

VMware Data Backup and Recovery

Data Domain Deduplication Storage Best Practices Guide

Abstract

VMware offers extraordinary benefits, but it can come at the cost of extra storage, backup resources and administrative challenges. Data Domain deduplication storage offers a way out by reducing redundant data across VMware data backups, operating at disk speeds, and providing cost-effective replication for fast DR using backup images. This paper reviews the best practices for architecting a backup/recovery/DR approach for VMware using Data Domain storage, regardless of which backup software or scripts are involved.

VMware Data Backup and Recovery

Data Domain Deduplication Storage Best Practices Guide

Table of Contents

1. INTRODUCTION	3
2. VMWARE INFRASTRUCTURE 3 (VI3) OVERVIEW	3
3. VMWARE INFRASTRUCTURE BACKUP ALTERNATIVES: SUCCESS METRICS	4
3.1 GETTING STARTED: BACKUP THE VM GUEST OS FILES AND VMDKS WITH A STANDARD BACKUP CLIENT WITHOUT VCB	5
3.1.1 BACKING UP THE VM GUEST OS FILES.	5
3.1.2 BACKING UP ESX AND VMDKS.	5
3.2 ENTERPRISE STANDARD PRACTICE: VMWARE CONSOLIDATED BACKUP (VCB) WITH COMMERCIAL BACKUP SOFTWARE	6
3.2.1 CONSOLIDATED BACKUP WITH DATA DOMAIN ..	6
3.2.2 PROXY INSTALLATION	6
3.3 ADVANCED BEST PRACTICE: SNAP / COPY VMDK IMAGES TO DATA DOMAIN FROM ESX OR VCB, RESTORE GUEST OS FILE COPIES VIA BROWSING	7
4. CONCLUSIONS	8

1. Introduction

This paper provides general information for using Data Domain deduplication storage systems to backup virtual machines deployed with VMware Infrastructure 3 (VI3). Topics include an overview of VI3 components and various methods of backing up data in VMDK images and Guest OS files for discrete recovery – either locally or remotely in a DR site. But the central focus is on which backup approach is best using specific criteria, and by extension, how can Data Domain deduplication storage systems be leveraged to assist in online restore and replication to a DR site.

VMware is the predominant method in the open systems world for creating multiple virtual machines in a physical computing system. By separating virtual systems from physical constraints, they become easier to manage. Consolidation helps take greater advantage of powerful server assets which can contribute to supporting green initiatives.

But VMware sites also tend to create more storage to manage and protect than their physical counterparts. By making it much simpler to multiply servers, it also becomes more likely that their storage footprints will multiply. This is significantly visible in backing them up. For example:

- ▶ Multiple similar VM environments, times multiple storage image versions for protection and DR, can equal much larger storage than when servers were more expensive to clone.

- ▶ In backing up VMware environments, restore needs may include both full VMDK images as well as individual file restores to the Guest OS. Backing up both VMDKs and Guest OS files can offer optimum protection, but the data is highly redundant. With normal backup target storage, it would mean that many more tapes or that much more disk storage capacity.

Fortunately, with Data Domain systems as the target storage for VMware environment backup, those similar files would be deduplicated at high speed. All of the common data sequences are pooled to the smallest reasonable storage footprint, for on site protection and replication to a DR site. Where most file system backups result in 10x-30x data reduction on a Data Domain system, VMDK-inclusive backups commonly offer 40x to 60x reduction. For more background on Data Domain's award winning deduplication storage system, please see also www.datadomain.com.

When using Data Domain storage with well-understood best practices for backup and taking snapshots in VMware, a deployment can simplify management of consistent images. Once stored, the images are ready to restore, locally or, with optimized deduplicated replication, at a remote DR site. This paper describes best practices for managing these synergistic technologies together.

While vendor references and script examples are included, these are examples only and should not be considered definitive or guaranteed by Data Domain.

2. VMware Infrastructure 3 (VI3) Overview

For a proper introduction to VMware components, terminology and use, please refer to the wide variety of information on www.vmware.com.

Briefly: VI3 is a bundled product with four components: Virtual Infrastructure Client (VI Client), License Server, Virtual Center Management Server (VC Server), and the ESX Server Console (part of the ESX Server).

An ESX host, which has a hypervisor kernel (vmkernel), runs VMs (Virtual Machine). The service console is actually just a VM itself, which is granted special privileges to access the configuration of the ESX host machine. The figure below shows an example of a typical VI3 deployment with 2 VMs.

Each VM has one or more of its own disk images (VMDK). The collection of VMDKs is managed by the VMFS file system and the ESX service console, a VM with special management privileges.

VI3 also includes a new package, VMware Consolidated Backup (VCB). VCB allows a more centralized approach to off-ESX-host backup, using a special purpose Windows proxy server that accesses ESX data independently. When deployed appropriately, VCB offers many scalability advantages for backing up VMs and VMDK files.

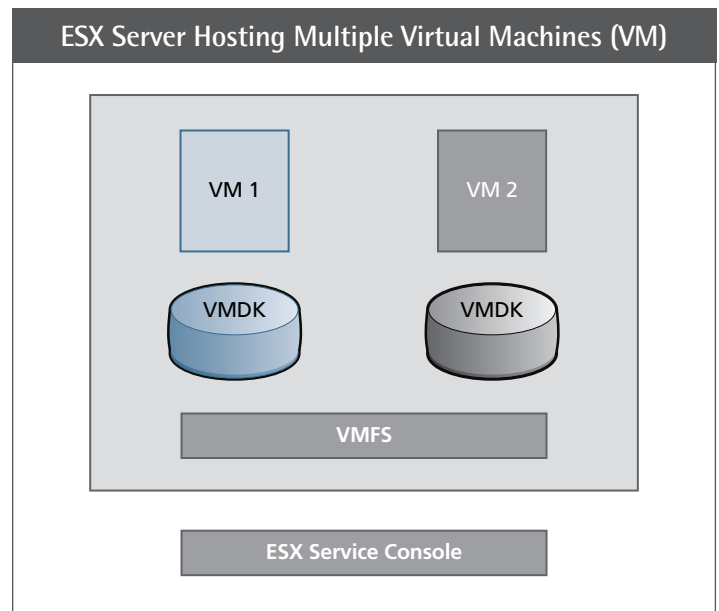


Figure 1. ESX Server Hosting Multiple VMs

A specialized backup proxy server has access to shared VMDK files and their Guest OS file systems. Using this proxy system, backup software can capture VMDK and Guest OS file backups with low application-VM impact and only moderate impact on the ESX server. It will require a proxy server.

3. VMware Infrastructure Backup Alternatives: Success Metrics

There are a wide variety of possible methods for backing up VMware data. To achieve best practices for your site, consider the following characteristics. You will want to find the right balance of simplicity, speed, cost and restorability in the answer you consider.

- ▶ **VMDK recoverability.** The method you choose should allow you to restore a consistent, whole VMDK components. Because the VMDK is a portable, “bare metal” unit of storage that incorporates all Guest OS settings, this is the simplest unit of recoverability for restoring a VM’s state. While normally viewed as a larger form of data, discouraging some admins from backing them up regularly, they offer huge deduplication effects on a Data Domain system. The impact of storing them more often can be substantially mitigated.
- ▶ **Guest OS file recoverability.** According to Gartner, about 80% of recoveries are file recoveries. Especially as VMware moves from lab functions to support of general server and desktop activity, single file restores for a given Guest OS will predominate. Depending on the approach, finding the files to restore can be supported either by a catalog/index using traditional backup software, or through a browsing method to find a file in a familiar namespace, similar to what would be done in conventional NAS system snapshots.
- ▶ **Impact on application VMs.** Except in very small deployments, it is important to minimize the backup impact on production VMs. While most backup software tries to limit server impact to <5% of system resource use, if you have 10 VMs on a single ESX host backing themselves up simultaneously using the same strategies it could use as much a substantial percentage of the physical host resources.
- ▶ **Impact on the ESX server running application VMs.** If the hosting ESX server gets slow, applications get slow, even if it’s not the fault of the Guest OS in a particular VM. Backup best practices in general encourage non-disruption of the application ecosystem.

- ▶ Backup performance. Backup windows matter as much in the virtualized world as in the physical world. Faster is better.
- ▶ Scalability. Make sure the approach you choose can scale to suit your needs as you add VMs and ESX servers.

VMware can make the choices seem complex, since there are a lot of potential alternatives. You could imagine backing up

- ▶ within the VMs,
- ▶ within ESX server running the VMs,
- ▶ from another ESX server (either using a shared storage fabric or through a virtual LUN over Ethernet in a related ESX server),
- ▶ with the V13 VMware Consolidated Backup proxy, or even
- ▶ backing up the underlying SAN or NAS store directly.
- ▶ by scripting file copies and/or snapshots, and/or
- ▶ by using commercial backup software.

While the possibilities seem endless, for most deployers there are really only a couple reasonable alternatives. The table below summarizes them, and they are explained in the following table.

It is important to note some methods that will not be discussed because of their limitations. These are not best practices.

- ▶ Backing up or replicating the backing store only. Most VMware deployments are backed by well-provisioned block storage, often on a SAN. It is possible to backup and restore just the block volumes of that storage. *This is not a best-practice because it is very tricky, or even unlikely, that all the necessary steps will be taken to snapshot a consistent image prior to backing it up.* It is best to back up within the logical image of ESX, the Guest OS and/or VCB.
- ▶ Using client-based deduplicating backup software within the VM, ESX or VCB. The point of VMware is to optimize how much work a server can do across more workloads. Further taxing them for backup is a step backwards. In addition, these specialized applications will not be generally able to support all datacenter backup needs, so it would create an orthogonal,

	METHOD		
	Backup Clients in Both VMs and Service Console	VCB: VMDKs Plus Guest OS Files	VMDK Snap/Copy to Data Domain System from Proxy
VMDK recovery	Yes, if in console	Yes	Yes
Guest OS file recovery	Yes, if in VM	Yes	Yes
File recovery method	Commercial backup software catalog	Commercial backup software catalog	Name-based browsing
Impact on application VMs	High	Low	Low
Impact on ESX server running app VMs	High	Medium	Low
Backup performance	Slow	Moderate	Fast
Scalability	Small deployments only	Scales well w/proxies	Scales well w/proxies
Guest OSs	Any	Windows only	Any

Table 1. Tradeoffs in VMware Data Protection

extra backup administration load. Standard backup software that understands VMware well is more than efficient enough for optimized backup/restore, when paired with a good deduplication system for storage and DR.

In the following section, references are made to commercial backup software, many of which have additional beneficial information on best practices for their use in a VMware environment. While they may evolve over time, these are noted as current papers for followup.¹

VMware:

- ▶ Pre VI3 infrastructure: www.vmware.com/pdf/esx_backup_wp.pdf
- ▶ 3.5 enhancements: www.vmware.com/files/pdf/vcb_35_new.pdf

Enterprise backup software provider examples:

- ▶ Symantec VERITAS: http://eval.symantec.com/mkt-ginfo/enterprise/white_papers/ent-whitepaper_veritas_netBackup_6.5_vmware_nov2007.pdf
- ▶ IBM TSM: www-1.ibm.com/support/docview.wss?uid=swg27009931&aid=1 (copy link, paste to browser)
- ▶ EMC Networker: www.emc.com/collateral/software/data-sheet/h3980-nw-vmware-ds.pdf
- ▶ CommVault Galaxy: www.commvault.com/pdf/CV_SolutionBrief_VMWare.pdf

Specialty VMware backup providers:

- ▶ Vizioncore: www.vizioncore.com/WhitePapers/1-2PunchBackupMethodology_WP_SF5.pdf
- ▶ Tomato AS (Europe-only as of this writing): www.tomato.no/

3.1 Getting Started

Backup the VM Guest OS Files and VMDKs with a Standard Backup Client Without VCB

This method is intuitive for a backup operator and sufficient for smaller deployments. It is very straightforward, uses the current backup solution for the rest of the data center but it does not scale well. It has two parts, each straightforward, and each suited to a different path to recovery. First, by treating each VM as a physical machine, use a conventional backup client in each VM to backup the Guest OS files for easy file recovery. Second, use a standard backup client in the special-purpose Linux VM that runs the Service Console, to backup ESX configuration data and VMDKs as complete system recover points.

3.1.1 Backing up the VM Guest OS Files

This technique is the most well known, and probably most straightforward way of backing up an operating system running within VMware. A standard backup client is simply installed on a VM. The Guest OS backup on the VM is then scheduled and file level backups are performed as with a physical machine. File restores are the same as for any standard client. The backup infrastructure can be shared as just one more component of a standard backup deployment, and it needs no additional hardware.

Using this Guest VM based approach also allows a simple solution for VMs that may be running an application (e.g., SQL, Oracle). The backup application being used will likely have a special agent for these situations that can be used to ensure a consistent point in time backup of the application.

Benefits. File level backups are typically done by the same technique as with other backup clients in the environment. Backups and restores occur just as if the Guest OS was hosted on a physical, not virtual, machine. As long as the Guest OS is supported by the backup software and VMware, the OS can be backed up safely. File level restores are supported. Data Domain fits well this method, providing excellent compression and “WAN optimized” replication for DR purposes. The technique is straightforward to implement. Method 1 also allows for the backup of not only the VM, but also application data that resides on different non-VM controlled storage partitions (such as “raw disk” of NAS).

Considerations. First, managing the sheer number of VMs that can quickly be created is difficult. Second, if a number of clients are backed up concurrently, backups can quickly overrun CPU, memory, and other resources on the ESX host. To avoid potential resource overload, IT managers have to be very selective about how many and which virtual systems they backup at the file level using this approach. The “Guest OS” technique simply does not scale or have a feasible management solution for medium to large VMware installations.

With this step by itself, you cannot backup an ESX host that contains the ESX service console, or the VMDK images. So it is very common to extend this method to include one more step.

3.1.2 Backing up ESX and VMDKs

Each VMware Guest OS has at least one disk file and other associated configuration files within a directory [DataStoreName]/VMname under /vmfs/volumes stored on the ESX host. VMware places the files on top of VMFS and adds a “vmdk” extension to each file. The files can be backed up and restored as standard files. Think of this as a whole VM image backup. For this configuration, a Linux client is installed on the VMware service console. Restoring a virtual machine requires restoring the individual VM folder that contains the VMDK disk files; that is, restoring an image of the VM from the point in time it was shutdown.

¹ These are listed for the benefit of readers only and do not represent any kind of warranty from Data Domain on their quality or accuracy.

Data Domain can take advantage of this method with excellent data reduction effects for local storage as well as WAN replication for DR purposes. Note that Service Console backup only enables the backup and restore of data that resides on the same storage area (such as “disk”) as the VM.

Benefits. This method works with all off the shelf backup software that supports Linux based (ESX) clients. It is easier to manage than the prior approach as only one Linux client exists for each ESX host. You can easily backup the entire ESX host itself if desired. It involves moderate load on the ESX server, but since larger files can be streamed without walking file system metadata to determine incremental changes, it can have less impact than a Guest OS file backup. This method can also help manage the load on any given ESX server because under normal circumstances the VMDK files will be backed up serially rather than concurrently.

Considerations. This method is essentially a disk image backup, which has inherent limitations for recovery granularity. No Guest OS file level restores are enabled from the backup software catalog. One other important consideration is the need to snapshot the VM system whose VMDK file and folder are being backed up; this requires appropriate commands or scripting to ensure image consistency. Separately, some backup applications may not support Linux. Finally, the process of scheduling VM shut downs, shutting down VMs, and scheduling the backups can be a management challenge as the number of VMs increase. However, scripting or the ability to shutdown all VMs on an ESX server at the same time can lessen the burden. If you add VMotion to the configuration then this option becomes even more challenging to manage, especially if you only want to back up selected VMs.

3.2 Enterprise Standard Practice

VMware Consolidated Backup (VCB) with Commercial Backup Software

This method uses a Windows 2003+ host as a backup proxy that is a source for backups. The storage to be backed up is shared with the backup proxy host. It may either use a storage array that is shared by the ESX servers and the proxy, or, as of ESX 3.5’s Virtual LUN construct, the proxy can use a LAN connection to another ESX server to access the storage. For Windows systems, the VCB proxy server can also mount the shared storage for backup access at the file system level if desired. It supports only Windows Guest OSs.

The Consolidated Backup agent, installed as a separate package on the Windows 2003 proxy host, takes a snapshot of an individual VMware virtual machine (or if scripted, multiple VMs), and copies the data to a temporary folder. The data is then used as the source for backups.

Just before the snapshot is made the `/usr/sbin/pre-freeze-script` or `C:\Windows\pre-freeze-script.bat` is run and once complete, `/usr/sbin/post-thaw-script` or `C:\Windows\post-thaw-script.bat` are run. Taking a snapshot only takes a few minutes so, as an example, you

could use the scripts to quiesce or stop your database. The entire snapshot and copy process needs to be understood to automate the process. Variables include size of the disk file and the amount of ESX server resources in use.

A VMDK “snapshot” involves the creation of special VMDK files with the extension `.redo`, which become the writable disk file, while the main `.vmdk` file is now closed for writing (and unlocked by VMFS) and therefore can be backed up from the service console.

While not shown, using the Data Domain system as a Virtual Tape Library is also possible.

Benefits. The most important benefit for VCB is that once the systems are quiesced and the VM snapshot taken, the VMs can be up and running during the actual backup, and the backup processing has minimal impact on their performance. Both VMDK images and Guest OS file images are available for copying to a safe place using scripting or backup software. Multiple virtual machines can be backed up at once with limited load on VMs or an ESX host. This approach is vastly simpler and more scalable, so most of the larger VMware deployments running VI3 or later tend to prefer it. Data Domain can take advantage of this method with excellent local storage deduplication and “WAN optimized” replication for DR purposes.

Considerations. The trade-offs when considering VMware consolidated backup are:

- ▶ Consolidated backup looks very promising, but currently may require some manual command line configuration and implementation. The process can be automated using VMware-supplied scripts. While most backup software vendors are committed to delivering simplified management, this is still evolving.
- ▶ If using shared LUNs from a SAN-attached storage device must have the same LUN numbers assigned to the proxy and to the ESX host. Both the proxy and ESX host must also be on the SAN.
- ▶ As of VI 3, the Consolidated Backup Framework requires a Windows 2003 host as proxy; no other OS is supported.

3.2.1 Consolidated Backup with Data Domain

Here is an example of how to setup a Consolidated Backup proxy and integrate with a Data Domain system. The example assumes that the backup software is installed on a separate host configured for file system backups and that the Data Domain system is properly installed. The connectivity assumes the layout in Figure 2.

3.2.2 Proxy Installation

A requirement is a Windows 2003 server with the VMware Consolidated Backup Framework installed that is directly plugged into a SAN along with one or more ESX servers. Both must directly see all the VM images sitting in the VMFS LUNs. The Windows

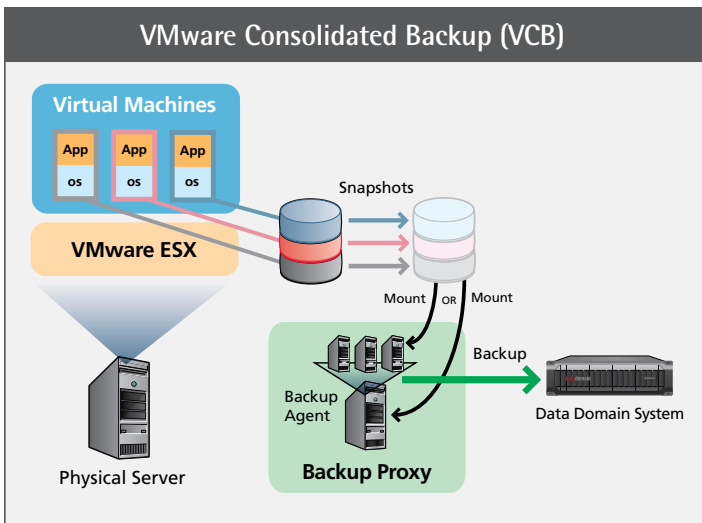


Figure 2. VMware Consolidated Backup with Data Domain

host is called a backup proxy, or proxy for short. The first time a Consolidated Backup occurs by the proxy host, the VMs must be powered on.

After configuration, the proxy talks to the ESX server(s), takes a snapshot, and makes a copy of it to locally attached disk. The process is manual and lends itself to scripting. VMware does provide scripts that let you tie in to some commercial backup products such as EMC/Legato NetWorker or Symantec/VERITAS NetBackup.

3.3 Advanced Best Practice

Snap/Copy VMDK images to Data Domain from ESX or VCB, Restore Guest OS File Copies via Browsing

In the prior methods, commercial backup software was used to identify, move, catalog and restore Guest OS files, VMDKs and even the ESX system itself. They provide complete protection in different levels of scalability. There is one other important approach that is normally used only with scripting. Some backup software providers, notably Tomato in Norway, have begun to offer it as a software package that links well into VMware infrastructure as well as Data Domain systems for local and remote/DR restores.

In this method, only VMDKs are actually copied for protection to the Data Domain system. This may be done from either a Service Console or from a VCB proxy. If these copied VMDK files are named with simple characteristics that allow them to be found easily, such as a combination of VM name, ESX server name, time and date, they can be stored to a network share on a Data Domain system (e.g., using NFS). This would enable simple VMDK recovery, locally or remotely. Once copied to the network share, the data is fully protected. Recovery of a given VMDK back to the ESX storage is just another file copy request.

But it doesn't stop there. It also enables file recovery, without Guest OS file backup. Using the random access properties of the Data

Domain system, a new VM can boot from the saved VMDK on the Data Domain System. Once running, a user can browse into its Guest OS file system, find the file in question, and just copy it back to a production VM's Guest OS. Using Storage VMotion, the whole data image of the VM may be migrated to a primary store if it requires a high rate of interaction (IOPS).

Benefits. This approach has the least impact on running VMs and ESX infrastructure, and it enables the fastest backups of all the methods because there is no requirement for Guest OS file system backups in order to restore discrete files. It is simpler to administer over time because there are fewer files to manage. From a help desk standpoint, the critical files to track are only the larger VMDKs. Unlike VCB, this approach supports any Guest OS, including Linux, Solaris and NetWare.

Considerations. Commercial backup software store the backup images on disk in their respective formats (e.g., tar). While some of these are well understood, they can obscure the internal data from view without use of the catalog or backup application for a restore. As a result, most backup packages, which are making progress on simplifying the scripting requirements of VCB, are not suitable for backups in Method 3. It must be scripted, using the conventions recommended by VMware for snapshotting and restarting VMs for consistent backups. A link to the papers discussing this from VMware's point of view is listed earlier in this section.

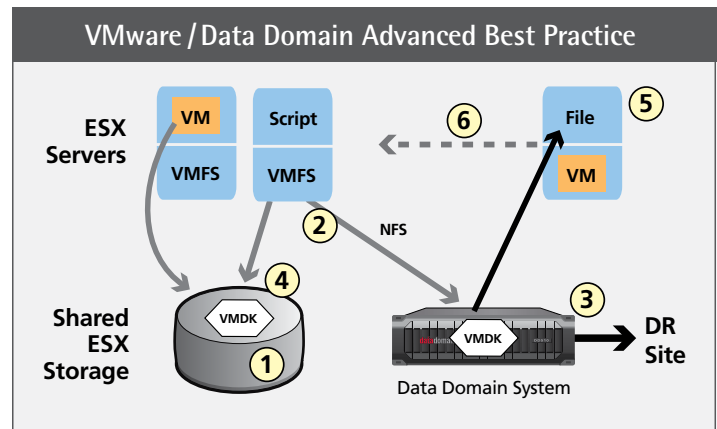


Figure 3. Snap/Copy VMDK Images to Data Domain from ESX or VCB. Restore Guest OS File Copies via Browsing.

- (1) Snap a VMDK Image
 - (2) Copy VMDK to Data Domain System
 - (3) Store / Replicate
 - (4) Restore VMDK from Data Domain System
 - (5) Restore File to Guest OS
 - (6) VMotion to Primary Store
-] Locally or in DR Site

4. Conclusions

In this paper, we presented the predominant approaches to link Data Domain deduplication storage systems to the VMware environment for simplicity, speed and safety of backup retention, recovery and replication.

To gain the benefits of VMware, storage – especially backup storage – often multiplies. Through efficient use of Data Domain storage, backup time can be well managed, and backup storage and replication bandwidth for DR can be brought back under control. While in normal file system backups and retention periods, Data Domain can offer 10-30x data reduction, a VMware environment can often result in 40-60x data reduction.

Be clear about your goals, and consider one of the following choices:

- ▶ To get started, consider using a traditional backup client in each Guest OS, though this will not scale well.
- ▶ For a professional IT deployment that will be supported by most backup software over time, for Windows virtualization, consider VIB VCB.
- ▶ For optimum scalability and efficiency, but with some scripting required in most cases, consider copying consistent VMDKs to a Data Domain system. For VM restore, copy back the VMDK, directly or through Storage VMotion. For file restore, boot the VM from the VMDK on Data Domain, and either copy the file to primary storage or again, use Storage VMotion while the VM is running.

There are a lot of possible choices, but only a few real best practices. For further reading, please visit www.datadomain.com/vmware.

Data Domain
2421 Mission College Blvd.
Santa Clara, CA 95054
866-WE-DDUPE; 408-980-4800
sales@datadomain.com
22 international offices: datadomain.com/company/contacts.html

Copyright © 2008 Data Domain, Inc. All rights reserved.

Data Domain, Inc. believes information in this publication is accurate as of its publication date. This publication could include technical inaccuracies or typographical errors. The information is subject to change without notice. Changes are periodically added to the information herein; these changes will be incorporated in new additions of the publication. Data Domain, Inc. may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time. Reproduction of this publication without prior written permission is forbidden.

The information in this publication is provided "as is". Data Domain, Inc. makes no representations or warranties of any kind, with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Data Domain and Global Compression are trademarks of Data Domain, Inc. All other brands, products, service names, trademarks, or registered service marks are used to identify the products or services of their respective owners.
WP-VMBP-0608