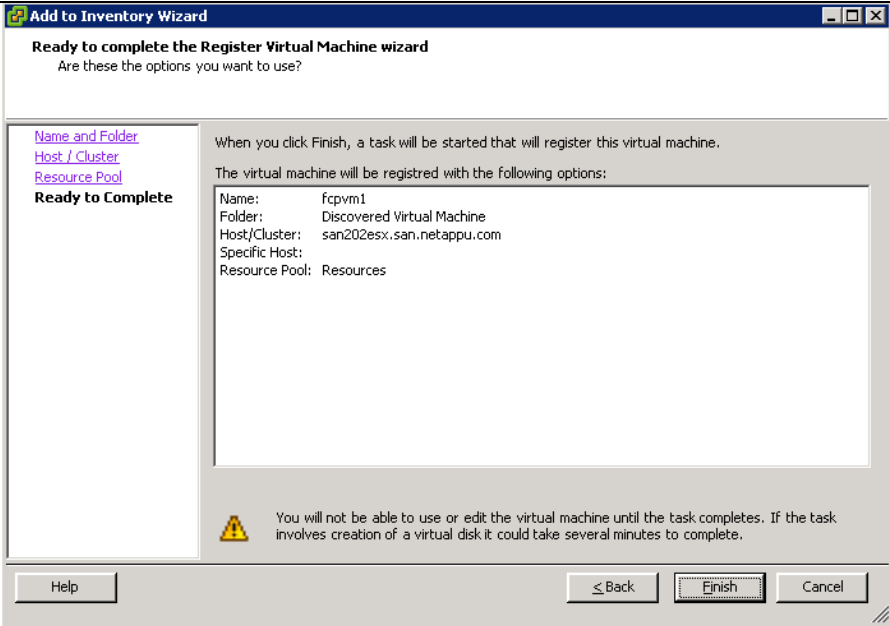
4 Under **LUNs > Manage** select the fcplun in the mirrored volume

- **Online** the LUN
- Then choose to Map the online LUN to your **stu#-fcp** igroup

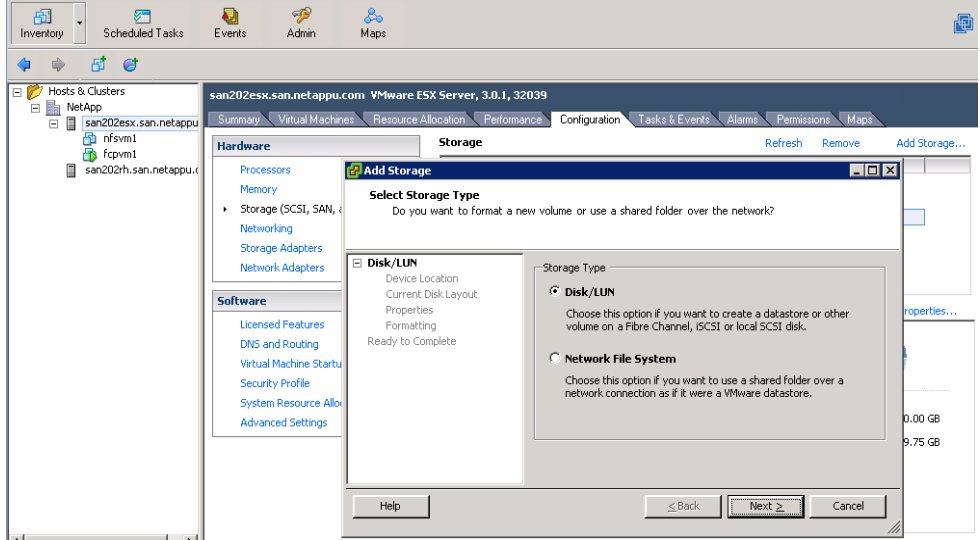
**Note:** When incrementing your LUN, use an exclusive LUN ID.

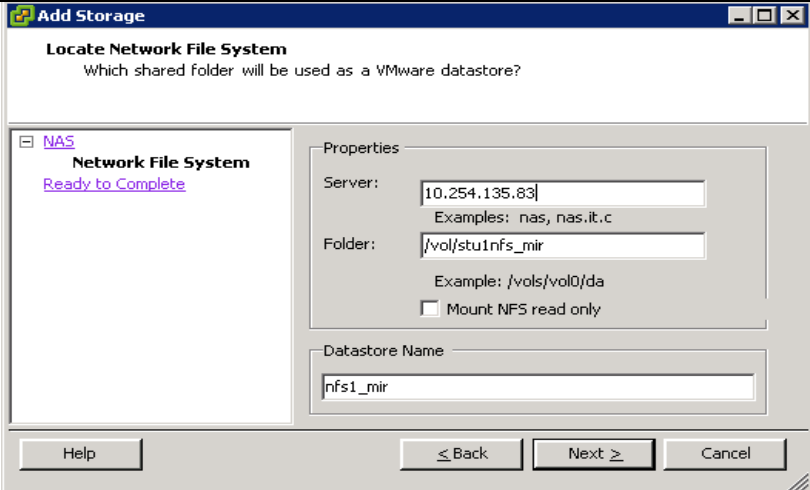




<p>5</p>	<p>Use the VI client to browse the just connected datastore by double clicking on <b>stu#fcp</b> under the storage heading.</p> <p><b>Note:</b> If the datastore is not available, go to the <b>Configuration tab &gt; adapters &gt;</b> and choose <b>rescan</b></p> <ul style="list-style-type: none"> <li>• Open the <b>stu#fcpvm1</b> folder then right click on the <b>fcpvm1.vmx</b> file and select <b>add to inventory</b></li> <li>• In the add to inventory wizard use <b>stu#fcpvm1</b> for the name and select the <b>discovered virtual machine</b> then choose <b>next</b></li> <li>• Select your ESX server and choose <b>next</b></li> <li>• Review the data and choose <b>finish</b></li> <li>• Power on <b>stu#fcpvm1</b></li> <li>• Log into the VM keep the old identifier and confirm that the file created in the previous lab is available.</li> </ul>	
<p>6</p>	<p>In a real world disaster scenario you would run in the DR environment for some time, then take a planned outage to mirror your data back to the primary data center and restore normal service. To shorten the time you will simply discard your replica and reconnect to your original datastores.</p> <ul style="list-style-type: none"> <li>• Power off <b>stu#fcpvm1</b></li> <li>• In the VI client, right click on <b>stu#fcpvm1</b> and select <b>Remove from Inventory</b>.</li> <li>• A confirmation box will pop up, click on <b>yes</b>.</li> </ul> <p>In FilerView, go to <b>LUNs &gt; Manage</b> select the LUN in the '<b>fcp_mir</b>' volume and choose <b>offline</b>.</p> <ul style="list-style-type: none"> <li>• Now go back to the LUN Manage window and <b>online</b> the original LUN in the <b>fcp</b> volume.</li> </ul> <p>Use the VI client to rescan the storage on your ESX servers.</p> <ul style="list-style-type: none"> <li>• You may need to run the process twice, once to find the LUN and a second time to discover the datastore.</li> <li>• Browse the just connected datastore, open <b>stu#fcpvm1</b> folder then right click on the <b>fcpvm1.vmx</b> file and select <b>add to inventory</b>.</li> <li>• You can use the original VM name for the wizard.</li> </ul>	

## Part III – Recover a VM from a NFS datastore

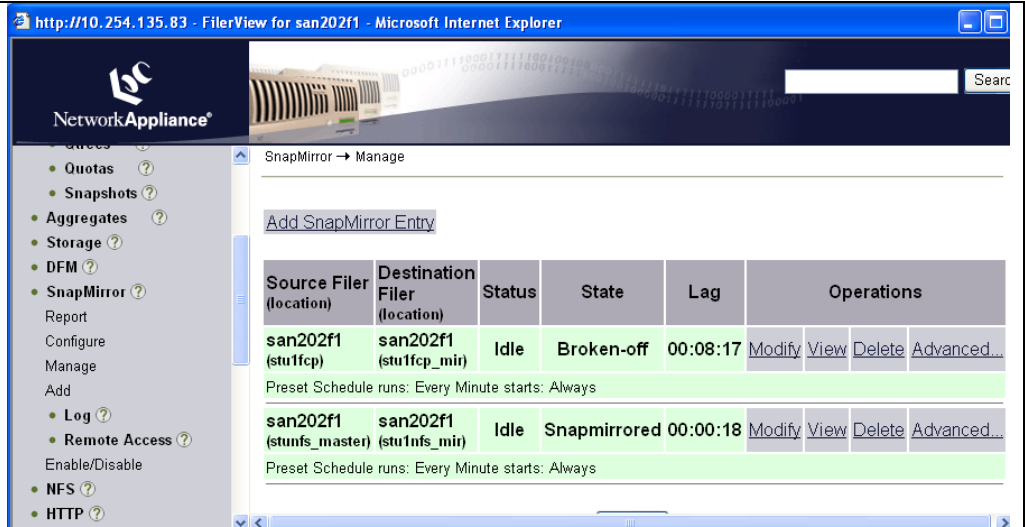
	<p>This exercise will be similar to the VMFS failover we completed in the previous section, but when using NFS datastores you can export and connect the datastore before you experience a failure or need to break the mirror. When using NFS datastores, you can connect and browse a mirror destination datastore at any time.</p>
<p>1 Use FilerView to verify that your volume has been exported</p> <ul style="list-style-type: none"><li>• Under the <b>NFS</b> option choose <b>Manage Exports</b></li><li>• The <b>/vol/nfs_mir</b> should be visible as an export. If not, then use the same export options as the source volume. **** <b>Make sure that Anon=0 is in the export options for both the source and mirror.</b> ****</li></ul> <p>To connect your NFS datastore choose your ESX server in the VI client.</p> <ul style="list-style-type: none"><li>• Click on the <b>configuration tab</b> and select the <b>Storage link</b> in the hardware section.</li><li>• Click on the <b>Add Storage</b> Link to open the Add Storage wizard.</li><li>• In the Select Storage Type page, select <b>Network File System</b> and click <b>next</b></li><li>• In the locate Network File System page you will need to provide the <b>IP address</b> of your storage controller and the path to your NFS volume, <b>/vol/nfs_mir</b></li><li>• Leave the Mount NFS Read Only option <b>unchecked</b>.</li><li>• Use the datastore name of <b>stu#nfs_mir</b> click <b>next</b></li><li>• Click <b>Finish</b> and verify that your NFS datastore is connected (you should see it appear on the storage list immediately)</li></ul>	

	
2	<p>You will be simulating a failure of a production datastore by disconnecting the VM and datastore.</p> <ul style="list-style-type: none"> <li>• In the VI client, right click on <b>stu#nfsvm1</b> and select <b>Remove from Inventory</b>.</li> <li>• A confirmation box will pop up, click on <b>yes</b>.</li> <li>• In the VI client, click on the <b>configuration tab</b>, then click on <b>storage</b>.</li> <li>• Right click on the <b>stu#nfs</b> datastore and select <b>remove</b>.</li> </ul>

3 Next you will break the SnapMirror relationship to make the destination writeable. Then you will discover your VM and reactivate them.

Use FilerView to break the SnapMirror relationship of your nfs volume that you created earlier in this lab.

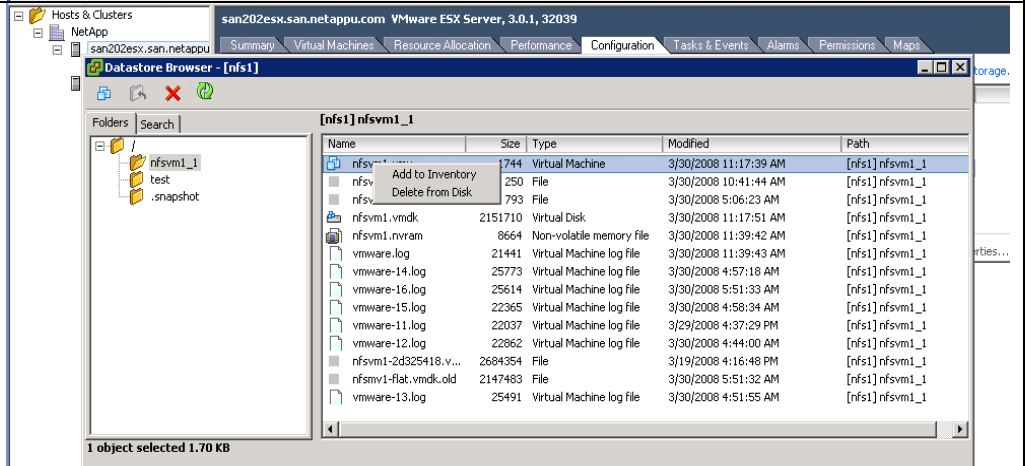
- Under **SnapMirror > Manage** select your mirror and go to **advanced options**
- Click on **Quiesce**, and then **break** the relationship



4 Now use the VI client to browse the just added **stu#nfs\_mir** datastore by double clicking on it

- Open the **stu#nfsvm1** folder then right click on the **nfsvm1.vmx** file and select **add to inventory**.
- You can use the original VM name **stu#nfsvm1** in the wizard, choose your assigned ESX server and accept the defaults for the rest of the screens and choose **finish** when complete.

- The VM will appear on the list
  - Power on **stu#nfsvm1**
- Log into the VM and confirm that the file created in the previous lab is available



5 In a real world disaster scenario you would run in the DR environment for some time, then take a planned outage to mirror your data back to the primary data center and restore normal service. To shorten the time you will simply discard your replica and reconnect to your original Datastores.

- Power off **stu#nfsvm1**
- In the VI client, right click on **stu#nfsvm1** and select **Remove from Inventory**.

- A confirmation box will pop up, click on **yes**.
- Click on the **configuration tab**, then click on **storage**.
- Right click on the **stu#nfs\_mir** datastore and select **remove**.

To reconnect your original NFS datastore, click on the **Add Storage Link** to open the Add Storage wizard.

- In the **Select Storage Type** page, select **Network File System** and click **Next**.
- In the **locate Network File System** page you will need to provide the **IP** address for your storage controller and the path to your original **NFS volume**.
- Leave the **Mount NFS Read Only** option **unchecked**.
- Use the original datastore name **stu#nfs**.
- Click **Finish** and verify that your NFS datastore is connected.
- Browse the just connected datastore, open the **stu#nfsvm1** folder then right click on the **nfsvm1.vmx** file and select **add to inventory**. You can use the original VM name of **stu#nfsvm1**





## LAB 7: FAS Deduplication of NFS and VMware

### Summary:

In this lab you will enable FAS deduplication on a datastore and observe the effect on storage consumption. FAS deduplication is effective in VMware environments because there can be large amounts of redundant data between two or more virtual machines in a datastore. You will be creating redundant data, and then executing FAS deduplication to eliminate the redundancy and allow the space to be better utilized.

## Part I –Create duplicate data

	Click Stream	Screen Shots																								
1	<p>In the VI client connect to your VirtualCenter server and select your ESX server.</p> <ul style="list-style-type: none"> <li>In the <b>Summary</b> tab under the <b>Resources</b> window, note the capacity and free space of the <b>stu#nfs</b> datastore.</li> </ul> <hr/> <p>In FilerView under the <b>Volumes &gt; Manage</b> heading you can see what is reported for volume usage and it will be the same as the ESX server.</p> <p>In the VI client, right click on the <b>stu#nfsvm1</b> virtual machine and click on <b>clone</b>.</p> <ul style="list-style-type: none"> <li>In the Clone VM wizard name the new <b>stu#nfsvm2</b></li> <li>For the virtual machine inventory location, choose the datacenter and click <b>next</b></li> <li>Select your ESX server and click <b>next</b>.</li> <li>Select the <b>stu#nfs</b> datastore and click <b>next</b>.</li> <li>Select <b>do not customize</b> and click <b>next</b>.</li> <li>Click <b>finish</b></li> </ul> <p><b>Note:</b> It will take several minutes for the clone to complete. Status can be viewed in the Recent Tasks bar at the bottom of the VI client</p> <p>Repeat the previous steps to create multiple VMs.</p>	<div data-bbox="1201 318 1797 748"> <p><b>Resources</b></p> <p>CPU usage: <b>31 MHz</b>  <small>4 x 2.992 GHz</small></p> <p>Memory usage: <b>563.00 M</b>  <small>2.00 GB</small></p> <table border="1"> <thead> <tr> <th>Datstore</th> <th>Capacity</th> <th>Free</th> </tr> </thead> <tbody> <tr> <td>iscsi1</td> <td>768.00 MB</td> <td>142.00 MB</td> </tr> <tr> <td>fcp1</td> <td>9.75 GB</td> <td>5.89 GB</td> </tr> <tr> <td>nfs1</td> <td>48.00 GB</td> <td>43.59 GB</td> </tr> </tbody> </table> <p>Network            VM Network</p> </div> <div data-bbox="1073 792 1927 1383"> <p><b>Clone Virtual Machine Wizard</b></p> <p>Ready to Complete New Virtual Machine</p> <p>Are these the options you want to use?</p> <p>When you click Finish, a task will be started that will create the new virtual machine.</p> <p>The virtual machine will be created with the following options:</p> <table border="1"> <tr> <td>Virtual Machine to Clone:</td> <td>nfsvm1</td> </tr> <tr> <td>Name:</td> <td>nfsvm2</td> </tr> <tr> <td>Folder:</td> <td>vm</td> </tr> <tr> <td>Host/Cluster:</td> <td>san202esx.san.netappu.com</td> </tr> <tr> <td>Datstore:</td> <td>nfs1</td> </tr> <tr> <td>Guest OS Customization Specification:</td> <td>None, do not customize guest OS</td> </tr> </table> <p><input type="checkbox"/> Power on the new virtual machine after creation</p> <p> You will not be able to use or edit the virtual machine until the task completes. If the task involves creation of a virtual disk it could take several minutes to complete.</p> <p>Help    &lt; Back    Finish    Cancel</p> </div>	Datstore	Capacity	Free	iscsi1	768.00 MB	142.00 MB	fcp1	9.75 GB	5.89 GB	nfs1	48.00 GB	43.59 GB	Virtual Machine to Clone:	nfsvm1	Name:	nfsvm2	Folder:	vm	Host/Cluster:	san202esx.san.netappu.com	Datstore:	nfs1	Guest OS Customization Specification:	None, do not customize guest OS
Datstore	Capacity	Free																								
iscsi1	768.00 MB	142.00 MB																								
fcp1	9.75 GB	5.89 GB																								
nfs1	48.00 GB	43.59 GB																								
Virtual Machine to Clone:	nfsvm1																									
Name:	nfsvm2																									
Folder:	vm																									
Host/Cluster:	san202esx.san.netappu.com																									
Datstore:	nfs1																									
Guest OS Customization Specification:	None, do not customize guest OS																									



## Part II – Enable FAS Deduplication in an NFS Datastore

	Click Stream	GUI View
	<p>Prior to beginning the deduplication process take note of the capacity and free space that is listed in the VI Client, and the Storage controller volume for the NFS datastore after cloning the VMs.</p> <p><b>stu#nfs</b> datastore (under resources) Capacity _____ Free _____</p> <p><b>nfs</b> (NetApp volume) Capacity _____ Free _____</p>	
<p>1 Using putty connect to your Storage Controller and enable FAS Deduplication on your NFS datastore volume</p> <ul style="list-style-type: none"> <li>To activate deduplication use '<b>sis on /vol/nfs</b>'</li> <li>Start a FAS deduplication scan with the command '<b>sis start -sf /vol/nfs</b>'</li> </ul> <p><b>NOTE:</b> Deduplication can take some time. You may want to move on to other activities and complete the following steps after some time has passed.</p> <ul style="list-style-type: none"> <li>As the Storage Controller de-duplicates the volume and recovers free space, use the '<b>df -g</b>', '<b>df -s</b>' and '<b>sis status</b>' commands or FilerView to observe the changes in used and free space.</li> <li>In the VI client check the datastore in the summary tab. Note that datastore free space is affected by the FAS deduplication operations.</li> </ul>	<pre> sis on /vol/stunfs_master SIS for "/vol/stunfs_master" is enabled. Already existing data could be processed by running "sis start -s /vol/stunfs_ma ster". san202f1&gt; sis start -fs /vol/stunfs_master Mon Mar 31 08:08:46 PDT [san202f1: wafl.scan.start:info]: Starting SIS volume sc an on volume stunfs_master. The SIS operation for "/vol/stunfs_master" is started. san202f1&gt; </pre>	
	<p>Prior to beginning the deduplication process take note of the capacity and free space that is listed in the VI Client, and the Storage controller volume for the NFS datastore after cloning the VMs.</p> <p><b>stu#nfs</b> datastore (under resources) Capacity _____ Free _____</p> <p><b>nfs</b> (NetApp volume) Capacity _____ Free _____</p>	



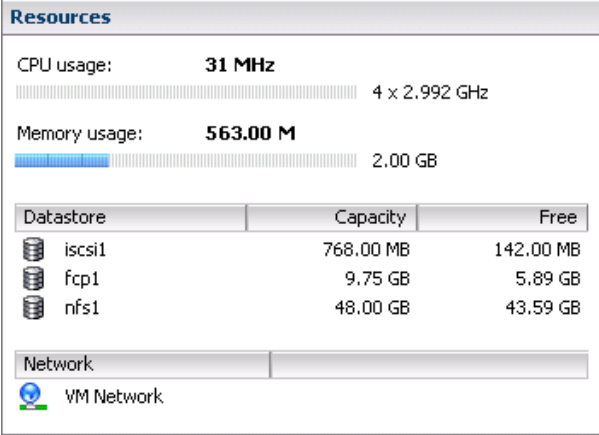


# LAB 8: Thin Provisioning & FAS Deduplication of VMFS Datastores

## Summary:

In this lab you will enable FAS deduplication on a VMFS datastore and observe the effect on storage consumption. Unlike a NFS datastore, in order for a VMFS datastore to take advantage of the space savings of FAS deduplication, the LUN on which the datastore resides must be thin provisioned. You will be creating redundant data, and then executing FAS deduplication to eliminate the redundancy and allow the space to be better utilized.

## Part I –Configure Thin Provisioning & Create Duplicate Data

	Click Stream	Screen Shots												
1	<p>In the VI client connect to your VirtualCenter server and select your ESX server.</p> <p>In the <b>Summary</b> tab under the <b>Resources</b> window, note the capacity and free space of the <b>stu#fcp</b> datastore.</p> <hr/> <p>Connect to your storage controller using putty and view the volume usage of your FCP datastore with '<b>df -g fcp</b>' Note the capacity and free space.</p> <hr/>	 <p>The screenshot shows the 'Resources' window in VMware vSphere. It displays CPU usage at 31 MHz (4 x 2.992 GHz) and Memory usage at 563.00 M (2.00 GB). Below this is a table of Datastore Capacity and Free space:</p> <table border="1"> <thead> <tr> <th>Datastore</th> <th>Capacity</th> <th>Free</th> </tr> </thead> <tbody> <tr> <td>iscsi1</td> <td>768.00 MB</td> <td>142.00 MB</td> </tr> <tr> <td>fcp1</td> <td>9.75 GB</td> <td>5.89 GB</td> </tr> <tr> <td>nfs1</td> <td>48.00 GB</td> <td>43.59 GB</td> </tr> </tbody> </table> <p>At the bottom, the Network section shows 'VM Network'.</p>	Datastore	Capacity	Free	iscsi1	768.00 MB	142.00 MB	fcp1	9.75 GB	5.89 GB	nfs1	48.00 GB	43.59 GB
Datastore	Capacity	Free												
iscsi1	768.00 MB	142.00 MB												
fcp1	9.75 GB	5.89 GB												
nfs1	48.00 GB	43.59 GB												

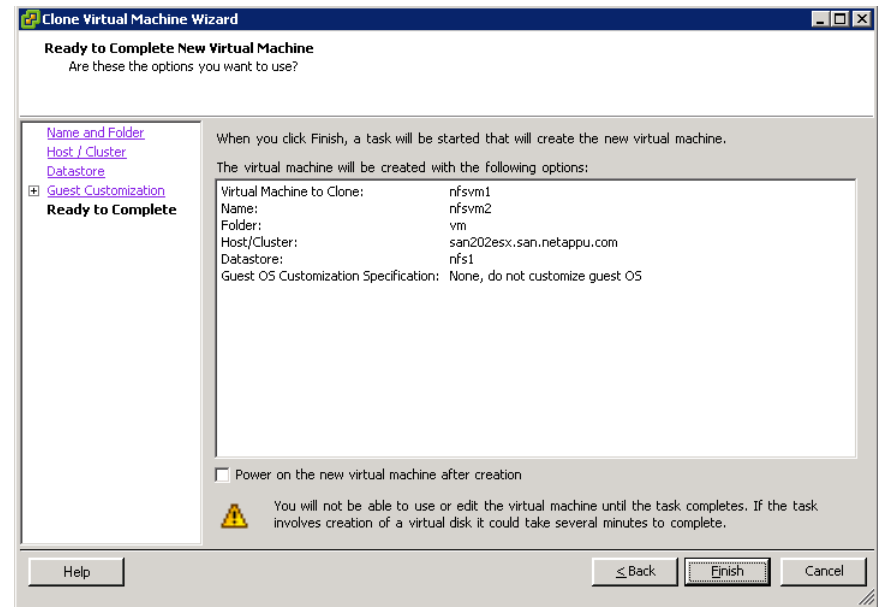
2 Disable the LUN reservation for your FCP LUN with the lun set reservation command. '**lun set reservation /vol/fcp/fcplun disable**'.

In the VI client, right click on the **stu#fcpvm1** virtual machine and click on **clone**.

- In the Clone VM wizard name the new **stu#fcpvm2**
- For the virtual machine inventory location, choose the datacenter and click **next**
- Select your ESX server and click **next**.
- Select the **stu#fcp** datastore and click **next**.
- Select **do not customize** and click **next**.
- Click **finish**

**Note:** It will take several minutes for the clone to complete. Status can be viewed in the **Recent Tasks** bar at the bottom of the VI client

Follow the previous steps to create multiple clones of the VM.



## Part II – Enable FAS Deduplication in the VMFS Datastore

	Click Stream	GUI View
1	<p>Prior to beginning the deduplication process take note of the capacity and free space that is listed in the VI Client, and the Storage controller volume for the FCP datastore after cloning the VMs.</p> <p><b>stu#fcp</b> datastore (under resources) Capacity _____ Free _____</p> <p><b>fcp</b> (NetApp volume) Capacity _____ Free _____</p>	
2	<p>Using putty connect to your Storage Controller and enable FAS Deduplication on your FCP datastore volume</p> <ul style="list-style-type: none"> <li>• To activate deduplication use '<b>sis on /vol/fcp</b>'</li> <li>• Start a FAS deduplication scan with the command '<b>sis start -sf /vol/fcp</b>'</li> </ul> <p><b>NOTE:</b> Deduplication can take some time. You may want to move on to other activities and complete the following steps after some time has passed.</p> <ul style="list-style-type: none"> <li>• As the Storage Controller deduplicates the volume and recovers free space, use the '<b>df -g</b>', '<b>df -s</b>' and '<b>sis status</b>' commands or FilerView to observe the changes in used and free space.</li> <li>• In the VI client, check the fcp datastore in the summary tab.</li> </ul> <p><b>Note:</b> There is no change of the datastore free space within VMware, but the space savings effects of FAS deduplication will be seen in the volume on the storage controller.</p>	<pre> san202f1&gt; san202f1&gt; sis on /vol/stulfcp SIS for "/vol/stulfcp" is enabled. Already existing data could be processed by running "sis start -s /vol/stulfcp". san202f1&gt; sis start -sf /vol/stulfcp Mon Mar 31 12:50:40 PDT [san202f1: waf1.scan.start:info]: Starting SIS volume scan on volume stulfcp. The SIS operation for "/vol/stulfcp" is started. san202f1&gt; █ </pre>







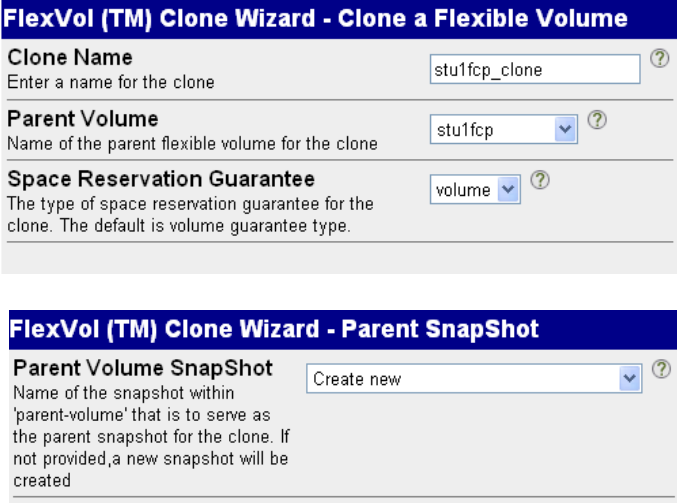


## LAB 9: Provisioning Datastores & VMs with FlexClone

### **Summary:**

Provisioning new Virtual Machines by cloning existing ones using VMware technology can be time consuming and generate a great deal of load on your ESX server and storage device. In this lab you will use FlexClone to rapidly provision new Datastores and virtual machines.

## Part I –Create a FlexClone of an Existing Datastore

	Click Stream	Screen Shots
1	<p>Use FilerView to create a FlexClone of your FCP datastore volume.</p> <ul style="list-style-type: none"> <li>In FilerView under <b>Volumes &gt; FlexClones</b> choose <b>create</b></li> <li>In the wizard use the name of <b>fcp_clone</b></li> <li>The parent volume should be <b>fcp</b> click <b>next</b></li> <li>Create a new Parent Volume SnapShot</li> </ul> <p>Online and map the newly cloned LUN.</p> <ul style="list-style-type: none"> <li>Choose <b>LUNs &gt; Manage</b> to view the cloned LUN</li> <li>Click on the link to the <b>LUN</b> and choose <b>Online</b> at the top of the page</li> <li>Choose the <b>Map LUN</b> option then <b>add groups</b>. Add the <b>fcp</b> group you created earlier in the labs</li> <li><b>Increment</b> the LUN id so it is not the same as the production LUN</li> </ul>	 <p><b>FlexVol (TM) Clone Wizard - Clone a Flexible Volume</b></p> <p><b>Clone Name</b> Enter a name for the clone <input type="text" value="stu1fcp_clone"/></p> <p><b>Parent Volume</b> Name of the parent flexible volume for the clone <input type="text" value="stu1fcp"/></p> <p><b>Space Reservation Guarantee</b> The type of space reservation guarantee for the clone. The default is volume guarantee type. <input type="text" value="volume"/></p> <p><b>FlexVol (TM) Clone Wizard - Parent SnapShot</b></p> <p><b>Parent Volume SnapShot</b> <input type="text" value="Create new"/></p> <p>Name of the snapshot within 'parent-volume' that is to serve as the parent snapshot for the clone. If not provided, a new snapshot will be created</p>

2

Now that the LUN has been cloned, mount the new datastore and VMs

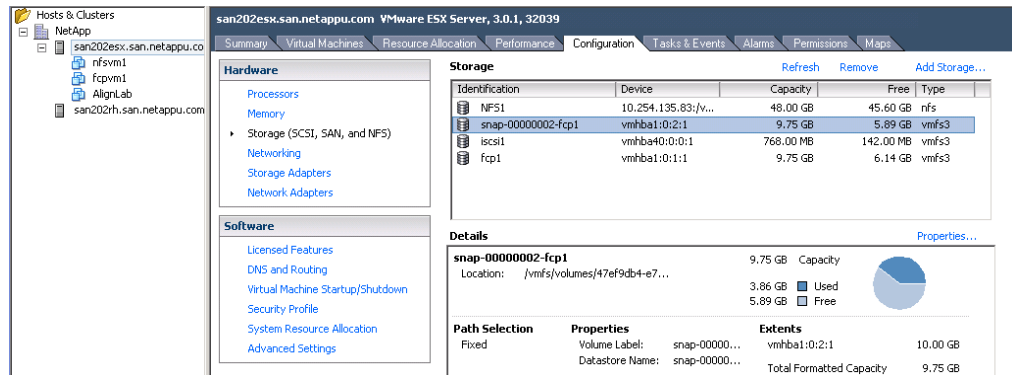
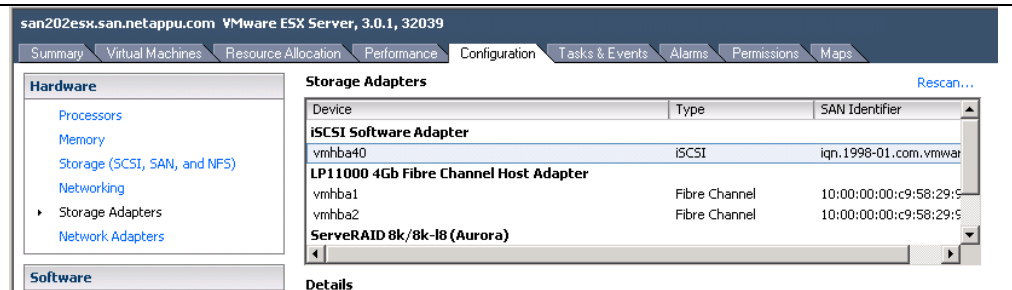
- In the VI client select your **ESX server** and click on the **Configuration** tab.
- Under **hardware** choose storage adapters and **rescan** the storage.

**NOTE:** You may need to run the process twice, once to find the LUN and a second time to discover the datastore.

To verify that the datastore has been discovered go to the **storage** heading under **hardware**.

- The datastore it will automatically be renamed to something different than the production datastore (it should be something like snap-00000001-FCP1).
- Right click on the datastore and rename it to **stu#fcp2**.
- Right click on the stu#fcp2 datastore and select **Browse Datastore**
- Open the **stu#fcpvm1** directory, right click on the **fcpvm1 .vmx** file and select **add to inventory**.
- In the window that opens, you will be asked to name the VM, for this exercise call the VM **stu#fcpvm3**
- For the virtual machine inventory location, choose your datacenter and click **next**.

You have now created a VM replica that is running on a zero space cloned LUN.

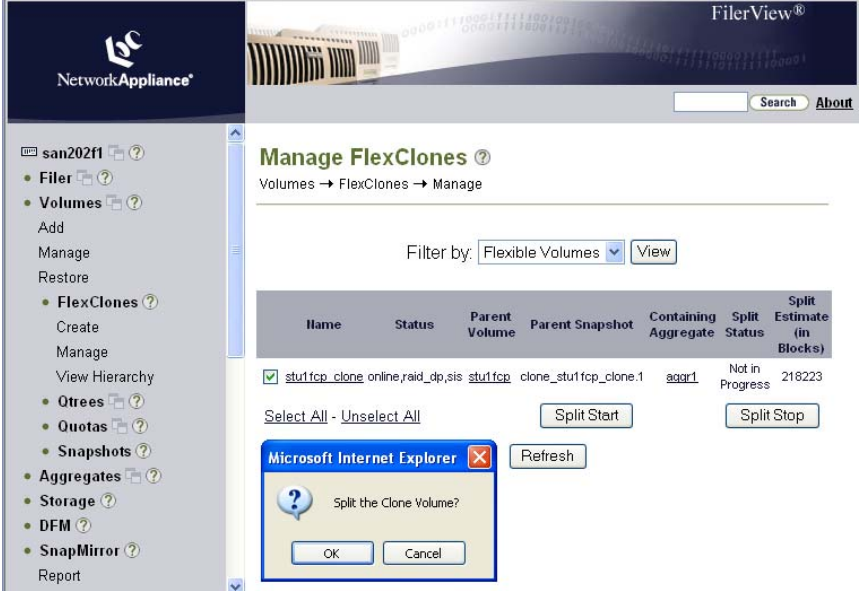


3

VMware on NetApp Solutions: LAB 9: Provisioning Datastores & VMs with FlexClone

© 2008 NetApp. This material is intended for training use only. Not authorized for reproduction purposes.

## Part II – Split FlexClone to Make the Datastore Permanent

	Click Stream	GUI View														
1	The FlexClone is still tied to a parent volume and its snapshot. If you want to use it as a permanent resource it should be split into its own permanent volume. You will see how easy it is to do that in the following exercise.															
2	<p>In FilerView go to <b>Volumes &gt; FlexClones &gt; Manage</b></p> <ul style="list-style-type: none"> <li>• <b>Check the box</b> next to the clone you created in the previous exercise and choose <b>start split</b>.</li> <li>• Choose <b>OK</b> to split the clone</li> </ul> <p><b>NOTE:</b> This operation may take a few minutes depending on the amount of data in the volume.</p> <p>When the FlexClone split operation is complete, delete the snapshot used to create the clone. The snapshot will be listed under the parent volume and contain clone in the name.</p>	 <p>The screenshot shows the NetApp FilerView interface. On the left is a navigation tree with 'FlexClones' selected. The main area displays 'Manage FlexClones' with a table of clones. One clone is selected, and the 'Split Start' button is highlighted. A dialog box from Microsoft Internet Explorer is open, asking 'Split the Clone Volume?' with 'OK' and 'Cancel' buttons.</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Status</th> <th>Parent Volume</th> <th>Parent Snapshot</th> <th>Containing Aggregate</th> <th>Split Status</th> <th>Split Estimate (in Blocks)</th> </tr> </thead> <tbody> <tr> <td>stuf1fcp_clone</td> <td>online_raid_dp_sis</td> <td>stuf1fcp</td> <td>clone_stuf1fcp_clone.1</td> <td>aggr1</td> <td>Not in Progress</td> <td>218223</td> </tr> </tbody> </table>	Name	Status	Parent Volume	Parent Snapshot	Containing Aggregate	Split Status	Split Estimate (in Blocks)	stuf1fcp_clone	online_raid_dp_sis	stuf1fcp	clone_stuf1fcp_clone.1	aggr1	Not in Progress	218223
Name	Status	Parent Volume	Parent Snapshot	Containing Aggregate	Split Status	Split Estimate (in Blocks)										
stuf1fcp_clone	online_raid_dp_sis	stuf1fcp	clone_stuf1fcp_clone.1	aggr1	Not in Progress	218223										



## LAB 10: Growing Datastores

### **Summary:**

In this lab you will see how you can increase the size of your datastores on a NetApp system simply and easily.

## Part I –Growing a VMFS Datastore

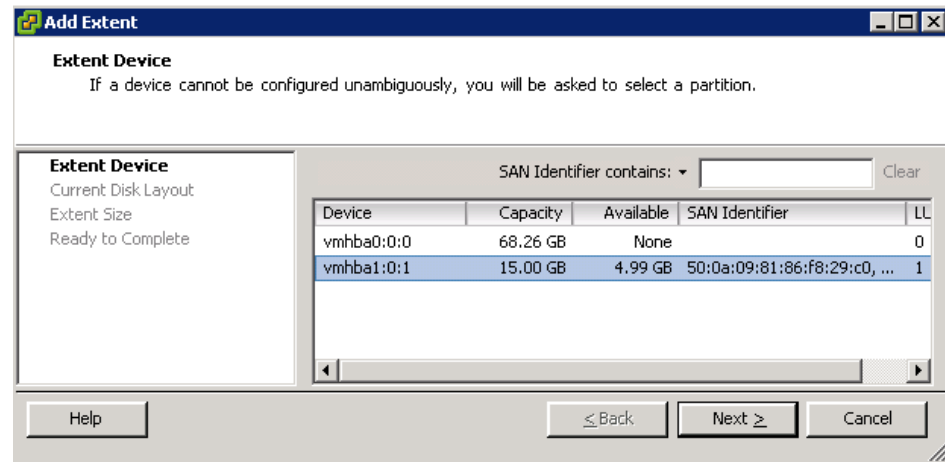
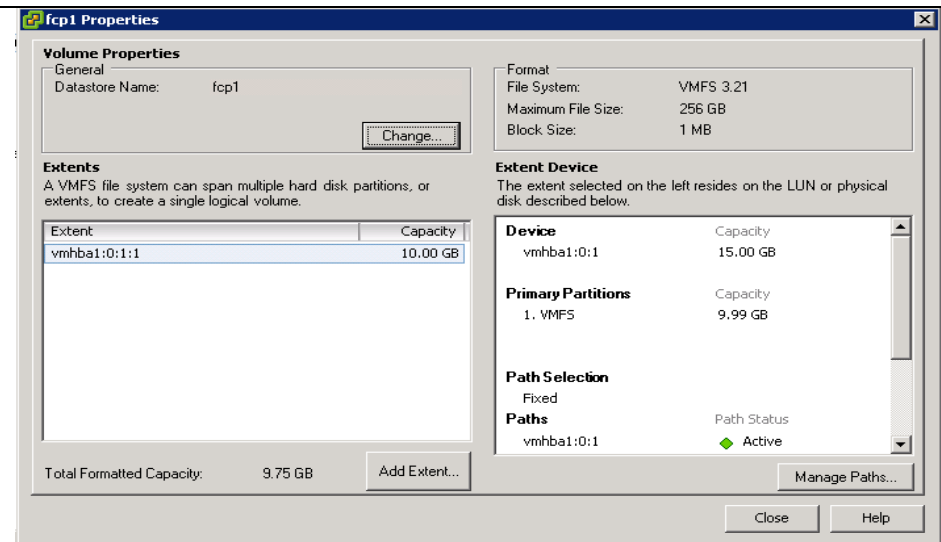
	Click Stream	Screen Shots
1	<p>You can grow an existing datastore in two ways, either by creating a new LUN or by resizing an existing LUN. Whichever you choose, the new space will be added to the datastore as an extent.</p> <p>Use FilerView to increase the size of your FCP LUN by 5GB.</p> <p><b>NOTE:</b> FilerView only allows this in Ontap 7.2.4 and above. If you are running a older version, use the command line as noted below.</p> <ul style="list-style-type: none"> <li>Go to <b>LUNs &gt; Manage</b> and click on your FCP LUN path</li> <li>In the Modify LUN window change the size field and choose apply</li> </ul> <p><b>OR:</b> you can also make this change from the command line using '<b>lun resize -f /vol/fcp/fcplun 15G</b>'</p> <p>Use the VI client to rescan the LUN on your ESX server.</p> <ul style="list-style-type: none"> <li>Select your <b>ESX</b> server and click on the <b>configuration</b> tab.</li> <li>In the <b>Hardware</b> section, click on the <b>Storage Adapters</b> link.</li> <li>Click on the <b>rescan</b> link and wait for the ESX server to rescan its adapters. You should now be able to see that ESX has detected the new size of your LUN.</li> </ul>	<div style="text-align: center;"> <h3>Modify LUN <span>?</span></h3> <p>LUNs → Manage → Modify</p> </div> <hr/> <div style="display: flex; justify-content: space-around; border-bottom: 1px solid #ccc;"> <span>[Manage LUNs]</span> <span>[Map LUN]</span> </div> <div style="display: flex; justify-content: space-around; border-bottom: 1px solid #ccc;"> <span>[Online]</span> <span>[Offline]</span> <span>[Delete]</span> </div> <p><b>Path:</b> /vol/stu1fcp/fcplun <span>?</span>  <small>The full path of the LUN, for example /vol/luns/lunOne. The LUN must be created in the root directory of a volume or a qtree.</small></p> <hr/> <p><b>Status:</b> online <span>?</span>  <small>Status of the LUN.</small></p> <hr/> <p><b>LUN Protocol Type:</b> VMware <span>?</span>  <small>Select the multiprotocol type for the LUN.</small></p> <hr/> <p><b>Description:</b> <input type="text"/> <span>?</span>  <small>An optional description of the LUN.</small></p> <hr/> <p><b>Size:</b> <input type="text" value="15"/> <span>?</span>  <small>The size of the LUN. (Readonly field). The current exact size is 16106127360 bytes.</small></p> <hr/> <p><b>Units:</b> <input type="text" value="GB (GigaBytes)"/> <span>?</span>  <small>A multiplier for the LUN size. (Readonly field).</small></p> <hr/> <p><b>Space Reserved:</b> <input checked="" type="checkbox"/> Space Reserved <span>?</span>  <small>Indicates whether this LUN is space reserved.</small></p> <hr/> <p><b>Serial Number:</b> C4mpcZHOKtx <span>?</span>  <small>LUN serial number.</small></p> <div style="text-align: right; margin-top: 10px;"> <input type="button" value="Apply"/> </div>





Now increase the size of the VMFS datastore by adding the extent.

- In the VI client click on **storage** under the **Configuration** tab

**Note** the size of the **stu#fcp** datastore.

- Select the **stu#fcp1** datastore and click on **Properties**.
- Click on 'Add Extent'
- In the Extent Device window, select the **15 GB device**.  
NOTE: You may have multiple devices in this list. The proper device will have 5GB in free space.
- Click **next**.  
NOTE: The warning message is referring to the new storage area, not the datastore being extended.
- Click **next**.
- Verify that **Maximize capacity** is checked and choose **next**
- Click **Finish**, when the operation is complete click **Close**.
- Note the new size of the **stu#fcp1** datastore.



Storage						<a href="#">Refresh</a>	<a href="#">Remove</a>	<a href="#">Add Storage..</a>
Identification	Device	Capacity	Free	Type				
 nfs1	10.254.135.83;/v...	48.00 GB	46.58 GB	nfs				
 iscsi1	vmhba40:0:0:1	768.00 MB	142.00 MB	vmfs3				
 fcp1	vmhba1:0:1:1	14.50 GB	10.64 GB	vmfs3				
 fcp2	vmhba1:0:3:1	9.75 GB	5.89 GB	vmfs3				



## Part II – Growing a NFS Datastore

Click Stream	
1	<p>Resizing an NFS datastore is even simpler than VMFS. You simply need to resize the volume, and refresh the view of the datastore in the VI client, and it is available.</p> <p>Open the VI client and connect to your ESX server.</p> <ul style="list-style-type: none"><li>• Select your <b>ESX</b> server and click on the <b>configuration tab</b>.</li><li>• Click on <b>storage</b>.</li></ul> <p>Note the size of the <b>stu#nfs</b> datastore. _____</p> <ul style="list-style-type: none"><li>• Use FilerView or the 'vol size' command to increase the size of your <b>NFS</b> volume by <b>5GB</b>.</li><li>• Go back to the storage screen of the VI client, right click on the <b>stu#nfs</b> datastore and select refresh.</li></ul> <p>Note the new size of the <b>stu#nfs</b> datastore _____</p>

