



Network Appliance and VMware ESX Server 2.5.x

Building a Virtual Infrastructure from Server to Storage

Vaughn Stewart and Michael Slisinger, Network Appliance, Inc.

May 2006, TR-3482

Abstract

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

Executive Summary

In recent years just about every company with an Information Systems department has begun some form of consolidation and virtualization effort with a goal of increasing asset utilization and reducing management and infrastructure costs. The virtualization marketplace is filled with solutions from just about every traditional vendor and a bevy of startups, but the company universally acknowledged as a leader in the virtualization space is VMware.

With the VMware ESX server, companies are able to create virtual infrastructures that allow for a separation between physical hardware and business applications. VMware deployments commonly provide administrators with a means to consolidate 8 to 12 physical servers to a single CPU on an ESX server. The savings in this consolidation provide a nearly instantaneous return on investment by lowering the total number of physical servers, network ports, floor space, maintenance contracts, and electricity required to run a data center operation.

Virtual infrastructures provide a fantastic solution to the challenge of a distributed server architecture; however, the native storage virtualization within VMware ESX server does not provide for the same benefits and hardware reductions as it does in the server space. It is common to see storage usage in an environment have the opposite result and actually increase after the implementation of a virtual infrastructure such as the VMware ESX server.

Enterprise-class ESX deployments require a shared form of storage, or a storage area network (SAN). Any consolidation effort inevitably requires greater uptime and fault tolerance from the consolidation platform. A failure in a disk subsystem can have a catastrophic impact to business within a virtual infrastructure since multiple systems would be impacted by such a failure. As we will discuss, the cost for implementing RAID levels that provide data protection beyond that of a single disk failure can come either with a price premium or with a performance penalty that renders the technology unusable for data center operations.

A second challenge commonly found in virtual infrastructures is the difficulty and/or inability to complete backup operations within a traditional backup window. This challenge is a natural byproduct of the disproportion of data to server bandwidth that is created when the physical servers are condensed by the virtual infrastructure. While not all VMware deployments will initially face backup challenges, it has quickly become a hot-button issue in the virtualization community. The following link lists the conference agendas from VMworld 2005. Note the number of times the word "backup" is on the agenda: www.vmware.com/vmtn/vmworld.

The contents of this white paper will demonstrate how integrating Network Appliance technologies into your virtual infrastructure can provide solutions to the challenges inherent with ESX deployments in the areas of storage utilization, cost-effective fault tolerance, and backups. With Network Appliance virtualized storage and data management solutions, administrators are able to make dramatic gains in these areas. Additionally, this paper will review backup solutions for VMware deployments that are not on NetApp storage.

Table of Contents

1. The Cost Of Data Protection.....	4
2. Storage Virtualization.....	5
2.1 Realizing Thin Provisioning.....	5
2.2 Extending Beyond Physical Storage Limits.....	7
3. VMware VIRTUAL MACHINE Backups.....	8
3.1 Backing Up Virtual Machines as Physical Servers.....	9
3.2 Backing Up Virtual Machines as Files.....	9
3.3 Enhanced Backups with Raw Device Mappings.....	10
3.4 Disaster Recovery of a Virtual Infrastructure Using NetApp Replication Technology.....	10
3.5 Making Sense of the Backup Options.....	11
4. Summary.....	12
Appendix A: Sample “Brute Force” Snapshot Backup Script (using ESX snapshots).....	12
Appendix B: Migrating from Virtual Disks to RDMs.....	14
Appendix C: Recovering a Single File from a NetApp Snapshot (RDM method).....	14
Appendix D: Recovering a VM from a NetApp Snapshot Copy (RDM only).....	15
References.....	16

1. The Cost Of Data Protection

The cost of data protection can be considered in two ways. First, there is the cost of the RAID level being implemented, specifically, how many additional hard drives are required in order to provide fault tolerance. This must be measured against the second way to consider the cost of data protection, which is to determine the financial cost to the business should data be lost. In the following paragraphs we will consider both forms of this question.

With any consolidation effort the consolidation platform must provide for a higher level of availability since the business impact due to a failure is magnified in relationship to the consolidation factor. Let us consider a common scenario in which a department has 10 servers and each has implemented RAID5. One of the servers has a disk drive failure, and during the RAID rebuild process a media error is found on one of the remaining drives. This server's RAID rebuild process fails and data is lost, requiring this data to be restored from a backup set.

Consider that you have deployed the same servers from the previous paragraph as virtual machines (VMs) in a virtual infrastructure. The 10 VMs store their data on SAN attached storage and the data is protected with RAID5. The impact of the same failure from the previous paragraph would have an impact 10 times greater in magnitude than before as now all 10 virtual machines have lost data and require a restore operation.

The previous paragraphs illustrate an example of the potentially negative side of any consolidation effort if the consolidating platform is not more reliable than the original distributed platform. It is with this understanding that many administrators look to provide a form of data protection that is more resilient than what their physical servers were deployed with (which typically is RAID5). Cost, performance, and storage utilization are also considerations when searching for the appropriate level of data protection for a virtual infrastructure.

Many administrators consider RAID10, which provides data protection against a double disk failure and (more likely) protects against encountering a media error during the RAID reconstruction process. RAID10 is also considered one of the highest performing forms of RAID technologies since it does not compute parity information when committing data writes. Even with the value of its data protection and high performance there is a significant cost to RAID10 because this technology requires an additional 100% overhead of physical disk storage (Nx2). This high cost is counter to a consolidation effort as RAID10 immediately decreases overall storage virtualization by 50%. For more information on RAID levels please see http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#RAID_10.

When looking at emerging technologies to provide fault tolerance that is on par with that of RAID10, administrators may consider implementing RAID6. RAID6 is an extension to RAID5 and, in order to fully understand it, we should review its origin. RAID5 stripes data and parity information across a set of disks, thus providing fault tolerance in the event of a failed disk drive. RAID6 extends the data protection provided with RAID5 by writing a second set of parity data. RAID6 can provide high storage utilization as it requires only a single drive beyond RAID5 (N+2). With the value of enhanced data protection and the cost savings of requiring only one additional drive with RAID6 there is a tradeoff in performance. RAID 6 builds off the parity calculations and additional writes of RAID5. Due to the negative performance impact of these additional calculations and write operations the storage industry is not seeing many data center deployments of RAID6. For more information on RAID levels please see http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#RAID_6.

Network Appliance addresses the question of the cost of data protection by offering RAID-DP™. RAID-DP exceeds the fault tolerance of RAID10 while maintaining the cost savings of RAID6 and also delivers exceptional performance. RAID-DP provides fault tolerance for the failure of any disk in the RAID set. This is not the case with RAID10 where if the same disk drive fails in both mirror sets all data is lost. RAID-DP only requires 12.5% of storage for RAID overhead (the default setting is 14+2 or N+2). In terms of performance

RAID-DP incurs an almost inappreciable performance penalty. For more information on RAID-DP performance and data protection capabilities please see www.netapp.com/library/tr/3298.pdf.

In summary, it is critical for a consolidation platform to provide higher availability as the impact of a failure is multiplied when compared to a distributed deployment. When choosing a storage platform for a virtual infrastructure one needs to consider the cost associated with the RAID technology and whether the RAID technology is in line with the virtualization goal of increasing asset utilization. As you can see from this section NetApp is unique in the storage industry as our RAID-DP technology provides all of the requirements of the ideal storage platform for virtual infrastructures.

2. Storage Virtualization

2.1 Realizing Thin Provisioning

Storage utilization goes beyond the cost and overhead required to provide fault tolerance. With every host that is connected to a SAN there are multiple layers of storage virtualization and in turn each layer has its own level of utilization. Typical storage environments include the RAID layer, a volume management layer, and a file system layer. In this section we will shift our focus from the RAID layer and review storage provisioning specific to a VMware ESX server environment. Before we begin, a bit of VMware storage background review is required.

VMware ESX server provides two means of provisioning storage to a virtual machine: virtual disks and raw device mappings of SAN LUNs (RDMs). Historically VMware ESX administrators have deployed virtual disks. This practice is the most common as it was the only storage option available before the release of ESX server version 2.5. With version 2.5 administrators were able to leverage the storage features inherent within SAN, including the ability to share storage with multiple hosts, thus enabling VMware VirtualCenter and VMotion technology, which allows virtual machines to be dynamically moved from one ESX server to another without service interruption.

In this section we will cover both disk options available to virtual machines beginning with virtual disks. With an ESX server the SAN administrator has to provision storage to the ESX server and this storage is formatted with the VMware file system (VMFS). The VMFS area represents the volume manager layer inside ESX. It is at this layer that an ESX administrator will create and assign virtual disk(s) to a VM. A virtual disk is a file that is presented to a VM as a local SCSI disk drive. When the virtual disk is created it allocates 100% of the required storage from the VMFS volume on which it resides (e.g., a 40GB virtual disk actually consumes 40GB of disk space).

In order to visualize how storage is consumed in this design, let us consider the following example. For this example we will follow common best practices such as limiting volume usage to 80% of capacity for optimal system performance. Consider that you have a dual CPU ESX server hosting 20 VMs. The VMFS volume is one terabyte (TB) in size and contains 20 virtual disks 40 gigabytes (GBs) in size for a total of 800GBs of utilized storage. Figure 1 reviews the design of our example.

VMWARE ENVIRONMENT	
Number of Virtual Machines	20
Size of Virtual Disk (GBs)	40
Stored Data per VM (GBs)	32

Figure 1)

With this example, if each VM proceeds to fill the virtual disk to 80% of capacity (32GBs), the storage utilization ranges from 32% to 57%, depending on the level of RAID protection. See Figure 2 for a comparison of the details.

	RAID10	RAID6	RAID-DP
Raw Storage	2,000	1,125	1,125
Useable Storage	1,000	1,000	1,000
VMFS Storage Allocated	1,000	1,000	1,000
Space Allocated for V-Disks	800	800	800
Data Stored Inside V-Disks	640	640	640
Raw Storage Utilization	32%	57%	57%

Figure 2)

NetApp allows administrators to leverage the power of storage virtualization by allowing thin provisioning of the VMFS volume that stores the virtual disks. With this solution NetApp provisions 1TB of storage to the ESX server but only consumes the amount of data that is actually written. If we implement thin provisioning with our example of 20 40GB virtual disks, the total storage utilization is 800GBs and achieves 100% storage utilization at the VMFS layer. As the VMFS layer needs additional storage it will be able to write data up to its original size of 1TB as if the storage was pre-allocated. In the example NetApp thin provisioning realized a storage utilization increase from 57% in the original configuration to 71%. See Figure 3 for a comparison of the details.

	RAID10	RAID6	RAID-DP
Raw Storage	2,000	1,125	900
Useable Storage	1,000	1,000	800
VMFS Storage Allocated	1,000	1,000	1,000
Space Allocated for V-Disks	800	800	800
Data Stored Inside V-Disks	640	640	640
Raw Storage Utilization	32%	57%	71%

Figure 3)

Let's consider this same environment but for this example we will move away from virtual disks and replace the VM storage with RDMS. With RDMS an administrator would create a 40GB LUN for each VM. Each LUN will be masked inside of ESX and then assigned to a virtual machine in the same manner as a virtual disk is configured. By utilizing RDMS with NetApp thin provisioning technology an administrator can obtain a much

greater level of storage virtualization. With NetApp thin provisioning with RDMs, each VM is allowed to write up to its maximum LUN capacity while consuming only the amount of data actually written to the disk. In our example we stated that each VM would consume only a maximum 80% of its local disk. With this design NetApp thin provisioning realized a storage utilization increase from 57% in the original configuration to 89% (640GB of useable storage is 89% of the raw 720GB). See Figure 4 for a comparison of the details.

	RAID10	RAID6	RAID-DP
Raw Storage	2,000	1,125	720
Useable Storage	1,000	1,000	640
VMFS Storage Allocated	n/a	n/a	n/a
Space Allocated for RDMS	800	800	800
Data Stored Inside RDMS	640	640	640
Raw Storage Utilization	32%	57%	89%

Figure 4)

Please note that most file systems do not immediately reclaim space from deleted files, leaving those blocks on disk and consuming empty block with new writes (this is why utilities such as Undelete are able to recover deleted files). Due to this behavior, thin provisioned LUNs will consume more space on the storage system than that reported used by the VM. NetApp offers tools, technologies, and engineering support to insure that your thin provisioning strategy will meet your expectations. Please be sure to review with your NetApp technical representative before implementing thin provisioning.

With NetApp thin provisioning, storage policies can be put into place, enabling the storage to automatically manage its size as the thin provisioned storage utilization grows over time. In summary, NetApp thin provisioning technology enables virtual infrastructures to leverage a storage platform that is unique in the storage industry with its ability to maximize the hardware utilization level of the VMware ESX server.

2.2 Extending Beyond Physical Storage Limits

One of the most exciting capabilities of a virtualized server infrastructure is the ability to quickly deploy virtual machines for a project that may only temporarily require server resources. Administrators frequently find themselves in need of physical server resources for such diverse tasks as development and QA environments, upgrade and patch testing, and disaster recovery exercises. Without a virtual infrastructure, administrators typically find it difficult to find appropriate resources for these nonproduction environments.

Within a virtual infrastructure, it becomes extremely easy to quickly deploy temporary virtual machines for any number of tasks. Unfortunately, in a traditional shared environment the cost of deploying the necessary storage resources for these temporary resources remain as high as the permanent ones.

When using NetApp storage technology within a virtual server infrastructure, it becomes possible to take advantage of NetApp LUN clone and volume FlexClone™ technologies to provide temporary storage resources in conjunction with temporarily provisioned virtual machines. When using these technologies common storage blocks between the temporary copy of data and the permanent copy consume no additional physical space on the storage system. Only the actual difference between the two will require its own storage resources.

To describe these capabilities we will build on the demonstrations used in the previous section. For this example the previously mentioned 20 virtual machines are Windows® 2003 servers performing a variety of tasks within an organization. A new service pack is released for the server platform and the system administrators and application owners would like to evaluate the impact of applying this service pack within their environment before moving forward with it in production.

In a virtual infrastructure with traditional shared storage, it is a relatively easy task to replicate the existing production environment and deploy 20 new virtual machines in a nonproduction environment in order to test applying the new service pack. While this will be inexpensive in terms of server resources it will be rather expensive in terms of storage resources, requiring a further 100% of deployed storage resources in order to make a second copy of the production virtual machines' virtual disks.

When the virtual infrastructure is used with NetApp storage, it is possible to use the LUN clone or volume FlexClone technologies to deploy temporary storage resources to go along with the temporary virtual machines. When initially deployed, these temporary copies require no additional physical storage resources in order to exist. Only as changes are made to the temporary copies will physical storage be required in order to store those changes.

Returning to the RDM demonstration, let's assume that after deploying the temporary environment, configuration changes and the application of the service pack to the virtual machines has generated 2GB worth of changed data on each of the temporary virtual machines. When using NetApp storage only 40GB of new physical storage in total will be required to store those changes as opposed to the 100% storage overhead required by other storage technologies. Please see Figure 5 for a comparison. Note: The numbers indicated below continue the use of RDMs, but the clone technologies work equally well with virtual disks.

	RAID10	RAID6	RAID-DP
Raw Storage	4,000	2,250	765
Useable Storage	2,000	2,000	680
VMFS Storage Allocated	n/a	n/a	n/a
Space Allocated for RDMs	1,600	1,600	1,600
Data Stored Inside RDMs	1,280	1,280	1,280
Raw Storage Utilization	32%	57%	167%

Figure 5)

As can be seen, using NetApp LUN clone or volume FlexClone technologies in a virtual infrastructure allows an administrator to temporarily experience storage utilization that can greatly exceed even the total physical storage allocated to the environment. For more information on these technologies please reference the following links: www.netapp.com/library/tr/3347.pdf and www.netapp.com/library/tr/3348.pdf.

3. VMware VIRTUAL MACHINE Backups

A VMware ESX server provides several options in terms of methods to back up the data served within each virtual machine. For the purposes of this paper we will focus specifically on the choices available in the area of completing a "hot" or "operational" backup. A hot backup is defined as a backup process that is completed while the VM is up and servicing requests. We will not be addressing "cold" or "offline" backups as the demand for these types of backups is rare by comparison to the previous method.

The greatest challenge facing VMware deployments is the disproportionate ratio of data to physical bandwidth, which is created as tens of physical servers are virtualized and consolidated onto a single server. We see that most administrators continue to back up each VM as if it was a physical server. The choice to backup each VM in this manner is logical, as it does not require any changes to backup operations; however, this option quickly uncovers the problems incurred by increasing the data to bandwidth ratio by failing to complete backups within the backup window. This challenge forces administrators to make a choice: they can continue backing up their virtual infrastructure as if it were physical and reduce the amount of consolidation or they can implement a storage vendor-based backup solution.

As stated earlier in this paper, backups are a major challenge and SAN-based backups are going to be enhanced with future versions of ESX server. This solution is referred to as a consolidated backup. We will review ESX backup options in depth and we will include the concept of a consolidated backup server that is available today with NetApp solutions.

3.1 Backing Up Virtual Machines as Physical Servers

With this method, backup operations remain the same as they were before the virtual infrastructure was implemented. The challenge with this method is that should an administrator want to drive server utilization up in the range of 8 to 12 servers per ESX CPU, the data to hardware bandwidth ratio becomes disproportional and consequently backups are unable to be completed within a traditional backup window. A common method of resolving this issue is to artificially capacitate the ESX server by reducing the total number of servers virtualized by that server. While this is an effective method of ensuring that backups are completed within their required window, it can dramatically reduce the savings realized by a virtual infrastructure by forcing the purchase of additional ESX servers.

Network Appliance provides a solution for backing up both physical and virtual machines with its Open Systems SnapVault® (OSSV) products. OSSV is a unique solution in the backup industry as it provides individual system backups that are stored on NetApp storage devices. With OSSV each virtual machine completes a full backup once, and for every subsequent backup only the block level changes are sent over the network. Implementing OSSV results in a significant reduction of data transferred in and out of the ESX server, allowing backups to complete quickly and eliminating the need for any artificial capacity limits. Restoration of complete VMs or individual files can be completed easily and in a traditional manner. In addition, OSSV is able to back up VMs that are not stored on NetApp storage systems, allowing easy integration into any existing ESX environment. For more information on backing up servers with OSSV, please refer to the following link: www.netapp.com/library/tr/3466.pdf.

3.2 Backing Up Virtual Machines as Files

Many administrators who have experienced the challenges of backing up VMs as physical servers have elected to back up their VMware environment by backing up the files that make up the VM (the virtual disks and configuration files). With this method traditional backup operations are modified for the ESX environment and storage solutions that provide for disk-based backups are required. Backing up the VM files produces quick, full VM backups; however, recovery of either a VM or an individual file can be cumbersome. In order to complete the restore of a VM the VMFS LUN must be restored to an alternate ESX server. This limit is imposed by the ESX server and applies to all storage vendors. For a full VM restore, the virtual disk files must be sent via ethernet from the recovery server to the production server. In order to perform an individual file restore the VMFS LUN is connected to an alternate ESX server, a recovery VM is created and connected to the disk, and the file to be restored is copied to the production VM via ethernet.

In addition to enabling disk-based backups, an administrator is required to select a disk backup technology. These technologies traditionally come in the following two flavors: mirror copies and copy out snapshots. With mirror copies the production data is mirrored to a second set of disks and the mirror is broken off at the time of backup. This method provides optimal performance but requires the cost of 100% additional storage for every backup copy. With copy out snapshot technology, the storage requirement is greatly reduced as

the storage system only copies out the block level changes to the data. However, copy out snapshot technology incurs a large negative impact on performance, and thus it is hard to find this technology deployed for applications such as VMware.

The inherent negative features of traditional disk-based backups do not apply to NetApp patented Snapshot™ technology. With Snapshot there is no performance penalty for taking Snapshot copies because the data is never moved as it is with copy out technologies, and the cost for Snapshot copies is only at the rate of block level changes and not 100% for each backup as it is with mirror copies. When combining NetApp Snapshot technologies with VMware ESX server, administrators are able to back up their entire virtual infrastructure within seconds and open up a number of other data management possibilities. Once taken, NetApp Snapshot copies can be backed up to tape and/or replicated to another facility with SnapMirror®, VMs can be restored almost instantly, individual files can be easily recovered, and clones can be instantly provisioned for test and development environments. For more information on NetApp Snapshot technology please refer to the following links: www.netapp.com/library/tr/3001.pdf and www.netapp.com/library/ar/ar1038.pdf.

3.3 Enhanced Backups with Raw Device Mappings

VMware recommends that administrators interested in advanced storage functionality under ESX server utilize raw device mappings. RDMs in virtual compatibility mode provide all of the functions of a virtual disk, including the ability to be used with VMotion, but RDMs are much more flexible and beneficial when they are integrated with storage management functions such as NetApp Snapshot, SnapRestore®, LUN Clone, and FlexClone. When an administrator integrates NetApp Snapshot technology with VMware RDMs, the challenges associated with individual file recovery are virtually eliminated. As each RDM is formatted with the native file system of the VM, the Snapshot of the RDM can easily be connected to either another VM or a physical system and accessed via Fibre Channel or iSCSI with NetApp SnapDrive®. Once the Snapshot RDM is connected, a user can simply copy the file to be recovered. See the Appendices of this document for the procedure for connecting to an RDM Snapshot with SnapDrive. For more information on VMware recommended usages for RDMs, please refer to the following link: www.vmware.com/pdf/esx25_rawdevicemapping.pdf.

Another powerful backup solution is available when using RDMs. Because Snapshot copies of RDMs can be fully accessed, it is possible for administrators to utilize SnapDrive to connect all backup RDMs to a consolidated backup server that has the sole function of copying all the data to tape. (As this server is not a production server it can run all day long in order to complete the backups.) An administrator considering this solution is required to have a consolidated backup server compatible with each version of file system stored in the RDMs. For example, a Windows 2003 server would be required for all NT 4.0, Windows 2000, and Windows 2003 RDMs. To back up EXT3 RDMs, a Linux® server would be required.

3.4 Disaster Recovery of a Virtual Infrastructure Using NetApp Replication Technology

As virtual infrastructure implementations mature and more and more mission-critical applications are run on virtual machines, site disaster recovery starts to become a larger issue within the backup and recovery space. The limitations of the tape medium can cause difficulty in a disaster recovery scenario, as the limitations of tape device data transfer speeds and the physical distance between a primary data center and its DR equivalent can potentially mean extended service outages in the event of a site disaster.

For an administrator storing VMware virtual machines on a NetApp storage system, SnapMirror replication technology can be used to dramatically reduce the impact of a site disaster to business processes. When using SnapMirror technology, a virtual infrastructure can be easily replicated over the wire to a remote data center. With this technology, recovering a virtual machine affected by a site disaster may be completed in minutes instead of the hours or days required by other storage solutions. For more information on SnapMirror technology, please see the following links: www.netapp.com/products/software/snapmirror.html and www.netapp.com/library/tr/3446.pdf.

3.5 Making Sense of the Backup Options

Since there are so many choices of backup options available within VMware environments, we have included the following chart to help summarize and compare each of them. Please see Figure 6.

ESX BACKUP TYPE	...VMS AS PHYSICAL SERVERS	...VMS AS FILES
Non NetApp Solutions	<p>Easy to implement, customer may already own software</p> <p>Backups run into bandwidth issues</p> <p>Artificially limits the amount of consolidation per server</p>	<p>Mirror copies require 100% additional storage per backup</p> <p>Copy out snapshots require a small amount of storage but have negative performance impact on production systems</p> <p>Individual file recovery is a challenge</p> <p>Individual VMFS VM recovery is a network copy operation to an alternate ESX server</p>
Solutions Integrated with NetApp Technologies	<p>OSSV backups only send block level changes after initial full backup</p> <p>Eliminates artificial consolidation limits</p> <p>Works on all storage platforms</p>	<p>Patented Snapshot technology provides high performance and space savings</p> <p>Backups complete in seconds</p> <p>Individual file recovery is a challenge</p> <p>Individual VMFS VM recovery is a network copy operation</p>
Solutions Integrated with NetApp & RDMS	<p>OSSV backups only send block level changes after initial full backup</p> <p>Eliminates artificial consolidation limits</p> <p>Works on all storage platforms</p>	<p>Patented Snapshot technology provides high performance and space savings</p> <p>Backups complete in seconds</p> <p>Individual files and VMs can easily be recovered in seconds</p>

Figure 6)

In summary, administrators have many options when considering how to back up their ESX server environment; each option provides its own unique benefits and limitations. Solutions available within ESX server, such as VMFS volumes and raw device mapping, only add to the complexity of choosing the appropriate backup and recovery solution. NetApp provides solutions for both backing up a virtual infrastructure as physical servers with OSSV and for file-based backups via our patented Snapshot technology. OSSV can be deployed with storage other than that from NetApp and can be layered with NetApp Snapshot technology for the most functional VMware backup and recovery solutions.

4. Summary

Many data centers are in the early stages of converting their physical infrastructure into a virtual infrastructure in which hardware utilization is greatly increased and applications are not tied to physical hardware. Network Appliance provides advanced storage virtualization and exciting technologies in the areas of advanced fault tolerance, thin provisioning, instantaneous storage cloning, and advanced backup and recovery solutions. NetApp systems are the ideal storage platform for a virtual infrastructure as they provide solutions unparalleled in the storage market to VMware challenges.

As products that support virtual infrastructures mature, they will inevitably begin to support additional storage technologies such as NAS and iSCSI. As a market leader in these spaces, NetApp will be able to continue to offer unique and innovative solutions within the virtual infrastructure market.

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

Appendix A: Sample “Brute Force” Snapshot Backup Script (using ESX snapshots)

The following script shows an example of one way to automate the snapshot backup of all your virtual machines running on a single ESX server. This script can be extended to put all of the VMs on multiple servers into “hot backup mode.” The total time to complete this script should be seconds.

```
#!/bin/sh
#
# take-hot.sh
#
# Example code to take a snapshot of all RDM VMs using REDO log capability
# provided with the vmware-cmd facility. It will maintain and cycle the
# last 4 snapshots.
#
# NOTE: VMware ESX does not ship with the RSH utility. The RSH protocol is
# An inherently insecure mechanism for remote computing, especially outside
# of a secure trusted environment. Alternate mechanisms utilizing ssh and
# the Perl based ONTAPI library are available for adaptation inside of a
# script like this one which may be developed for production-quality
# implementation. For the purposes of this demonstration, a viable copy of
# the RSH binary was obtained from a running instance of Linux and is used
# herein to facilitate a quick and easily readable code example.
#
# This sample code is provided AS IS, with no support or warranties of any
# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# 05.12.2006 Vaughn Stewart, Network Appliance
```

```

#
# -----

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

# location where Vmware LUNs reside
FILER=filer50
VOL=vmware_luns
# location of VM config files
VM_CFG_PATH=/vmfs/vmx/

echo ' Starting ---> `date`

# rename and delete old snapshots (maintain 4 backups)
RSH -l root $FILER snap delete $VOL backup4
RSH -l root $FILER snap rename $VOL backup3 backup4
RSH -l root $FILER snap rename $VOL backup2 backup3
RSH -l root $FILER snap rename $VOL backup1 backup2

# set up the REDO log file on the boot drive
for i in `ls $VM_CFG_PATH`
do
vmware-cmd $i addredo scsi0:0
done

# take a new snapshot of the volume
RSH -l root $FILER snap create $VOL backup

# apply the REDO and get rid of the log file
# <freeze>=0 <level>=0 <wait>=1
for i in `ls $VM_CFG_PATH`
do
vmware-cmd $i commit scsi0:0 0 0 1
done

echo ' Finished ---> `date`

```

Appendix B: Migrating from Virtual Disks to RDMs

For ESX server administrators who are interested in the benefits of raw device mappings and would like to know how to import a current virtual disk into a RDM, please refer to the document at www.vmware.com/pdf/esx25_rawdevicemapping.pdf.

Appendix C: Recovering a Single File from a NetApp Snapshot (RDM method)

This process involves restoring the NetApp Snapshot copy on either a physical or virtual host via either the Fibre Channel protocol or via an iSCSI initiator (software or hardware). The iSCSI software initiator is the simplest method. As VMware ESX server deployments traditionally host Windows servers and their LUNs are formatted with NTFS, this document describes this procedure from a Windows host.

1. Right click on My Computer and select Manage.
 - a. Expand Storage
 - i. Expand SnapDrive
2. Create a new iSCSI connection (if required).
3. In the right pane, right click and select Connect to Disk.
 - a. The Connect Disk Wizard will launch.
 - b. Enter the UNC path to the share on the VMware LUNs volume.
 - i. See Figure 7.

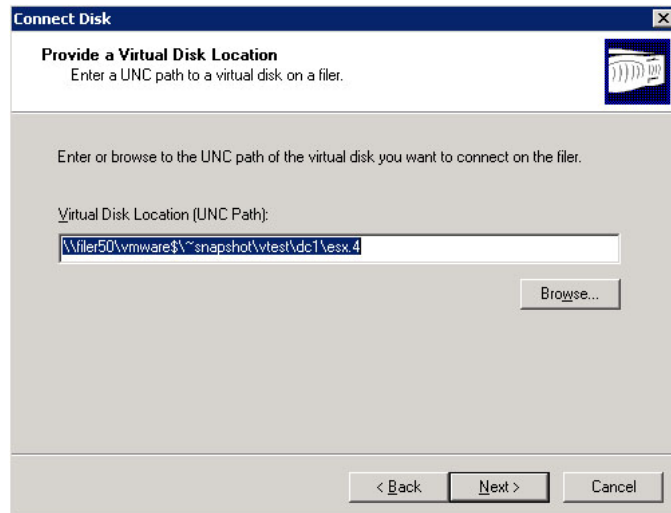


Figure 7)

- c. Select Dedicated Disk.
- d. Assign a Drive Letter.
- e. Select an Initiator (FCP or iSCSI).
- f. Finish.

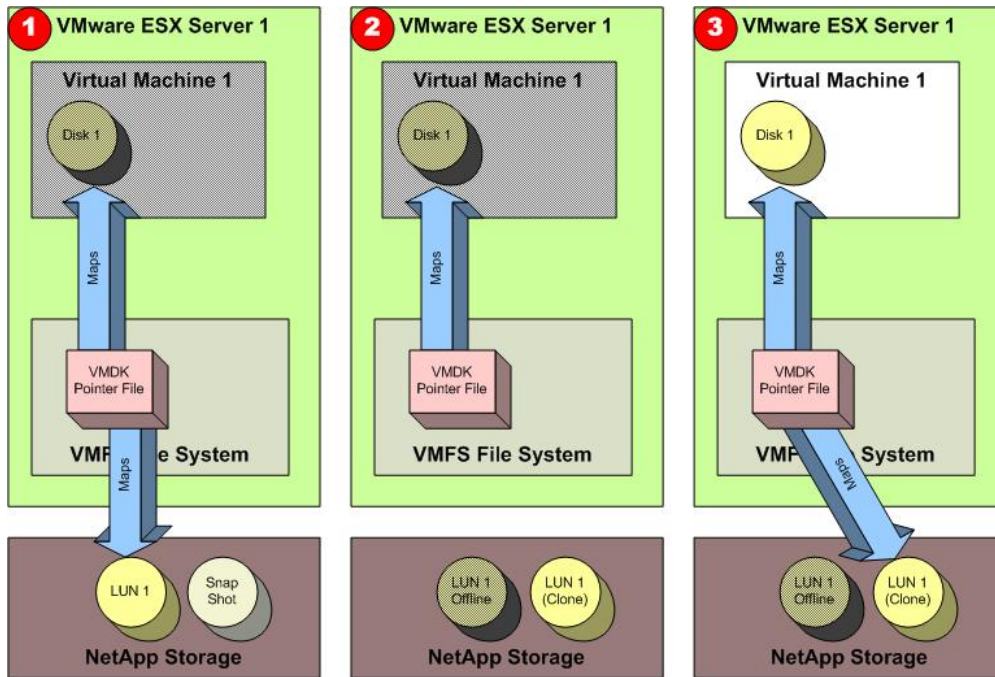
4. Open Windows Explorer and browse to the drive you connected.
 - a. Browse to the directory or file to recover.
 - i. Select Copy.
 - b. Browse to a share on the VM.
 - i. Paste the copied directory or files.
5. Open SnapDrive and delete the LUN.

Appendix D: Recovering a VM from a NetApp Snapshot Copy (RDM only)

This section details the process required to restore a virtual machine instantaneously via NetApp Snapshot technology. Figure 6 below represents a high level view of the process listed below.

RDM Recovery Overview

1. Power off the VM.
2. Rename and offline the LUN to be restored, and clone the LUN from a Snapshot.
3. Power on the VM.



RDM Recovery Process

1. Stop the virtual machine.
 - a. From the ESX server console select the stop button (red square) for the VM.
2. Offline and rename the current version of the LUN in use.
 - a. From the system console verify the igroup and ID of the LUN to be restored.

- i. lun show -m (LUN path)
 - b. From the system console rename and offline the LUN.
 - i. lun move (original LUN path) (new or renamed LUN path)
 - ii. lun offline (offlines or disables the LUN)
3. Clone the original LUN from a recent Snapshot, bring it online, and map it.
 - a. From the system console
 - i. lun clone create (original LUN path) -b (original LUN path) (Snapshot name)
 - ii. lun online (LUN path)
 - iii. lun map (LUN path) (igroup) (ID)
4. Start the virtual machine.
 - a. From the ESX server console select the start button (green triangle) for the VM.
 - b. Note: If you restore the LUN to the original LUN ID you do not need to rescan the FC bus.
5. Validate that the restore is to the correct version.
 - a. Log into the server and verify that the system was restored to the proper timeframe.
 - b. If the system was not restored to the correct timeframe then repeat the process and select a different Snapshot.
 - c. If the restore was correct, proceed to Step 6.-
6. Delete the original LUN and split the clone into a whole LUN.
 - a. From the system console delete the original LUN.
 - i. lun destroy -f (renamed LUN path)
 - b. From the system console split the LUN clone into an independent LUN.
 - i. lun clone split start (LUN path)

References

TR-3428 Network Appliance & VMware ESX Server

Instantaneous Backup & Recovery Including Single File Restoration with NetApp Snapshot Technology

www.netapp.com/library/tr/3428.pdf

TR-3393 Using Network Appliance Snapshot Technology with VMware ESX Server

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

TR-3401 Using a Network Appliance SAN with VMware to Facilitate Storage and Server Consolidation

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

Wikipedia RAID definitions and explanations

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

VMware VMworld Conference Sessions Overview

www.vmware.com/vmtn/vmworld

TR-3001 A Storage Network Appliance

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

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

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

VMware Using Raw Device Mapping

www.vmware.com/pdf/esx25_rawdevicemapping.pdf

TR-3466 Open Systems SnapVault (OSSV) Best Practices Guide

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

TR-3347 FlexClone Volumes: A Thorough Introduction

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

TR-3348 Block Management with Data ONTAP® 7G: FlexVol™, FlexClone, and Space Guarantees

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

SnapMirror software overview

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

TR-3446 SnapMirror Best Practices Guide

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



© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, DataFabric, FAServer, FilerView, NetCache, NearStore, SecureShare, SnapManager, SnapMirror, SnapRestore, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, ContentDirector, ContentFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, MultiStore, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW NetApp on the Web, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SnapVault, SohoCache, SohoFiler, SpinMirror, SpinShot, SpinStor, The evolution of storage, Vfiler, VFM, Virtual File Manager, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.