NETAPP TECHNICAL REPORT

# NetApp SAN Solutions for VMware

NetApp
August 2008

TR-3704

TABLE OF CONTENTS

# 1 EXECUTIVE SUMMARY

Many companies are adopting virtualization technology as a means to further server and storage consolidation goals. The intent of this report is to illustrate the strength of NetApp Fibre Channel and iSCSI storage protocols as the supporting infrastructure for VMware deployments. Of particular focus will be the integration of block storage and VMFS, NetApp Host Utilities Kits, SnapDrive, Storage Management & Application Integration (SMAI), and replication of virtual machines to provide DR and HA solutions.

The NetApp SAN solution offers users of VMware an excellent platform for hosting VM datastores. Using NetApp technology, VMware datastores can be consolidated onto more efficiently scaled storage architecture accessed across a Fibre Channel or iSCSI SAN. This dramatically simplifies storage provisioning, increases asset utilization, and allows a company to take advantage of NetApp data protection and disaster recovery functionality.

The options presented to the user for iSCSI and/or FC attached shared storage and for VMFS and RDM storage management will be explored in the pages that follow. This report will outline the relative advantages of choosing either RDM or VMFS datastores and the various circumstances which would lead one to select one versus the other.

NetApp Host Utilities for managing iSCSI and FC LUNs within the guest OS present significant value added features when used in a VMware environment.

The paper will also cover NetApp's SnapManager for Virtual Infrastructures. This section will cover the integration of VMFS and NetApp Snapshot technology to improve storage usage efficiency and recoverability of VMDK files.

NetApp data protection and retention technologies compliment the DRS and HA features available from VMware VI3. This section will explore SnapMirror, SnapVault, MetroCluster and VMware DRS and HA integration.

# 2 VMWARE STORAGE ARCHITECTURE

## 2.1 SAN ARCHITECTURE OVERVIEW

VMware ESX Server can access FC and iSCSI shared storage and present the underlying disk capacity to virtual machines through virtual disk adapters within each VM. The storage is accessed as raw disk mapping or VMFS type volumes. This presentation is transparent to the guest OS but may have significant impact on the configuration options to be chosen when the disks are made available to the guest OS. It is important to understand when each of these options are selected what the implications are on the use of volume types, host utilities and NetApp SnapManager.

### 2.1.1 VMware Storage Components

VMware Infrastructure requires storage to be managed at two levels. Globally at the VMware ESX Server layer and locally at the Guest VM layer. NetApp provides for the management of storage at both layers and presents a considerable value in its approach to presenting and managing storage at each layer of the virtual infrastructure.

At the ESX layer, LUNs are presented from the NetApp FAS array by either FC or iSCSI and are added to the VMware datastore as RDM or VMFS-3 volumes. The storage in the datastore can span multiple LUNs within the storage subsystem or even multiple arrays. The datastore is a logical pool of storage made up of volumes with the same formatting which is then allocated in portions to virtual machines as virtual disks. Within the containing datastore the virtual disk is treated as a file (.vmdk) for easier and more flexible management of the virtual machine storage than with traditional server and storage architectures.

For further reference, a detailed examination of VMware storage virtualization concepts and options can be found in the ***VMware SAN System Design and Deployment Guide,*** available from VMware.

### 2.1.2 Virtual Machine Storage Access

The presentation of the partitioned storage is accomplished through internal virtual SCSI adapters within ESX. Each virtual machine may contain up to four (4) virtual adapters. The virtual disk partition may be used for either the guest OS boot volume or as additional storage for a particular virtual machine.

It is important to determine the type of storage access desired by the virtual machines prior to creating the supporting virtual disks. A comparison of the features available with VMFS and RDM are explained in the following sections.

## 2.2 VMWARE FILE SYSTEM (VMFS) OVERVIEW

Virtual Machine File System (VMFS) Datastores are the most common method of deploying storage in VMware environments. VMFS is a clustered file system that allows LUNs to be accessed simultaneously by multiple ESX Servers running multiple VMs. The strengths of this solution are that it provides high performance and the technology is mature and well understood. In addition, VMFS provides the VMware administrator with a fair amount of independence from the storage administrator, because once storage has been provisioned to the ESX Servers, the VMware administrator is free to use the storage as needed. Most data management operations are performed exclusively through VMware VirtualCenter.

When a LUN is presented from the NetApp FAS array and formatted as a VMFS volume the user is able to take advantage of a number of features available only from NetApp for VMware environments. VMFS also provides some management advantages versus using RDM as well.

The VMFS-3 file system contains a number of enhancements over previous versions which improved upon the scaling potential and SAN interoperability. In relation to SAN deployments,

the most important enhancement to VMFS-3 is the inclusion of a Logical Volume Manager (LVM). The LVM in VMFS-3 allows for automatic resignaturing of Snapshot volumes for enabling dynamic use of Snapshots within the ESX environment.

One key differentiator when using VMFS is that a datastore may span multiple NetApp LUNs without impacting the consistency of the host-based replication scheme applied to the datastore. As the datastore is managed by VMware it may be managed more atomically from within the ESX environment. This feature enables for multiple datastore creation from array presented LUNs for easier management. The virtual disks can be dynamically assigned without the process of manually creating connections to each individual VM as well. This works to simplify the management of the storage infrastructure and provides for a greater variety of operational use and recovery options for the VMware Infrastructure.

When using VMFS, SnapManager for Virtual Infrastructure may perform individual VM Snapshot restores. Using SMVI a Snapshot of the entire datastore may be created and, when restored, either the entire datastore restored or individual *vmdk* files can be restored.

One caution when using VMFS is that the datastore is serving the I/O demands of many VMs, and this design doesn't allow a storage array to identify the I/O load generated by an individual VM. The VMware administrator must perform I/O load monitoring and management, which has traditionally been handled by storage administrators. VMware VirtualCenter allows the administrator to collect and analyze this data. NetApp extends the I/O data from VirtualCenter by providing a mapping of physical storage to VMs, including I/O usage and physical path management with VMInsight in NetApp SANscreen®. For more information about SANscreen, see http://www.netapp.com/us/products/managementsoftware/sanscreen-vm-insight.html.

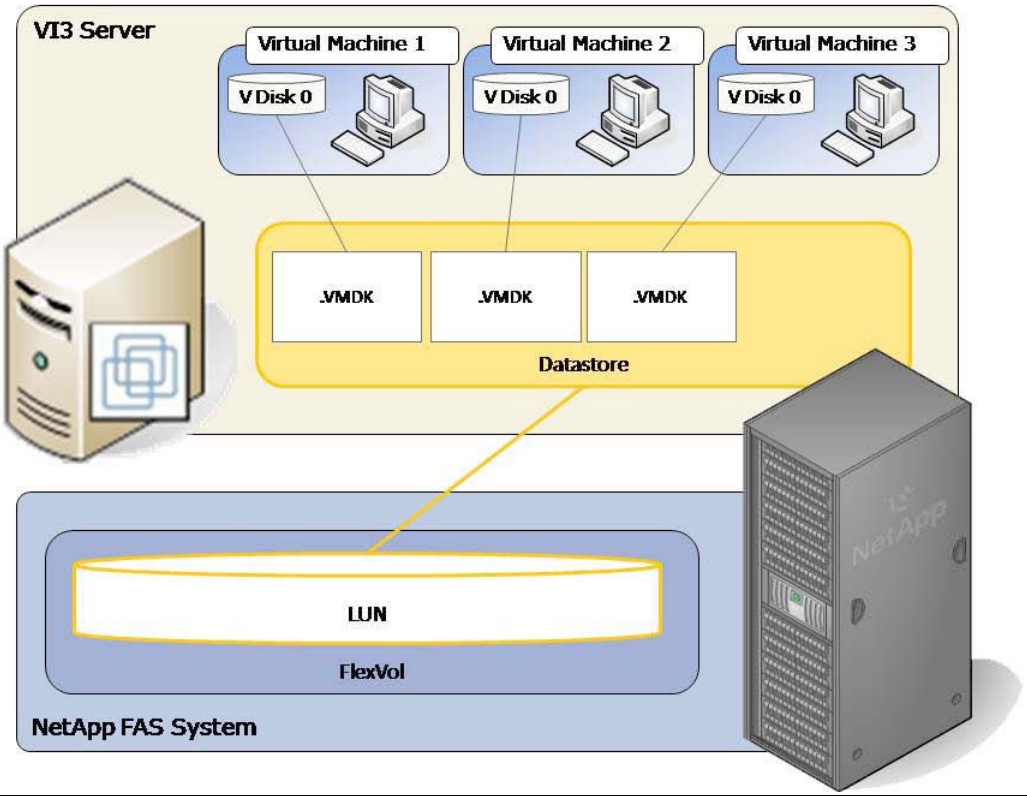

## 2.3 RAW DEVICE MAPPING

Support for raw device mapping (RDM) was introduced in VMware ESX Server 2.5. Unlike VMFS datastores, which provide storage as a shared, global pool, RDMs provide LUN access directly to individual virtual machines. ESX acts as a connection proxy between the VM and the storage array. The core strength of this solution is support for virtual machine and physical-to-virtual-machine host-based clustering, such as Microsoft® Cluster Server (MSCS). In addition, RDMs provide high individual disk I/O performance; easy disk performance measurement from a storage array; and easy integration with features of advanced storage systems such as SnapDrive®, VM granular Snapshots, SnapRestore, and FlexClone®. The core weakness is that it does add some administrative overhead into the configuration.

The challenges of this solution are that VMware clusters may have to be limited in size, and this design requires ongoing interaction between storage and VMware administration teams. Each *vmdk* file has a direct I/O channel to a dedicated LUN. This storage model is analogous to providing SAN storage to a physical server, except for the storage controller bandwidth, which is shared.

RDMs are available in two modes; physical and virtual. Both modes support key VMware features such as VMotion, and can be used in both HA and DRS clusters. The key difference between the two technologies is the amount of SCSI virtualization that occurs at the VM level. This difference results in some limitations around MSCS and VMsnap use case scenarios. For more information about raw device mappings over Fibre Channel and iSCSI, see the *VMware ESX Server 3i Configuration Guide.*

Using an RDM type datastore can be useful when configuring LUNs for systems which will be dynamically changing often. Test and development systems where the VMs will be frequently remapped or shifted between distributed virtual environments may find it advantageous to use RDMs. One NetApp advantage in an RDM environment is the ability to use FlexClone to create multiple read-writable Snapshot copies of the VM datastore and present those clones to multiple systems independently.

One caution when configuring the system to use RDM is in setting the OS LUN-type at the time of LUN creation. When then the volume is created in the NetApp array, it should not contain an OS specific format type (i.e. Linux, Windows, etc.) but should be left to configuration per VM at the guest OS layer by SnapDrive. It is important to address this potential issue at the time the VMDK is created. SnapDrive will dynamically create a LUN from the NetApp array and assign the manual connection to the VM. However, it should be noted that the LUN must be created with the appropriate “LUN-type” properties as the intended guest OS. Failure to do so may result in LUN misalignment with negative operational implications for the system. Taking into account these predeployment considerations when using an RDM type storage configuration is important to a successful implantation.

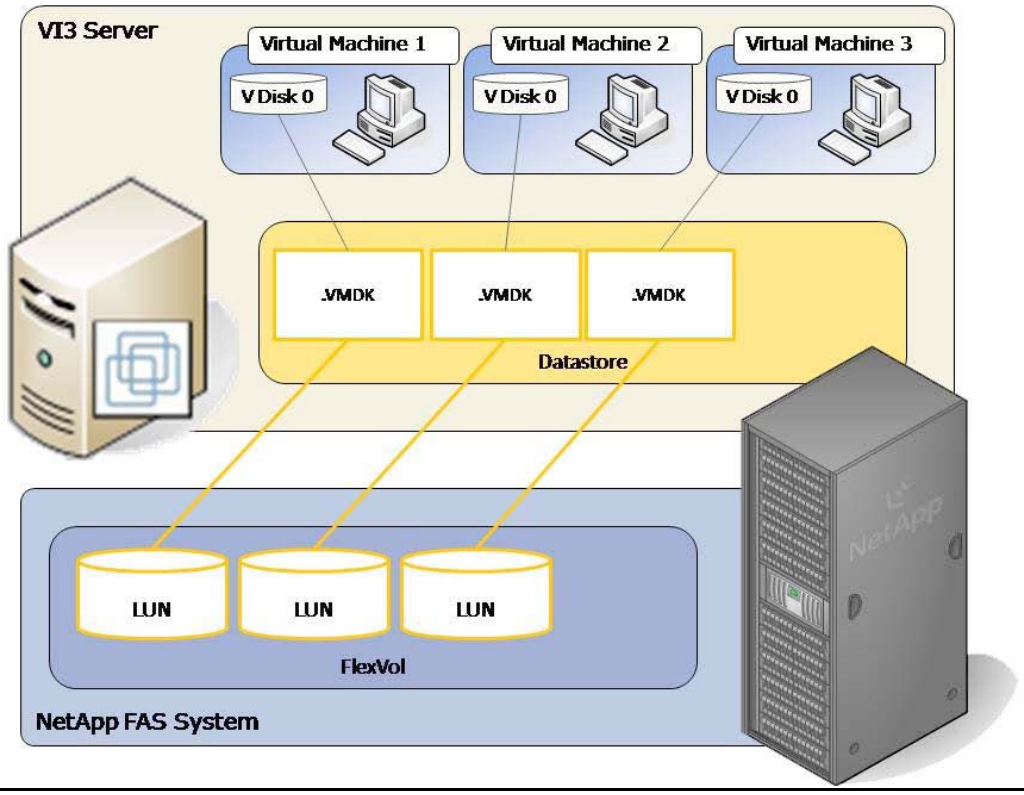


## 2.4 RDM OR VMFS: WHICH IS RIGHT FOR YOU?

This guide seeks to present the options available to VMware administrators when using NetApp SAN storage and to further clarify points illustrated in TR3428. Choosing the correct infrastructure is an individual exercise dependent upon the specific environment and under what workload and service level requirement the system is to be deployed. This guide does not seek to recommend or endorse any one deployment method as superior to another architecture because each environment is unique. It is also important to understand that it is possible to deploy multiple storage formats for unique datastores in environments with mixed workloads. This adds flexibility to the architecture when accommodating multiple types of server workload.

## 2.5 STORAGE MANAGEMENT AND BOOT FROM SAN

For further reference, a detailed examination of VMware storage virtualization concepts and options can be found in the ***VMware SAN System Design and Deployment Guide,*** available from VMware.

VMware SAN Implementation considerations

## 2.6 IDENTIFYING APPLICATION DEPENDENCIES

The selection of the storage architecture for the VMware Virtual Infrastructure is multifaceted. As has been explored in the previous chapter the type of LUN to be used for the datastore is one facet of designing the supporting infrastructure. Another is the I/O channel configuration of the SAN. From the transport protocol to the multipath options there are a number of considerations to be weighed and used to determine the best architecture for a given environment.

## 2.7 I/O CHANNEL CONFIGURATION OPTIONS

A feature comparison of FC and iSCSI protocol:

| Capability/Feature | FCP | iSCSI |
|---|---|---|
| Format | VMFS or RDM | VMFS or RDM |
| Max Datastores or LUNs | 256 | 256 |
| Max Datastore size | 64TB | 64TB |
| Max running VMs per Datastore | 32 | 32 |
| Available link speeds | 1, 2, 4Gb | 1, 10Gb |
| Protocol overhead | Low | Moderate under high load |
| **Backup Options** | | |
| VMDK image access | VCB | VCB |
| VMDK file level access | VCB, Windows® only | VCB, Windows only |
| NDMP granularity | Full LUN | Full LUN |
| **VMware Feature Support** | | |
| VMotion | Yes | Yes |
| Storage VMotion | Yes | Yes |
| VMware HA | Yes | Yes |
| DRS | Yes | Yes |
| VCB | Yes | Yes |
| MSCS support | Yes, via RDM | Yes, via RDM |
| Resize Datastore | Yes, but not in production | Yes, but not in production |
| **NetApp Integration Support** | | |
| Snapshot copies | Yes | Yes |
| SnapMirror® | Datastore or RDM | Datastore or RDM |
| SnapVault® | Datastore or RDM | Datastore or RDM |
| Data Deduplication | Yes | Yes |
| Thin provisioning | Datastore or RDM | Datastore or RDM |
| Open Systems SnapMirror | VM | VM |
| FlexClone | Datastore or RDM | Datastore or RDM |
| MultiStore® | No | Yes |

When provisioning LUNs for access via FC or iSCSI, the LUNs must be masked so that only the appropriate hosts can connect to the LUNs. With a NetApp FAS system, LUN masking is handled by the creation of initiator groups. NetApp recommends creating an igroup for each VMware cluster. NetApp also recommends including in the name of the igroup the name of the cluster and the protocol type (for example, DC1_FCP and DC1_iSCSI). This naming convention and method simplify the management of igroups by reducing the total number created. It also means that all ESX Servers in the cluster see each LUN at the same ID. Each initiator group includes all of the FCP worldwide port names (WWPNs) or iSCSI qualified names (IQNs) of the ESX Servers in the VMware cluster.

NetApp recommends using dedicated physical resources for storage traffic whenever possible. With IP storage networks, this can be achieved with separate physical switches or a dedicated storage VLAN on an existing switch infrastructure. It is also recommended that iSCSI traffic be separated from other IP network traffic by implementing a separate network or VLAN for the virtual machine traffic used for such features as VMotion.

One of the challenges of configuring VMware ESX networking for IP storage is that the network configuration should meet these three goals simultaneously:

- Be redundant across switches in a multiswitch environment
- Use as many available physical paths as possible
- Be scalable across multiple physical interfaces

To enable ISCSI connectivity, the ESX Server requires a special connection type, referred to as a VMkernel port, along with a service console port. NetApp recommends that each ESX Server should have two service console ports, and the second port should be configured on the same vSwitch as the VMkernel port. The VMkernel network requires an IP address that is not currently in use on the ESX Server.

For further examination of these and other topics and configuration best practices refer to the reference section at the end of this document for additional documentation.

# 3 MANAGING THE STORAGE INFRASTRUCTURE

## 3.1 MANAGING CHANGE: PROVISIONING STORAGE FOR VMWARE

### 3.1.1 NetApp FlexVols: Thin provisioned storage

With traditional storage provisioning and the storage blocks are preallocated and assigned to a server—or, in the case of VMware, a virtual machine. It is also a common practice for server administrators to overprovision storage in order to avoid running out of storage and incurring the associated application downtime when expanding the provisioned storage. Although no system can be run at 100% storage utilization, there are methods of storage virtualization that allow administrators to address and oversubscribe storage in the same manner as with server resources (such as CPU, memory, networking, etc.). This form of storage virtualization is referred to as ***thin provisioning***.

Traditional provisioning preallocates storage; thin provisioning provides storage on demand. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. It is important when thin provisioning and oversubscribing storage that monitoring policies be implemented to ensure free space for each virtual machine is continuously available in order to guaranty service level requirements are met. Without the addition of physical storage, if every VM requires its maximum possible storage at the same time, there will not be enough storage to satisfy the requests. Establishing and monitoring storage utilization thresholds for opportunity to add more storage "*just-in-time"* is an important function in a thin provisioned environment.

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

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

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

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

### 3.1.2 VMware: Thin provisioned Virtual Machines

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

By default in ESX 3.5, virtual disks preallocate the storage they require and in the background zero out all of the storage blocks. This type of VMDK format is called a "***zeroed thick VMDK***."

VMware provides a means to consume less storage by provisioning VMs with thin-provisioned virtual disks. With this feature, storage is consumed on demand by the VM. When using thin provisioned data stores and NetApp storage there is no performance or operational penalty which may be experienced when implementing this VMware feature with other vendors' storage arrays.

Thin-provisioned VMDKs are not available to be created in the Virtual Infrastructure client with VMFS Datastores. To implement thin VMDKs with VMFS, you must create a thin-provisioned VMDK file by using the vmkfstools command with the –d options switch. By using VMware thin-provisioning technology, you can reduce the amount of storage consumed on a VMFS datastore.

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

NetApp offers storage virtualization technologies that can enhance the storage savings provided by VMware thin provisioning. These technologies offer considerable storage savings by increasing storage utilization with deduplication redundant data and thin provisioning VMFS and RDM LUNs. Both of these technologies are native to FAS arrays and don't require any configuration considerations or changes to be implemented with VMware.

### 3.1.3 Data Deduplication

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

Illustration of typical storage consumption in a VI3 deployment:

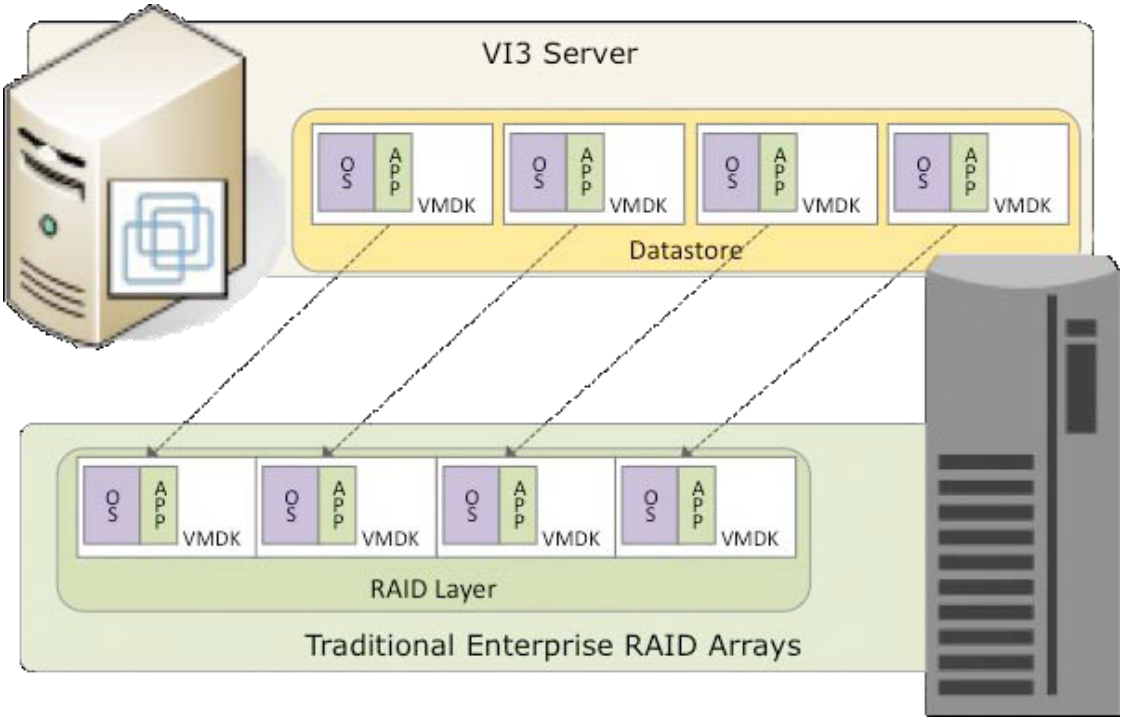


NetApp offers a data deduplication technology called *NetApp Deduplication.* With NetApp Deduplication, VMware deployments can eliminate the duplicate data in their environment, enabling greater storage utilization. Deduplication virtualization technology enables multiple virtual machines to share the same physical blocks in a NetApp FAS system in the same manner

that VMs share system memory. It can be seamlessly introduced into a virtual infrastructure without having to make any changes to VMware administration, practices, or tasks. Deduplication runs on the NetApp FAS system at scheduled intervals and does not consume any CPU cycles on the ESX Server.

An example of the impact of deduplication on VMDK storage utilization:

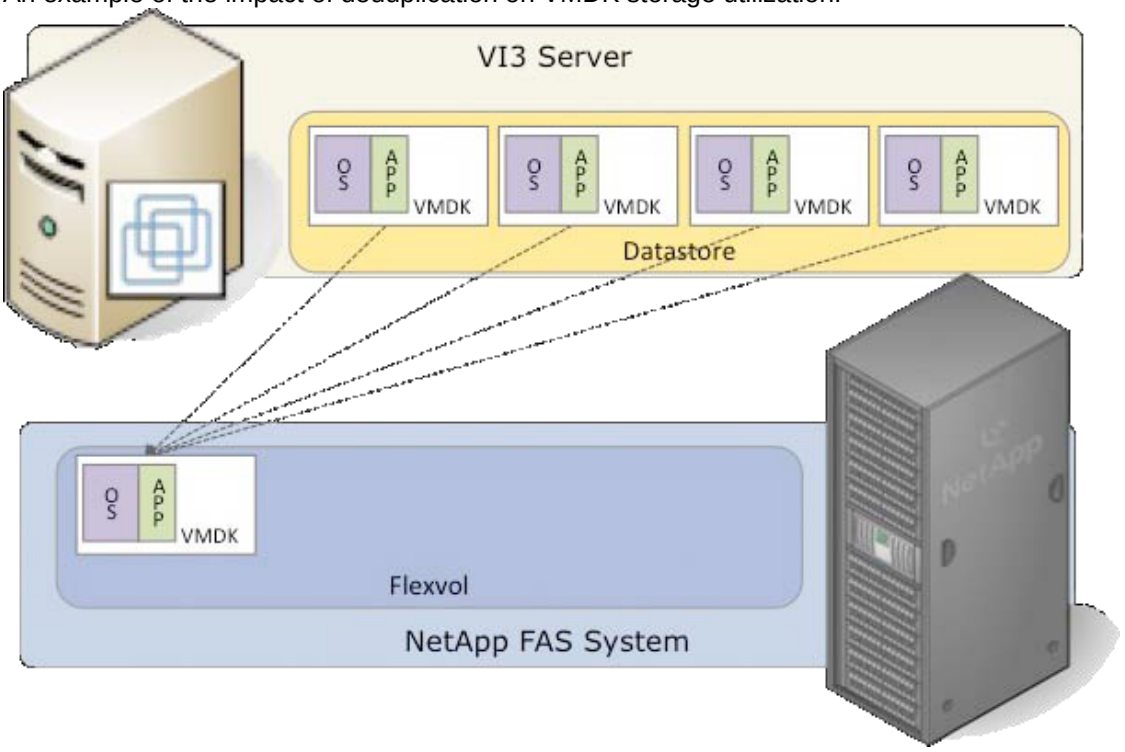


### 3.1.4 Deduplication with VMFS and RDM LUNs

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

To recognize the storage savings of deduplication with LUNs, you must enable NetApp LUN thin provisioning. In addition, although deduplication reduces the amount of consumed storage, the VMware administrative team does not see this benefit directly, because their view of the storage is at a LUN layer, and LUNs always represent their provisioned capacity, whether they are traditional or thin provisioned.

For more deduplication best practices, including scheduling and performance considerations, see TR 3505.

## 3.2 NETAPP HOST UTILITIES KITS (HUK/HAK)

NetApp provides this utility for simplifying the management of ESX nodes on FC SAN. This utility is a collection of scripts and executables that is referred to as the FCP ESX Host Utilities for Native OS. Using the NetApp host utilities will ensure that the implementation is configured to optimally and reliably function. NetApp highly recommends using this kit to configure settings for FCP HBAs.

One of the components of the Host Utilities is a script called config_mpath. This script reduces the administrative overhead of managing SAN LUN paths by using the procedures previously described. The config_mpath script determines the desired primary paths to each of the SAN LUNs on the ESX Server and then sets the preferred path for each LUN to use one of the primary paths. Simply running the config_mpath script once on each ESX Server in the cluster can complete multipathing configuration for large numbers of LUNs quickly and easily. If changes are made to the storage configuration, the script is simply run an additional time to update the multipathing configuration based on the changes to the environment.

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

For more information about the FCP ESX Host Utilities for Native OS, see
http://now.netapp.com/NOW/knowledge/docs/hba/fcp_esx/fcpesxhu31/pdfs/install.pdf

# 4 MAINTAINING RESILIENCY IN A VIRTUAL INFRASTRUCTURE

## 4.1 REPLICATION AND RECOVERY WITH NETAPP

One of the main benefits of virtualization for disaster recovery is independence of the recovery process from the recovery hardware. Because virtual machines encapsulate the complete environment, including data, application, operating system, BIOS, and virtualized hardware, applications can be restored to any hardware with a virtualization platform without concern for the differences in underlying hardware. The physical world limitation of having to restore to an identical platform does not apply. Not only does hardware independence allow IT managers to eliminate manual processes associated with adjusting drivers and BIOS versions to reflect the change in platform, it also eliminates Windows® registry issues and plug-and-play issues. By leveraging the hardware independence of VMware virtual machines, customers no longer need to worry about the need for identical hardware at their DR sites, which can significantly reduces the cost and complexity of regional DR. VMware enterprise customers actively take advantage of VMware consolidation benefits for their production and staging servers. These consolidation benefits are even greater for the failover hardware, because customers can consolidate servers at the primary data center to fewer physical servers at their disaster recovery centers.

Another benefit of VMs that helps to ease the complexity of DR is the VMware flexible networking features. Because VMware handles VLANs on its virtual switches, entire complex network environments can be isolated, contained, or migrated very easily with little setup at the DR site.

The advantages of VMware in the realm of disaster recovery and high availability are enhanced greatly when NetApp FAS systems and solutions are used in conjunction with VMware. With Snapshots, SnapMirror, and MetroCluster backups and disaster recovery are further simplified in iSCSI or Fibre Channel environment. With the ability to backup, recover, and create offsite duplicates of entire VMware ESX VMFS datastores, the rapid recovery of entire VMware Infrastructures or individual VM's becomes both a manageable and simple process. With this in mind, both recovery time objectives (RTOs) and recovery point objectives (RPOs) are reduced for organizations running in VMware Fibre Channel or iSCSI environments. By reducing both of these objectives both the cost and risk associated with a potential disaster are significantly mitigated. NetApp FAS systems also work in conjunction with VMware High Availability (HA) to minimize the downtime incurred by hardware failure. Furthermore, when Snap Manager for Virtual Infrastrutures (SMVI) is brought into play in a VMware Fibre Channel or iSCSI environment the ability to perform rapid backup and recovery of individual VM's becomes a reality.

### 4.1.1 Snapshot Technology

VMware Virtual Infrastructure 3 introduced the ability to create Snapshot copies of virtual machines. Snapshot technologies allow the creation of point-in-time copies that provide the fastest means to recover a VM to a previous point in time. NetApp has been providing customers with the ability to create Snapshot copies of their data since 1992, and although the basic concept of a Snapshot is similar between NetApp and VMware, you should be aware of the major differences between the two, and when you should use one rather than the other.

VMware Snapshots provide simple point-in-time versions of VMs, allowing quick recovery. The benefits of VMware Snapshots are that they are easy to create and use, because they can be executed and scheduled from within VirtualCenter. VMware suggests that the Snapshot technology in ESX should not be leveraged as a means to back up Virtual Infrastructure. For more information about native VMware Snapshots, including usage guidelines, see the ***VMware Basic System Administration Guide*** and the ***VMware Storage/SAN Compatibility Guide for ESX Server 3.5 and ESX Server 3i.***

NetApp Snapshot technology can easily be integrated into VMware environments, where it provides crash-consistent versions of virtual machines for the purpose of full VM recovery, full VM cloning, or site replication and disaster recovery. This is the only Snapshot technology that does

not have a negative impact on system performance. VMware states that for optimum performance and scalability, hardware-based Snapshot technology is preferred over software-based solutions. The disadvantage of this solution is that it is not managed within VirtualCenter, requiring external scripting and/or scheduling to manage the process. For details, see the **VMware *Basic System Administration Guide*** and the ***VMware ESX Server 3i Configuration Guide.***

### 4.1.2 Data Layout for Snapshot Copies

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

Virtual machines should have their swap files, pagefile, and user and system temp directories moved to separate virtual disks residing on separate datastores residing on NetApp volumes dedicated to this data type. In addition, the ESX Servers create a VMware swap file for every running VM. These files should also be moved to a separate datastore residing on a separate NetApp volume, and the virtual disks that store these files should be set as independent disks, which are not affected by VMware Snapshots.

For example, if you have a group of VMs that creates a Snapshot copy three times a day and a second group that creates a Snapshot copy once a day, then you need a minimum of four NetApp volumes. For traditional virtual disks residing on VMFS, each volume contains a single LUN; and for RDMs, each volume contains several RDM formatted LUNs.

### 4.1.3 VMware HA with NetApp Data protection

VMware Infrastructure changes the way that information systems are designed. Featuring such advanced capabilities as migration of virtual machines between any virtualization platforms, Snapshot™ copies, automated restart on alternate hosts in a resource pool, and VMotion, VMware Infrastructure creates environments where outages are limited to brief restarts at most. For a continuous availability solution to guard against application or hardware failure, VMware HA provides easy-to-use, cost-effective protection for applications running on virtual machines. In the event of server failure, affected virtual machines are automatically restarted on other physical servers in a VMware Infrastructure resource pool that have spare capacity.

As a result of the advantages listed above, VMware HA minimizes downtime and IT service disruption while eliminating the need for dedicated standby hardware and installation of additional software. VMware HA provides uniform high availability across the entire virtualized IT environment without the cost and complexity of failover solutions tied to either operating systems or specific applications.

In iSCSI and Fibre Channel environments, VMware HA is further enhanced when used in conjunction with NetApp storage. NetApp RAID-DP provides increased reliability over other RAID technologies by improving the resiliency of disk arrays, allowing double disk failures to occur within the same RAID group and still providing data availability to the VMware infrastructure. Also, NetApp has an active-active controller design to ensure data availability. Active-active controllers provide simple automatic transparent failover to deliver enterprise class availability to NetApp storage. Due to RAID-DP and the active-active design, storage availability is virtually assured during normal business operations. With storage availability assured due to RAID-DP and active-active design, VMware HA is free to easily migrate and power on virtual machines down due to server hardware failures. By assuring storage availability in iSCSI and Fibre Channel environments, NetApp and VMware HA work in conjunction to assure that critical virtual machines achieve maximum run time in a VMware ESX environment.

### 4.1.4 SnapMirror

SnapMirror software is the value leader in the industry when it comes to disaster recovery (DR). Its simplicity and flexibility make it affordable for customers to deploy a DR solution for more of their application infrastructures than would be possible with competitive alternatives. SnapMirror supports synchronous replication limited to metro distances, ensuring zero data loss; semi-synchronous replication that supports recovery point objectives (RPOs) in seconds with minimal impact on the host application; and asynchronous replication, which is the most cost-effective solution that can meet RPOs ranging from 1 minute to 1 day. Its functionality and configuration flexibility enable SnapMirror to support multiple uses, including disaster recovery, data distribution, remote access, data migration, data replication, and load balancing.

If an incident occurs that makes the entire campus unavailable, NetApp SnapMirror provides long-distance replication to protect against such incidents. Operating either asynchronously or synchronously, SnapMirror utilizes NetApp Snapshot copies to make replication both easy and efficient.

In iSCSI and Fibre Channel VMware ESX environments, SnapMirror allows for off-site replication of VMFS data stores hosted on NetApp storage. Based on the RPO of a customer, VMFS data stores can be synchronously or asynchronously enabled. In the event of a disaster SnapMirror would allow an entire VMware infrastructure to be up and operational in a short period of time at a location hundreds or thousands of miles away.

### 4.1.5 MetroCluster

MetroCluster is a unique, cost-effective, synchronous replication solution for combining high availability and disaster recovery in a campus or metropolitan area, to protect against both site disasters and hardware outages. MetroCluster provides automatic recovery for any single storage component failure, and single-command recovery in case of major site disasters, ensuring zero data loss and making recovery possible within minutes rather than hours. When combined with VMware iSCSI and FC protocols and VMFS an incredibly high level of data protection and disaster recovery can be achieved decreasing recovery time objectives (RTOs) in the event of a disaster which eliminate a VMware infrastructure at the main site.

Virtual Infrastructure deployed in conjunction with a NetApp FAS systems using MetroCluster has an additional built-in level of robustness. Any virtual machine that resides on a NetApp FAS system can survive a crash of the server hardware that runs the VM, and can be restarted on another ESX Server at the original or an alternate campus location. Utilizing a NetApp FAS system's replication technology, a VM can be replicated and restored anywhere in the world, whether it's cross-campus or cross-country, with little IT staff intervention.

### 4.1.6 SnapManager for Virtual Infrastructures

SMVI is a backup solution for VMware ESX 3.02 and above that addresses the growing problem associated with backing up individual VM's. SMVI uses NetApp Snapshot™ technology to instantaneously backup VMware VMFS data stores. It then provides granular restoration of individual virtual machines or entire datastores as needed within minutes. By providing this functionality SMVI reduces recovery time objectives (RTOs) and recovery point objectives (RPOs) for organizations inside of their VMware infrastructures while minimizing the load on ESX infrastrucutes.

SMVI allows for automated data protection by granting the ability to assign backup policies to individual VM's or their associated datastores. Also, SMVI is integrated with SnapMirror™ replication to replicate snapped VM's to an off-site location. With these features SMVI provides hot backups, rapid restores, and disaster recovery solutions for VMware virtual infrastructures that allow for availability, scalability, performance, and reliability in a FC or iSCSI VMFS environment.

# 5 CONCLUSION

VMware Virtual Infrastructure offers customers several methods of providing storage to virtual machines. NetApp FAS storage arrays provide customers the option of connecting their storage by the most effective protocol for a specific environment. Together, these storage options give customers flexibility in their infrastructure design, which in turn provides cost savings, increased storage utilization, and enhanced data recovery.

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

# 6 REFERENCES


***VMware SAN System Design and Deployment Guide***
VMware, March 2007 and July 2008
http://www.VMware.com/pdf/vi3_san_design_deploy.pdf

***VMware iSCSI SAN Configuration Guide***
VMware, February 2008
http://www.VMware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_iscsi_san_cfg.pdf

***TR3428 NetApp and VMware Virtual Infrastructure 3 Storage Best Practices***
NetApp – M. Vaughn Stewart, Michael Slisinger & Larry Touchette
http://www.netapp.com/library/tr/3428.pdf

***TR3446 SnapMirror Best Practices Guide***
NetApp – Darrin Chapman & Srinath Alapati
http://www.netapp.com/library/tr/3446.pdf

***TR3348 Block Management with Data ONTAP 7G:***
***FlexVol, FlexClone, and Space Guarantees***
NetApp – Jim Lanson
http://www.netapp.com/library/tr/3348.pdf

***TR3606 High Availability and Disaster Recovery for VMware Using NetApp SnapMirror and MetroCluster***
NetApp – Jim Lanson & Srinath Alapati
VMware – Eric Hardcastle
http://www.netapp.com/library/tr/3606.pdf

***TR3393 Using NetApp Snapshot Technology with VMware ESX Server***
NetApp – Brian Casper
http://www.netapp.com/library/tr/3393.pdf

***SnapManager for Virtual Infrastructures***
NetApp – John Lockyer, et al.


***TR3704 was compiled, edited and authored in part by Jeremy LeBlanc TME-SiSBU and Jack McLeod TME-VGIBU of NetApp.***