



Network Appliance and VMware ESX Server 3.0 Building a Virtual Infrastructure from Server to Storage

Vaughn Stewart and Michael Slisinger, Network Appliance
October 2006 | TR-3515

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 by using Network Appliance™ technology.

TABLE OF CONTENTS

Executive Summary.....	3
Performance, Data Protection, and Storage Utilization.....	4
Storage Virtualization: Realizing Thin Provisioning.....	8
Storage Virtualization: Extending Beyond Physical Storage Limits	10
Summary.....	13
References	14

EXECUTIVE SUMMARY

In recent years, just about every company with an Information Systems department has begun some form of consolidation and virtualization effort with the goal of increasing asset utilization while 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 that is universally acknowledged as a leader in the virtualization space is VMware.

With the release of the VMware Virtual Infrastructure 3.0 Suite, companies can decouple business applications from physical server hardware, which in turn reduces operational costs and provides a much more flexible and dynamic infrastructure. By reducing the amount of physical servers, network ports, floor and rack space, maintenance contracts, and electricity required to run a data center operation, companies can actualize the return on investment from consolidation efforts within months. Common virtual infrastructure deployments provide consolidation ratios of 8 to 12 physical servers to a single ESX Server. Multiple ESX Servers are grouped together to form ESX clusters and data centers hosting upwards of hundreds of virtual machines.

Virtual infrastructures are a fantastic solution to the challenges of a distributed server architecture. However, the native storage virtualization capabilities shipped with VMware ESX Server do not provide the same benefits and hardware reductions as those seen in the server space. Many customers have experienced an increase in storage requirements after implementing their virtual infrastructure. The reasons for this increase are many, including but not limited to a requirement for a shared storage platform, inefficiencies in the multiple layers of storage virtualization, overprovisioning, and challenges with backups that can lead to inefficient disk-to-disk backup solutions.

This technical report demonstrates how integrating Network Appliance technologies in a virtual infrastructure can solve the unique challenges inherent with ESX deployments in the areas of storage utilization, fault tolerance, and backups. With Network Appliance virtualized storage and data management solutions, customers can make dramatic gains in these areas. This report also reviews backup solutions for VMware deployments that are not on NetApp storage.

PERFORMANCE, DATA PROTECTION, AND STORAGE UTILIZATION

With every consolidation effort, the consolidation platforms must meet a new set of business challenges that are unique to virtual infrastructures. When considering the acquisition of a storage system, it's important to understand the impacts on disk I/O performance, data protection, and storage utilization. To begin with, a storage system must at a minimum provide the aggregated disk I/O performance of the combined distributed platforms being consolidated. Virtual infrastructures can apply a significant I/O load on disk subsystems. This load is a result of the VMware default storage design. With VMware VMFS datastore, multiple virtual disks (or VMDK files) are stored, which means that multiple virtual machines are concurrently accessing the file system. VMFS datastores are notorious for being extremely random in their read and write requirements. Failure to provide a robust storage system also has a negative impact in areas outside of serving VM data requests. These negative impacts may be experienced in areas such as backing up VM data to tape. For details, see the *Virtual Infrastructure 3 SAN Configuration Guide* at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf

In addition, a consolidation platform needs to provide a high level of availability, because the business impact of a failure is magnified in direct proportion to the consolidation factor. Consider a common scenario where a department has 20 servers. To protect the data stored on each server, RAID 5 has been implemented. One of these servers incurs a disk drive failure, and during the RAID rebuild a media error is found on one of the surviving drives. This server's RAID rebuild process fails and data is lost, requiring this data to be restored from a tape-based backup set. This process usually takes a significant amount of time and effort.

Suppose that you have deployed the same 20 servers as virtual machines in a virtual infrastructure. The VMs store their data on a shared storage platform, and the data is protected with RAID 5. The impact of the same failure just described would be 20 times greater in magnitude, because now all 20 virtual machines have lost data that must be restored.

The cost of data protection should be considered in two ways. First there is the acquisition cost of the RAID level being implemented; specifically, how many additional hard drives are required to provide fault tolerance. Second, this cost must be measured against the cost of impact to business operations if data is lost. The following paragraphs consider both forms of this question.

The previous paragraphs illustrate 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 seek to deploy a form of data protection that is more resilient than what their physical servers were deployed with (typically, RAID 5). Cost, performance, and storage utilization are also considerations when searching for the appropriate level of data protection for a virtual infrastructure.

Many administrators consider RAID 10 (RAID 1+0), which provides data protection against a double disk failure and (more likely) protects against encountering a media error during the RAID reconstruction process. RAID 10 is a nested RAID technology that stripes data over pairs of RAID 1 mirrors. RAID 10 is considered to be one of the highest performing forms of RAID technologies, because it 10 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 RAID 10, because this technology requires an additional 100% overhead of physical disk storage (Nx2). This high cost is counter to a consolidation effort; the use of RAID 10 immediately decreases overall storage virtualization by 50%. For more information on RAID levels, see Figure 1 and http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#RAID_10.




Figure 1 An Example of a RAID 10 Group

In considering emerging technologies to provide fault tolerance that is on a par with RAID 10, administrators may consider implementing RAID 6 or RAID 50 (RAID 5+0). Both technologies are extensions to RAID 5, which stripes data and parity information across a set of disks, providing fault tolerance in the event of a single failed disk drive.

RAID 6 extends the data protection of RAID 5 by writing a second set of parity data. RAID 6 provides high storage utilization, because it requires only a single drive beyond RAID 5 (N+2). RAID 6 provides enhanced data protection and the cost savings of requiring only one additional drive, but there is a tradeoff in performance. RAID 6 requires double the number of parity calculations and additional writes associated with RAID 5. Because of the negative performance impact of these additional calculations and write operations, the storage industry is not seeing many data center deployments of RAID 6. For more information on RAID levels, see Figure 2 and http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#RAID_6.




Figure 2 An Example of a RAID 6 Group

RAID 50 is a nested RAID technology that appears to extend the data protection provided by RAID 5 by striping data (RAID 0) across RAID 5 groups. On the surface, RAID 50 offers excellent data protection in the event of multiple drive failures. However, RAID 50 suffers from the same issue as its predecessor: As disk drive densities increase, the number of physical imperfections increases exponentially. With RAID 50, if two drives fail in a single RAID 5 group, or more likely a single drive fails and during the RAID reconstruction a media error is encountered, the data is lost. The impact of this failure goes beyond the data loss of the individual RAID group, because with RAID 50 the data loss affects the RAID 0 stripes written across all of the RAID 5 groups. For more information on RAID 50, see Figure 3 and http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#RAID_50.




Figure 3) An Example of a RAID 50 Group

Network Appliance RAID-DP™ uniquely addresses the challenge of providing the highest level of data protection while requiring a minimal amount of storage. Introduced in 2004, RAID-DP exceeds the fault tolerance of RAID 10 while maintaining the cost savings found with RAID 6 and RAID 50. RAID-DP provides fault tolerance for the failure of any two disks (data or parity) within a RAID group. In addition, RAID-DP incurs an almost inappreciable performance penalty. For more information on RAID-DP performance and data protection capabilities, see Figure 4, http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks#Double_parity, and <http://www.netapp.com/library/tr/3298.pdf>.




Figure 4 An Example of a RAID-DP Group

Table 1 summarizes the characteristics of each RAID technology. For the storage utilization example, assume that 5TB of usable storage is provisioned. From the data in Table 1, the advantages of RAID-DP are easy to see. RAID-DP provides the highest level of storage utilization, data protection, and performance of any of the RAID technologies reviewed in this document.

Table 1) Comparison of RAID technologies.

	RAID 10	RAID 6	RAID 50	RAID-DP
Raw storage	10000	5625	5625	5625
Useable storage	5000	5000	5000	5000
Maximum storage utilization	50%	88%	88%	88%
Number of drive Failures with no data loss	2	2	1	2
RAID performance level	High	Low	Medium	High

In summary, it is critical for a consolidation platform to provide high availability, because the impact of a failure is multiplied when compared to a distributed deployment. When choosing a storage platform for a virtual infrastructure, it's important 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. The cost savings, reliability, and scalability of RAID-DP make it the best form of data protection that can be implemented in a consolidation effort. NetApp is unique in the storage industry in providing all of the requirements of the ideal storage platform for virtual infrastructures.

STORAGE VIRTUALIZATION: REALIZING THIN PROVISIONING

Storage utilization goes beyond the cost and overhead required to provide fault tolerance or just provisioning storage. With every host connected to a fabric-attached storage array, there are multiple layers of storage virtualization and management, which in turn have their own level of utilization. Typical storage environments include the RAID layer, a volume management layer, and a file system layer. This section reviews storage provisioning that is specific to a VMware ESX server environment. When using virtual disks with an ESX Server, the storage administrator has to provision storage to the ESX Server. This provisioned storage is formatted with the VMware ESX File System (VMFS). The VMFS area represents the volume manager layer inside ESX. At this layer, the ESX administrator creates and assigns virtual disks to virtual machines (VMs). Virtual disks, or VMDK files, are flat files that are presented to VMs as SCSI disk drives connected to a local SCSI bus.

To visualize how storage is consumed in this design, consider the next example, which follows common storage best practices, such as limiting volume usage to 80% of capacity for optimal system performance. Suppose that you have a number of ESX Servers hosting a total of 100 VMs. The VMFS datastore is 5TB in size and contains 20 virtual disks of 40GB each, for a total of 3.2TB of written data. Table 2 summarizes this design.

Table 2) Example VMware environment.

Number of virtual machines	100
Size of virtual disk (GB)	40
Data stored per VM (GB)	32
Total data written to disk (GB)	3200

In the area of storage utilization, NetApp thin provisioning technology provides unmatched efficiencies in advanced storage virtualization. By using NetApp thin provisioning, customers can dramatically increase their storage utilization without sacrificing performance. Thin-provisioned storage has been provisioned just like traditional storage, but it is not consumed until data is written. With traditional models, storage is preallocated, or reserved in advance of any data actually being written. Once storage has been provisioned, it becomes inflexible and any excess in the provisioning becomes, in essence, wasted space waiting for data to someday be stored in it.

Consider the example environment, the VMFS datastore that serves 100 virtual disks. All of the RAID solutions previously described provision 5TB of useable storage to the ESX Servers. However, with NetApp thin provisioning the actual amount of storage consumed is 3.2TB, which is exactly the amount of data that has been written by the 100 VMs. Implementing NetApp thin provisioning in this example increased storage utilization to levels unobtainable with other storage technologies. Table 3 compares the details given in the examples, and Figures 5, 6, and 7 present this information graphically.

Table 3) The Impact of Thin Provisioning

	RAID 10	RAID 5, RAID 6, or RAID 50	RAID-DP w/ Thin Provisioning
Raw storage	10000	5625	3600
Useable storage	5000	5000	3200
VMFS storage allocated	5000	5000	5000
Space allocated for virtual disks	4000	4000	4000
Data stored in virtual disks	3200	3200	3200
True storage utilization	32%	57%	89%

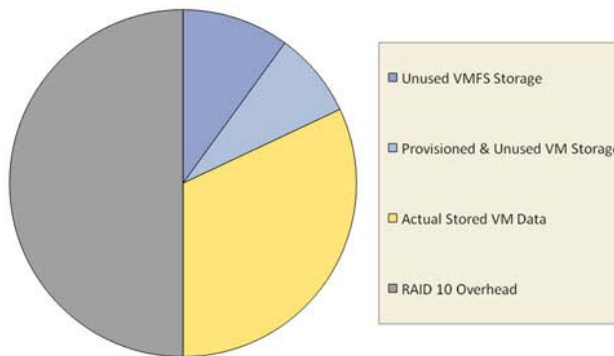


Figure 5) RAID 10 utilization.




Figure6) RAID 5, 6, or 50 utilization.




Figure 7) RAID-DP utilization with thin provisioning.

The VMware ESX Server provides an additional means of provisioning storage to virtual machines, including datastores and raw device mappings of Fibre Channel SAN and iSCSI LUNs (RDMs). NetApp advanced storage virtualization technologies, including thin provisioning, apply just as well to these other storage options. 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.

It's important to 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.) Because of this behavior, thin provisioned LUNs consume more space on the storage system than is being reported as used by the VM. NetApp offers tools, technologies, and engineering support to ensure that your thin provisioning strategy conforms to best practices in order to meet your consolidation goals.

Note: Be sure to consult with your NetApp technical representative before implementing thin provisioning.

Although it is an option and not a requirement for VMware deployments, NetApp thin provisioning technology enables virtual infrastructures to leverage a storage platform that is unique in the storage industry because of its ability to maximize its storage utilization level. It's easy to compare and align the high virtualization rates available with NetApp systems to the goals of any VMware deployment.

STORAGE VIRTUALIZATION: 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 require server resources only temporarily. 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, it is typically difficult for administrators to find resources for these nonproduction environments.

In a virtual infrastructure, it's 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 for these temporary resources remains as high as the cost of the permanent storage.

When using NetApp storage technology in a virtual server infrastructure, it's possible to take advantage of NetApp LUN clone and volume FlexClone™ technologies to provide temporary storage resources in conjunction with temporarily provisioned virtual machines. With 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 requires its own storage resources.

To describe these capabilities, this section builds on the demonstrations in the previous section. For this example, suppose that the previously mentioned 100 virtual machines are Microsoft® Windows® 2003 servers performing a variety of tasks in an organization. A new service pack is released for the server platform, and the system administrators and application owners want to evaluate the impact of applying this service pack in their environment before moving forward with it in production.

In a virtual infrastructure with traditional shared storage, it is relatively easy to replicate the existing production environment and deploy 100 new virtual machines in a nonproduction environment in order to test applying the new service pack. Although this approach is inexpensive in terms of server resources, it is 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 technology 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 is physical storage required in order to store those changes.

Returning to the example cited in the previous sections, assume that, after deploying the temporary environment, configuration changes and applying the service pack to the virtual machines has generated 2GB of changed data on each of the temporary virtual machines. When using NetApp storage, only 200GB of new physical storage in total is required to store those changes, as opposed to the 100% storage overhead required by other storage technologies, including VMware built-in cloning technologies. See Table 4 for a comparison.

Table 4) The Impact of Flexclones

	RAID 10	RAID 6	RAID-DP
Raw storage	20000	11250	3825
Useable storage	10000	10000	3400
VMFS storage allocated	10000	10000	10000
Space allocated for virtual disks	8000	8000	8000
Data stored in virtual disks	6400	6400	6400
Raw storage utilization	32%	57%	167%

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

VMWARE VIRTUAL MACHINE BACKUPS AND DISASTER RECOVERY

Completing backups of virtual infrastructures can be a major challenge for many customers. Backups were the main topic addressed in numerous presentations at VMworld 2005 (<http://www.vmware.com/vmtn/vmworld>). In summary, the challenges are due to the disproportionate ratio of data to physical bandwidth, which is created by consolidating physical servers to a single virtual server. With ESX there are several methods to complete a backup. However, it is important to note that each solution provides its own set of pros and cons. In order to determine the best backup strategy for your environment, it is critical to identify your company's backup and recovery goals to determine which solution best aligns with those goals.

Although VMware provides several methods of backing up the data served in each virtual machine, this paper focuses 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. This paper does not address "cold" or "offline" backups, which are rare in comparison to hot backups.

This section considers the following backup methodologies: traditional file based, storage based, and consolidated backup server. It also describes how NetApp functionality can extend the capabilities of these technologies.

Traditional File-Based Backups

A traditional backup (also known as a file-based backup) is one in which each VM is backed up and restored as if it were a physical server. Backing up each VM in this manner is ideal from an operations standpoint, because procedurally no changes are required. Virtual machines are handled exactly like the physical servers in the environment. The challenge with this method is that, by its nature, the lowest level of granularity of a traditional backup is at the file level. In addition, traditional backups are very redundant in their function. This type of backup process typically attempts to complete a full backup of the entire infrastructure on some schedule, usually once a week.

Because of the disproportionate amount of data being addressed by each individual ESX Server, the ability to back up all of the data stored in an operational backup window is difficult at best. Many customers have found that the only way to meet their backup window is to implement alternative backup solutions such as Storage Based Backups or VMware's Consolidated Backup.

Storage-Based Backups

Many administrators who have experienced the challenges of backing up VMs as physical servers have elected to backup their VMware environment by backing up the files that make up the VM (the virtual disk files and configuration files). Backing up this data directly to tape drives results in the same challenges as those described for traditional file-based backups. There is generally too much data behind each physical server to back up in a traditional backup window.

To maintain high utilization ratios, many customers have asked their storage vendors to implement some form of storage-based backup for their virtual infrastructure. With this method the virtual machines are placed in a hot backup mode; the virtual disks are locked and all new data is written to temporary log files. Once in this state, the virtual disks are backed up. When the virtual disks have been successfully backed up, the locks are released and the contents of the temporary files are flushed back into the virtual disks.

Disk-based backup practices include copying the VMDK file from the production disk to a second set of disks, or—for customers who want a faster operation—some form of split mirror backup technology. Although both of these solutions provide a much faster backup than backing up directly from the production system to tape, both solutions require 100% additional storage for every backup, and that storage needs to be completed and kept online. This requirement for additional storage is so counter to the utilization goals associated with VMware deployments that it should not be considered. Some storage vendors offer copy-out snapshot technologies as alternatives to the 100% additional storage required with split mirror technologies. The I/O overhead required with copy-out snapshot technologies, and the subsequent performance impact, prevent these solutions from being implemented.

The inherent negative features of traditional disk-based backups do not apply to the NetApp patented Snapshot™ technology. With NetApp technology there is no performance penalty for taking Snapshot copies, because the data is never moved, as it is with copy-out technologies. The cost for Snapshot copies is only at the rate of block level changes, not 100% for each backup as with mirror copies. By combining NetApp Snapshot technology with VMware ESX server, administrators can back up their entire virtual infrastructure in seconds and open up a number of other data management possibilities. The NetApp Snapshot copies can be backed up to tape and/or

replicated to another facility with NetApp SnapMirror® or SnapVault, VMs can be restored almost instantly, individual files can be quickly and easily recovered, and clones can be instantly provisioned for test and development environments. OSSV can be incorporated along with Snapshot copies for a very robust solution. For more information on NetApp Snapshot technology, see <http://www.netapp.com/library/tr/3001.pdf> and <http://www.netapp.com/library/ar/ar1038.pdf>.

Consolidated Backups

The VMware Virtual Infrastructure 3.0 suite introduced an additional method of backup, in which the workload of backups is moved from the production ESX Server to a standalone Windows server whose sole purpose is to connect to storage-based backups and back up their contents to tape. This solution is called a VMware consolidated backup (VCB). The problem of not being able to complete the back up of all VM data within a designated backup window is solved with a consolidated backup, because with this solution a nonproduction server can send the backup data to tape, taking as long as needed. Note that although the I/O load is no longer affecting the production ESX Servers, the backups are being drawn from and will still affect the production disk subsystem. Therefore you should take the time to properly size your storage solution.

There are some limitations to this design. First, VM data will have a separate backup and recovery process from the physical infrastructure. Second, your backup software must not utilize the archive bit as a method for computing incremental backups, and at the time this paper was published only Fibre Channel connectivity and Windows VMs were supported. For more information on VMware consolidated backups, see http://www.vmware.com/pdf/vi3_consolidated_backup.pdf.

Administrators who find the features of consolidated backups desirable and who would like to expand on the solution provided by VMware should consider that NetApp SnapDrive® software can provide consolidated backup solutions for RDMs, which can be fully accessed via Fibre Channel or iSCSI.

Disaster Recovery of a Virtual Infrastructure Using NetApp Replication Technology

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

Administrators who are storing their VMware virtual machines on a NetApp storage system can use NetApp SnapMirror replication technology to dramatically reduce the impact of a site disaster on business processes. With 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 can be completed in minutes instead of the hours or days required by other storage solutions. Customers can leverage the replicated copy of their virtual infrastructure for uses such as test and development or tape archiving. For more information on SnapMirror technology, see <http://www.netapp.com/products/software/snapmirror.html> and <http://www.netapp.com/library/tr/3446.pdf>.

SUMMARY

Many enterprises are in some stage of either upgrading their existing storage or migrating to a virtual infrastructure. Network Appliance provides advanced storage virtualization technologies and storage solutions 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, providing solutions to VMware challenges that are unparalleled in the storage market.

As products that support virtual infrastructures mature, customers will inevitably begin to leverage and require support for additional storage technologies such as NAS, iSCSI, and SATA. As a market leader in these spaces, NetApp will continue to offer unique and innovative solutions in 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. Also, expertise is required to solve user-specific deployments. Contact your local Network Appliance representative to speak with one of our VMware solutions experts.

Comments on this technical report are welcome. Please contact the authors [here](#).

REFERENCES

- TR3428 Network Appliance and VMware ESX Server
Instantaneous Backup and Recovery, Including Single File Restoration with NetApp Snapshot Technology
<http://www.netapp.com/library/tr/3428.pdf>
- TR3466 Open Systems SnapVault (OSSV) Best Practices Guide
<http://www.netapp.com/library/tr/3466.pdf>
- TR3347 FlexClone Volumes: A Thorough Introduction
<http://www.netapp.com/library/tr/3347.pdf>
- TR3348 Block Management with Data ONTAP 7G: FlexVol, FlexClone, and Space Guarantees
<http://www.netapp.com/library/tr/3348.pdf>
- SnapMirror Software Overview
<http://www.netapp.com/products/software/snapmirror.html>
- TR3446 SnapMirror Best Practices Guide
<http://www.netapp.com/library/tr/3446.pdf>
- TR3001 A Storage Network Appliance
<http://www.netapp.com/library/tr/3001.pdf>
- Total Cost Comparison: IT Decision-Maker Perspectives on EMC and Network Appliance Storage Solutions in
Enterprise Database Environments
<http://www.netapp.com/library/ar/ar1038.pdf>
- VMware Introduction to Virtual Infrastructure
http://www.vmware.com/pdf/vi3_intro_vi.pdf
- VMware Server Configuration Guide
http://www.vmware.com/pdf/vi3_server_config.pdf
- VMware SAN Configuration Guide
http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf
- VMware Virtual Machine Backup Guide
http://www.vmware.com/pdf/vi3_vm_backup.pdf
- VMware VMworld Conference Sessions Overview
<http://www.vmware.com/vmtn/vmworld>
- ESX Server 3.X Systems
http://www.vmware.com/pdf/vi3_esx_vmdesched.pdf
- Wikipedia RAID Definitions and Explanations
http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks



www.netapp.com

© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, SnapDrive, SnapMirror, and SnapVault are registered trademarks and Network Appliance, FlexClone, RAID-DP, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Microsoft and Windows are registered trademarks of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.