

Using VMware® ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery



Table of Contents

Introduction..... 3

Disk Structure of ESX Server 3

Virtual Machine Disks..... 4

Disk Files and Redo Logs 4

Raw Device Mappings..... 4

Accessing and Managing Virtual Disk Files 4

What to Back Up on ESX Server..... 5

Treating Virtual Machines as Physical Machines 5

Treating Virtual Machines as Files on ESX Server 6

Treating Virtual Machines as Files on Shared Storage 6

Backing Up the Service Console..... 7

Backing Up Virtual Machines as Physical Machines 7

Implementation Steps 7

Creating Backup Copies of Virtual Machine Files on ESX Server 8

Built-in VMware File Operations 8

Using vmkfstools 9

Using vmsnap and vmres 9

Using a Backup Agent on the Service Console 9

Backup Off-Loading..... 12

SAN Image..... 12

Server-Based Replication 12

Using Specialized ESX Server for Backup 12

Appendix A 14

Appendix B: Using vmkfstools 14

Glossary 17

Using VMware® ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery

Introduction

Backup, restoration, and disaster recovery are among the most crucial elements of datacenter management. VMware® ESX Server System and VMware Virtual Infrastructure provides many different capabilities to improve these processes each with its own advantages and disadvantages. Different environments favor different solutions. The flexibility provided by Virtual Infrastructure, allows ESX Server backup and restore procedures into your existing methodologies and procedures. Given the unique nature of VMware technology, though, new and more advantageous ways of approaching these critical tasks may present themselves.

This tech note describes what resources should be backed up on an ESX server and explains all the options available for that backup, including advantages and disadvantages of each option. Lastly, the tech note offers some recommendations based on these trade-offs.

Disk Structure of ESX Server

ESX Server uses VMware ESX Server File System (VMFS) for storage of virtual machines. VMFS is a high-performance file system on physical SCSI disks and partitions, optimized for storing large files such as virtual disk images and the memory images of suspended virtual machines. ESX Server 2 uses VMFS-2, (earlier versions of the product used VMFS-1, which was more limited in its capabilities.) VMFS-2 volumes can span multiple partitions, across the same or multiple (up to 32) LUNs or physical disks. A VMFS-2 volume is a logical grouping of physical disk partitions, which may also be called "physical extents". Because the files stored on the VMFS may exceed 2GB in size, they cannot always be accessed using the same tools as files on a standard ext2, ext3, FAT, or NTFS file system.

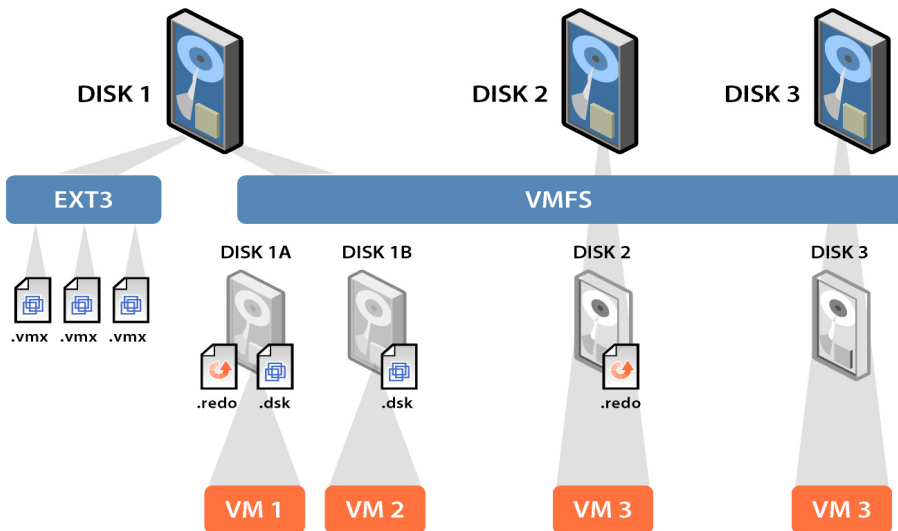


Figure 1

Virtual Machine Disks

As noted above, the disk files of virtual machines are stored on the VMFS file system. These files are in a special format and generally use .disk or .vmdk file extensions. The disk files can comprise all the information the virtual machine stores on the virtual disk or be a symbolic link from a VMFS to a raw LUN when raw device mappings (RDM) are used.

Disk Files and Redo Logs

The simplest use case has the disk files comprising all the information in a virtual machine. In this scenario, redo logs make performing backups much easier. Redo logs make performing backups much easier. In the default state, a virtual machine disk is simply a single file. All changes to that disk are written directly and immediately to that .vmdk file. However, ESX Server can, instead, use what is known as a redo log. When a redo is added to a .vmdk file, that base disk file becomes static and unchanging. All writes are “trapped” in the redo log. This state is represented by file name: if the base disk is called disk.vmdk, the redo log will be called base.vmdk.REDO. A disk file may have a maximum of two redo logs, which is a total of three files: base.vmdk, base.vmdk.REDO, and base.vmdk.REDO.REDO. The format of the redo log is a bitmap record of changes to the disk.

Redo logs are , therefore, useful for disk snapshots. When a disk is represented by the two files disk.vmdk and base.vmdk.REDO, disk.vmdk reflects the state of the drive at the time the disk snapshot was taken while base.vmdk.REDO is a bit-by-bit map of changes to the hard drive since that time.

Raw Device Mappings

With RDM, VMware now has the most flexible storage capability for virtual machines in the industry. This mapping allows all of the features of VMware Virtual Infrastructure to be used in conjunction with raw SAN LUNs. The mapping file—the file that is used to connect the raw LUN to the virtual machine—is what is referenced in the virtual machine’s configuration. The information about the the raw LUN is stored within the mapping file, providing a consistent location for the virtual machine to find its disk across ESX servers even if the LUN is presented differently to each.

There are two modes for RDMs: virtual compatibility and physical compatibility. Virtual compatibility mode allows a mapping to act exactly like a virtual disk file, including the use of redo logs. Physical compatibility mode allows direct SCSI access to the device being mapped for those applications that need lower level disk access and control. In both cases, data is stored on the LUN or SCSI device, not on the disk file.

In both modes, an RDM file in a VMFS volume manages metadata for its mapped device. There is a one-to-one mapping between mapping files and mapped devices. The mapping file is presented to the VMware Service Console as an ordinary disk file, available for file system operations. To the virtual machine, the ESX Server presents the mapped device as a locally attached SCSI device.

In physical compatibility mode, RDM provides minimal SCSI virtualization of the mapped device. In this mode, the VMkernel passes all SCSI commands to the device with one exception: the Report LUNs command is virtualized so that the VMkernel can isolate the LUN to the virtual machine that owns it. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical compatibility mode is useful when there is a need to run SAN management agents or other SCSI target-based software in the virtual machine. This mode is also useful for virtual-to-physical clustering for cost-effective high availability.

Accessing and Managing Virtual Disk Files

Virtual disk files on the VMFS are accessible through ESX Server Service Console, ESX Server MUI, VirtualCenter, and VMware SDK.

From the service console, files can be viewed and manipulated on VMFS volumes under the /vmfs directory with ordinary file commands, such as ls and cp. Although mounted VMFS volumes may appear similar to any other file system, such as ext3, VMFS is, as noted, primarily intended to store large files, such as disk images. The ftp, scp, and cp commands can be used for copying files to and from a VMFS volume as long as the host file system supports these large files. The nfs protocol is subject to a 2G file size limitation and should not be used.

Additional file operations are enabled through the vmkfstools command. This command supports the creation of a VMware ESX Server file system (VMFS) on a SCSI disk and can be used to create, manipulate, and manage files stored in VMFS volumes. This command is also used to list files on the VMFS volume, add a redo log, commit a redo log, and export .disk files into other formats.

The vmsnap and vmres scripts that automate many of the common backup and restore tasks are run from the service console. This will be discussed in a later section of this document.

Importing and exporting disk files can also be done through the ESX Server MUI by copying the files from VMFS mount and pasting them to a partition running ext3 file system (the file system common for Linux).

What to Back Up on ESX Server

When thinking about what to backup on the ESX Server the user may want to identify the items that need to be address for recovery. With ESX Server, there are three major components you may need to back up:

- Virtual disks
- Virtual machine configuration files
- The cconfiguration of the ESX Server itself

Note that Aall the information normally backed up in the enterprise infrastructure, including the operating system, applications, and data, is included in the virtual disks.

Because a virtual machine is just like a physical machine, one possible approach is to back it up in the same manner as a physical machine, using backup software running inside a virtual machine. Although backing up virtual machine configuration files goes beyond backing up a physical infrastructure, doing so makes recovery easier. This method allows for traditional incremental and differential backups, though, as will be noted later, there are some disadvantages.

With only a few files encapsulating an entire virtual machine, it is very simple to get back to a previous known state at a known time. Thus, another possible approach is to back up all the files on the ESX Server that make up a virtual machine.

This approach also allows for easy check pointing without any additional third-party software or hardware. Two levels of redo logs can be created, allowing maintenance of multiple snapshots. The use of redo logs allows for hardware-independent snapshotting of virtual machines, yielding true point-in-time copies without the use of SAN features.

Yet another possibility is off-line backups where files encapsulating virtual machines are accessed and backed up without loading the ESX Server that the virtual machines normally run on. Storage replication must be used to assist in this process.

Each approach to backing up virtual machines has its own advantages. Available backup methods also depend on the selected format of the virtual machine disks. In many cases, it is best to employ a combination of both methods for the best results.

ESX Server Service Console is installed with the ESX Server installation and rarely changes. However, it may still be a good practice to back up the console to streamline the recovery process.

Treating Virtual Machines as Physical Machines

Virtual machines provide complete guest operating systems (OS) on virtualized hardware. These OS installations can be backed up in the same ways as their physical counterparts. First, one may attach backup hardware to the virtual machines and install backup servers in them. Alternately, one may install a backup agent within each VM, and back up data over the network to other backup servers. Finally, one might copy data manually or with a script to another machine. Backing up a virtual machine in this way is precisely like backing up a physical machine.

Advantages

The major advantage of this approach is that you can use consistent methodology to back up all the servers in your data center, whether physical or virtual. This approach also allows you to do a file-level backup and restoration and gives more flexibility in the choice of backup software. Backup software that has backup modules specific to the applications can be used to quiesce the applications for the duration of the backup. This method also makes incremental and differentials easier.

Disadvantages

The major disadvantage of this approach is that you cannot take advantage of the encapsulation of virtual machines into one or more discrete files. By backing up individual files within the guest operating system as files, you do not have the ability to back up and restore the virtual machine as a whole. Also, note that backup loads of 5% may be acceptable on one physical host but multiply that by number of guests and now you have real load.

Treating Virtual Machines as Files on ESX Server

This approach takes advantage of the service console's ability to see each virtual machine's virtual disk as a file. ESX Server creates one file per virtual machine disk device (with redo logs stored separately). These .dsk or .vmdk files can be backed up, effectively backing up an entire virtual hard drive in a single pass. This approach is not possible if you are using RDM disks for your virtual machines. To take advantage of this method, you must add and commit REDO logs to the base disk of each VM as the base disk is backed up.

Advantages

This approach takes a very simple view of backups and restorations that allows you to back up or restore an entire virtualized server in one step. With disk snapshots and redo log manipulation, you can do near-line backups as explained in further detail in Appendix A.

Disadvantages

The main disadvantages of this approach are that, although greatly simplified compared to a traditional backup file-level restorations. If you have extremely large

virtual machines, you might have to restore tens of gigabytes of data to restore a single file. Because this backup process treats the virtual disk as a whole and is not application aware, the backups created through such a process are crash consistent. Finally, the need to access large (>2GB) virtual disk files on VMFS may limit your choice of qualified backup software or necessitate additional intermediate processing. Note that all supported backup vendors can handle > 2GB files See Appendix B for details.

Treating Virtual Machines as Files on Shared Storage

When virtual machine files reside on shared storage, it is possible to use SAN-based imaging or an independent backup server to back up virtual machine files without creating an additional load on the ESX Server where the virtual machines normally run.

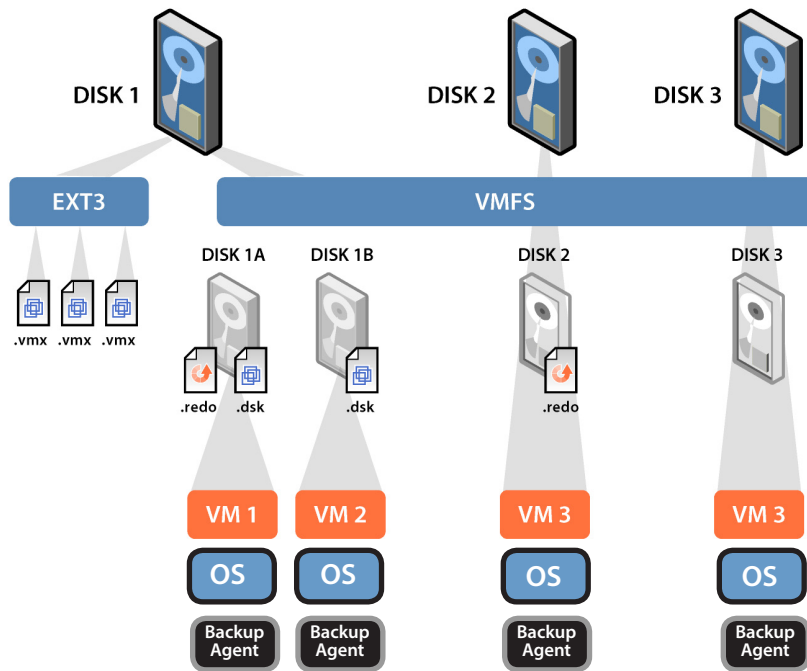


Figure 2

Backing Up the Service Console

The service console is a specialized virtual machine running Linux, which provides a management interface (both command prompt and Web-based) to the ESX Server and the virtual machines running on it. The service console is nearly stateless, however, having it safely backed up may save time in a disaster-recovery scenario. You may not find it necessary to back up the service console as frequently as the virtual machines and their associated data.

Backing Up Virtual Machines as Physical Machines

Compatible with the following virtual driver formats:

- Virtual disks (all formats)
- RDM disks (all formats)

Issues to Consider

The following issues need to be considered before you decide on a backup method:

- Whether the backup server is on a physical or virtual machine
- The network configuration
- Type of long term storage

Implementation Steps

- 1) Setup.
 - a) Install the backup agent of your choice on each virtual machine to be backed up.
 - b) Schedule the backups and manage the tapes as directed

by the documentation for your backup program.

c) Configure your backup server or node. If the backup server is installed in a virtual machine on the ESX Server, the following steps are required.

- i) Attach the tape drive or library to a SCSI port on the ESX Server.
- ii) Through the ESX Server Web management interface, assign the tape drive to a virtual machine.

d) Configure the virtual machine to use the tape drive or library, installing the appropriate drivers and backup server of choice.

2) Ensure that networking is configured for access between the backup server and virtual machines to be backed up. If both virtual machines to be backed up and the backup server are on the same ESX Server, you may use a private virtual network switch to connect them to each other.

3) Install the backup agent on all virtual machines to be backed up.

4) Backup and restore.

- a) Follow the instructions for the backup software you installed.

Example architecture: an example of the architecture described above implemented with Legato Networker

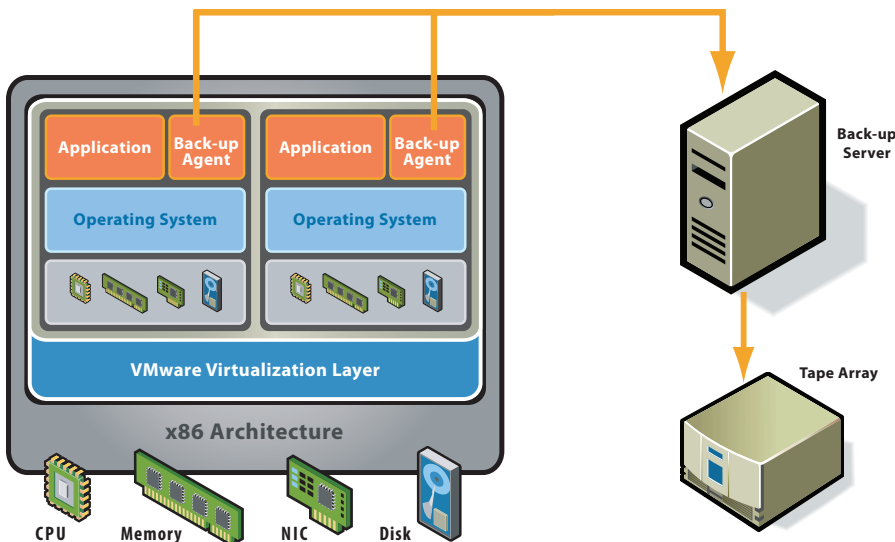


Figure 3

Creating Backup Copies of Virtual Machine Files on ESX Server

Each virtual machine is represented by up to three virtual disk files that live on the VMFS and the virtual machine configuration file which resides on the file system managed by the ESX Server Service Console. Disk files are frequently larger than 2GB, and not all backup programs can directly access its files. With this stated, most current backup product on the market today do not have this issue.

By default, Virtual disk files cannot be backed up while the virtual machine is powered on because, during this time, the disk file is open and being written to by a virtual machine. Powering off or suspending the virtual machine closes the virtual disk file and makes it safe to back up. This may not be a tenable solution for most enterprise situations. Fortunately, with the REDO log capability, one can add a REDO log to the VM. All new writes are trapped in the REDO log, making the base virtual machine disk file static and available for back up. Doing so gives you a point-in-time, crash-consistent snapshot of the virtual disk to back up. Alternatively, you can take a snapshot the virtual machine. This captures all new writes to a redo log, making the base virtual machine disk file .dsk available for back up. Doing so gives you a safe and consistent snapshot of the virtual disk to back up.

Issues to Consider

- Can you shut down the virtual machine for the duration or to take a snapshot?
- Do you prefer using backup tools or built-in VMware file operations?
- Is your backup tool of choice supported when running inqualified with the service console?

Built-in VMware File Operations

Compatible with the following virtual driver formats:

- Virtual disks only (all formats)

Virtual disks and other virtual machine files, such as the configuration file, logs, and memory, can be manipulated from ESX Server Service Console command line through scripting API commands and VMware Virtual Infrastructure SDK.

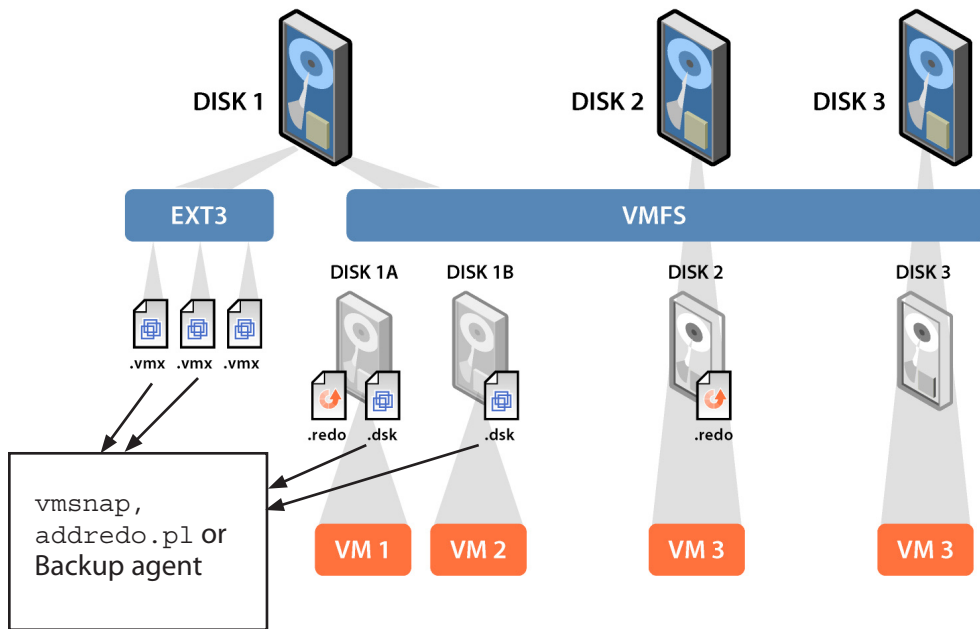


Figure 4

Using vmkfstools

VMware provides a tool called `vmkfstools`, which handles files larger than 2 GB and guarantees good read and write performance to the VMFS partition. The `vmkfstools` commands are documented in the *ESX Server Administration Guide*¹. `vmkfstools` can be used to convert virtual disks from ESX Server format (also known as COW) to the format used by GSX Server and VMware Workstation and vice versa. ESX Server disks are optimized for use on the VMFS, while GSX and Workstation disks must be in a less-efficient format due to the limitations of other file systems. `vmkfstools` allows portability between those formats.

One side effect of the conversion is a potential saving of disk space. For performance reasons, ESX Server virtual disk files are always the size of the space allocated. When `vmkfstools` exports these files into the other format, it only copies the data that appears to be in use on the virtual hard drive. To enable this behavior, use the shrink option in the VMware Tools control panel on the guest OS before exporting the disk. To open the VMware Tools control panel, double-click the VMware Tools icon in the virtual machine's system tray for Windows, or run `vmware-toolbox` in Linux. The Shrink option does not change the size of the virtual disk on the VMFS volume, but it zeros out all unused parts of the virtual disk file. This enables `vmkfstools` to export disks that are the size of the used space rather than the allocated space. This can save time, network bandwidth, and tapes. This mode cannot be used on redo logs, which are already effectively in the converted form.

Note: The `vmkfstools -export` command cannot be used to segment redo files. In practice, it means that redo files should be kept smaller than 2GB to ensure compatibility if you plan to back them up. Alternatively, you may choose to commit the redo log before exporting virtual disks for backup or for transfer to VMware Workstation or VMware GSX Server.

In addition, `vmkfstools` can write to standard output as well as to a file. By putting `vmkfstools` into backup scripts and piping its standard output into backup commands that read from standard input, the need for intermediate storage of virtual disk files is eliminated.

Using vmsnap and vmres

The `vmsnap` and `vmres` Perl scripts available from VMware as a stand-alone download use scripting API commands. As of ESX Server 2.5 these are shipped with the product. You may want to explore `addredo.pl` as well for these functions. The `vmsnap` script performs ESX Server snapshots of virtual disks and makes backups of disks, config, log, and `nvram` files with little or no virtual machine downtime. This is also called a "hot" backup. It may then export the virtual disk into a `.vmdk` format

where each of the virtual disk files is subdivided into segments smaller than 2GB. Such segments can be easily copied to a drive with a commonplace operating system, NAS, or archiving media.

Because snapshotting virtual disks and advanced disk modes, such as "undoable" or "append," uses the same underlying mechanism, only disks that are initially in the persistent mode can be archived with the `vmsnap` script. This script operates on running virtual machines only. You can specify a local or remote archiving directory or pipe the script output as an input of specialized archiving software. When you specify a local archiving directory, `vmsnap` skips the step of subdividing the files into 2GB segments.

If you use VMware Tools to shrink the base virtual disk before exporting it, you export smaller files although the VMFS files remain the same size. Having these smaller files can speed up backup and restoration if you are using only a fraction of the disk space allocated for the virtual machine.

Note: Redo logs are exported as a single file.

The `vmres` script supports restoration of the previously archived virtual machine to either the same or a different ESX Server.

A detailed description of the command options for `vmsnap` and `vmres` scripts is available in Appendix A.

Using a Backup Agent on the Service Console

There are a number of backup agents that have been qualified for use with the ESX Server service console. The most current list is available at:

http://www.vmware.com/pdf/esx_backup_guide.pdf

Any backup software capable of handling files over 2GB can be used to back up virtual machine files, including virtual machine configuration files (`.vmx`), the memory state of suspended virtual machines (`.vmram`), log files, the virtual disks (`.disk`, `.vmdk`, `.REDO`), and the service console.

Advantages

This is similar to a standard network backup situation and should fit into most methodologies easily. This method can also be used to back up the service console.

Disadvantages

The main disadvantage of this method is that it places a heavy load on the service console because the service console has to run a backup agent and has to write all the virtual machine data over the network. In addition, this can create a heavy network load as the virtual disk files grow. You might need a private backup network to handle the load.

¹<http://www.vmware.com/support/esx25/doc/admin/index.html>

If you are backing up to a local tape drive, network is not an issue. You should not run a backup server directly on the ESX Server Service Console. Instead, run the backup server in one of the virtual machines and assign the SCSI tape drive to that virtual machine as described above.

In addition, unless vmsnap tool is used as an intermediate step, virtual machines need to be powered off for the duration of the backup.

Implementation Steps

1. Install the backup agent on the service console.
2. Install the backup server on a separate platform, which could be a physical or a virtual machine.
3. Ensure network connectivity between the service console and the backup server.
4. Make sure the VMFS is mounted on the service console.
5. Schedule the backups and manage the tapes as directed by the documentation for the backup program.
6. Schedule the virtual machine power-down and power-up cycle to match the backup schedule.

Example Architecture: using vmsnap and Veritas Netbackup

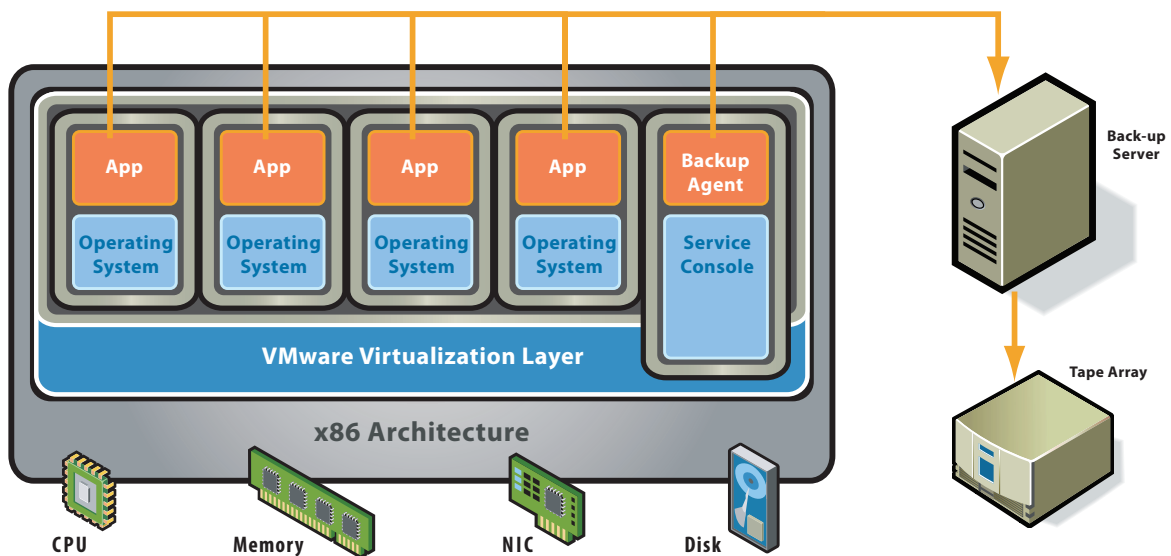


Figure 5

You can create backups of virtual machines without shutting them down using the following procedure:

1. Make sure that the Netbackup DataCenter server is configured either in a physical machine or in a virtual machine and that there is network connectivity between the server and the ESX Server Service Console.
2. Install Netbackup Linux RedHat 2.2 client off the Netbackup DataCenter Linux CD ROM.²
3. Run `vmsnap.pl` to create a local backup of the virtual machine.

```
[root@localhost VM1winXPPro]# vmsnap.pl -c /root/vmware/VM1winXPPro/winXPPro.vmx
-d /vmimages/winXPProbackup -l
vmsnap: VM config file is '/root/vmware/VM1winXPPro/winXPPro.vmx'
vmsnap: Local Directory is '/vmimages/winXPProbackup'
vmsnap: This VM will only be backed up locally.
vmsnap: Disks found that are in use:
vmsnap: Disk found: scsi0:0 (vmhba0:1:0:1:VM1_Empty.vmdk)
vmsnap: 1: scsi0:0
vmsnap: Creating the RedoLogs for cfg : /root/vmware/VM1winXPPro/winXPPro.vmx
vmsnap: scsi0:0:Adding REDO
vmsnap: scsi0:0: BACKUP STARTING NOW...
local dir : /vmimages/winXPProbackup has 656960k space/vmfs/vmhba0:1:0:1/VM1_Empty.vmdk size : 2048000 k
```

In this example, you are backing up a virtual machine with the configuration file located in `/root/vmware/VM1winXPPro` locally to a directory `/vmimages/winXPProbackup` without segmenting the disk files. The backup includes a virtual machine configuration file, log files, and the disk file.

4. In the Netbackup DataCenter administrative console, specify a new policy that points out the ESX Server hosting the virtual machines to be backed up.
5. Enter the fully qualified path to the files of vmsnap output, in this case:
`/vmimages/winXPProbackup/WinXPPro/*`
6. Start Manual Backup by right-clicking on the new policy you just created.

² Refer to <http://www.vmware.com/pdf/veritasnetbu.pdf> for more details on configuration of Veritas Netbackup with ESX Server.

Backup Off-Loading

While each virtual machine is encapsulated in a few files that are visible and can be backed up from the ESX Server Service Console, doing so creates a noticeable load on the service console and can affect performance of the production server. It is possible to minimize impact on the production system by off-loading backup to specialized dedicated systems. This can be done with the snapshot functionality of storage area networks, using server-based replication or using built-in snapshotting for virtual machine disks.

SAN Image

Compatible with the following virtual driver formats:

- Virtual disks (all formats)
- RDM disks (all formats)

If your virtual disk files are stored on a SAN, you can use features supplied by your SAN vendor to create a copy of your production LUN, containing all virtual disks. These copies can then be sent to your backup media. With this method, you do not have to use virtual machine snapshotting functionality during the backup process because the SAN snapshot guarantees consistency. Snapshotting and replicating SAN volumes requires the use of layered applications with SAN. ESX Server is compatible with some SAN layered applications but not all as standards are still evolving. You will want to check the Support Matrix for the latest supported applications.

Server-Based Replication

Unlike SAN-based replication, this approach requires a replication agent in each of the virtual machines. Because of this, the replication agent has visibility into the file system and the application on the virtual disk and is capable of incremental replication where only the files that have been modified since the time of the previous replication are copied. On the other hand, you cannot take advantage of the encapsulation of virtual machines into few files. A replication server and replication targets are required.

Compatible with the following virtual driver formats:

- Virtual disks (all formats)
- RDM disks (all formats)

Using Specialized ESX Server for Backup

ESX Server's built-in snapshot function can also be used in the backup off-load scenario. You can create a snapshot of a virtual disk located on shared storage by adding a redo log. Thereafter, ESX Servers can access the virtual disk snapshot. A designated backup ESX Server can then access the disk snapshot and back it up using a backup agent installed in the backup ESX Server Service Console.

Advantages

Consistent methodology is used for all SAN-stored data, including VMFS volumes. The backups do not overload the production servers' service consoles or their networks as the heavy workload is carried by the dedicated backup ESX Server, you have the option of using attached tapes. This method can also be used with virtual disks in RDM format in both physical and virtual compatibility modes.

Disadvantages

The above backup methods are only possible if shared storage is in place. In addition, image-based backup is not application aware. Thus, the backup image is "crash consistent." That is, the image is equivalent to the image left on the system if the power is unexpectedly cycled. Some transactions in progress may not complete. In addition, image-based backup does not allow for file-based restoration. This can increase restoration time.

Implementation Steps

- 1. Using the SAN management software, schedule snapshots on the disk backend.
- 2. Schedule the backups of the SAN snapshots and manage the tapes as directed by the documentation for your backup program.

Example Architecture: using SDK snapshot interface and helper ESX Server

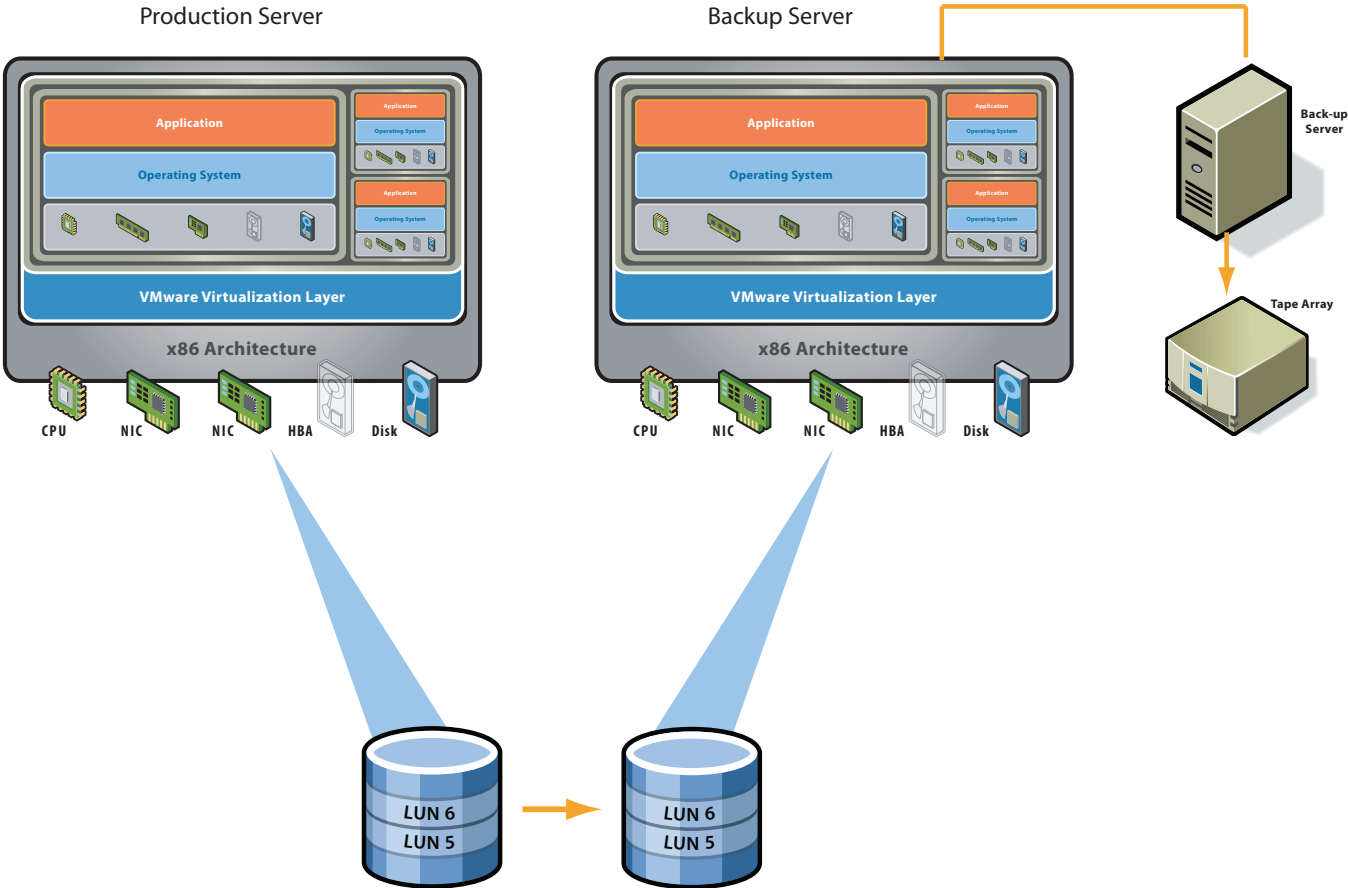


Figure 6

Appendix A

Vmsnap and vmres description and command options

Vmsnap DESCRIPTION

This tool must be run as the “root” user or as the owner of the virtual machines to be backed up. This tool enables a hot backup of a virtual machine’s virtual disk through the use of redo logs as the snapshots (disk persistent mode).

For maximum flexibility, the .dsk file is converted to .vmdk files. These .vmdk files are copied to the Archive Server and stored in /vmimages/archivebackup/\$guest or a user-specified location where \$guest is the name of the virtual machine.

The associated .cfg file (ESX Server 1.5.x) or .vmx file (ESX Server 2.x) is copied to the Archive Server.

The default local_dir is /vmimages/local-backup and default remote_dir is /vmimages/archivebackup.

Command-Line Options

vmsnap [-a server_name] [-c config_file] [-d local_dir] [-R remote_dir] [-g] [-h] [-l] [-m] [-r]

-a server_name	Specify an archive server.
-c config_file	Specify a virtual machine configuration file to use for vmsnap.
-d local_dir	Specify the local directory for backup.
-R remote_dir	Specify the remote directory for backup on the archive server.
-g	List all available guest virtual machines for backup.
-h	Help.
-V	Version.
-l	Perform local backup only.
-m	Generate the man page for this program.
-r	Commit the redo logs in case they are already present.
vmres	Command options.

Appendix B: Using vmkfstools

Example: Using vmkfstools

This example illustrates migrating a virtual machine’s virtual disk file from ESX Server to VMware GSX Server or VMware Workstation and then migrating the virtual disk back to ESX Server.

```
vmkfstools -e winXP.vmdk vmhba0:6:0:1:winXP.vmdk
```

The preceding command exports the winXP.vmdk virtual disk file to one or more .vmdk files of maximum size 2GB that you can use as a virtual disk in a virtual machine on GSX Server or Workstation. The resultant winXP.vmdk files can reside on a VMFS volume, or an /ext2, /ext3, or NFS file system.

The following example imports a GSX Server or Workstation virtual disk file into the VMFS volume on the specified SCSI device.

```
vmkfstools -i winXP.vmdk vmhba0:6:0:1:winXP.vmdk
```

By contrast, if you are importing directly into a raw partition, the example becomes:

```
vmkfstools -i winXP.vmdk vmhba0:6:0:1
```

Listing the Files on the VMFS of the Specified Device

```
vmkfstools -l vmhba0:2:0:0
```

This command illustrates listing the contents of the file system, including redo logs, virtual disk files, and swap files on target 2 of SCSI adapter 0.

Note: You can also do most of the vmkfstools operations through the VMware Management Interface.

vmkfstools Command Syntax

Note: You must be logged in as the root user to run the vmkfstools command.

Specifying a VMFS Volume or File

The format for the vmkfstools command when specifying a VMFS volume or file is:

```
vmkfstools <options> <path>
```

where <path> is an absolute path that names a directory or a file under the /vmfs directory.

For example, you can specify a VMFS volume by a path such as:

```
/vmfs/vmhba1:2:0:3
```

You can also specify a single VMFS file:

```
/vmfs/lun1/rh9.vmdk
```

vmkfstools Options

This section includes a list of all the options used with the `vmkfstools` command. Some of the tasks in this section include options that are suggested for advanced users only. These advanced options are not available through the VMware Management Interface.

Basic vmkfstools Options

Basic options are common tasks that you may perform frequently. You may also perform them through the management interface.

Lists the Attributes of a VMFS Volume or an RDM

```
-P --querypartitions <VMFS_volume_name>
```

```
-P --querypartitions <VMFS_volume:fileName>
```

For a `VMFS_volume_name`, the listed attributes include the VMFS version number (VMFS-1 or VMFS-2), the number of physical extents (partitions) comprising the specified VMFS volume, the volume label (if any), the UUID (if any), and a listing of the SCSI device names of all the physical extents comprising the VMFS volume.

For a `VMFS_volume:fileName`, the listed attributes include the `vmhba` name of the raw disk or partition, corresponding to the mapping referenced by `fileName`, and any identification information for the raw disk.

Exports the Contents of the Specified File on the Specified SCSI Device to a Virtual Disk on the File System of the Service Console

```
-e --exportfile <dstFile>
```

After the export, you may transfer the virtual disk to another server machine and import it to a SCSI device on the remote machine. If your virtual disk has redo logs, you have the following options:

- If you use the `exportfile` option on the base virtual disk, only the base virtual disk is exported. Any uncommitted redo logs are not exported but can be copied separately.
- If you use the `exportfile` option on a ESX Server redo log, the exported virtual disk contains the redo log, any previously created redo logs, and the base virtual disk. That is, the newly created exported virtual disk appears as if the redo logs were committed to its base virtual disk.

Note: However, your original source redo logs and base virtual disk remain unchanged.

- If you want to export your redo logs and base virtual disk separately, then use the `exportfile` option to export the base virtual disk and the `cp` command to export each redo log separately.

Use the combination of `exportfile` and `importfile` together to copy VMFS files to remote machines. The virtual disk should take less space than the full size of the VMFS file because the virtual disk does not include zeroed sectors of the VMFS file.

Imports the Contents of a VMware Virtual, Plain, or Raw Disk on the Service Console to the Specified File on the Specified SCSI Device

```
-i --importfile <srcFile>
```

This command is often used to import the contents of a VMware Workstation or VMware GSX Server virtual disk onto a SCSI device. You may also run this command to import a virtual disk that was created by exporting the contents of a disk from another SCSI device.

Note: The destination device must have space for the entire size of the virtual disk, even if it is mostly free space, as the complete contents of the source disk are copied.

Caution: The `vmkfstools` command may fail when attempting to import plain disks created with GSX Server version 2.5 or earlier. If `vmkfstools` returns an error when importing a plain disk, see the *ESX Server Administration Guide*.

Lists the Files on the File System on the Specified Device

```
-l --list
```

```
-h --human-readable
```

```
-M --verbosemappings
```

The output includes permissions, sizes, and the last modification time for redo logs, virtual disk files, and swap files. You can use the `-h` option to print the sizes in an easier-to-read format, for example, 5KB 12.1MB and so on. The `-M` option lists the `vmhba` name that corresponds to each RDM.

Sets the Name of the VMFS on the Specified SCSI Device

```
-S --setfsname <fsName>
```

You can see the VMFS name by running the `vmkfstools` command with the `-l` option: `vmkfstools -l`.

Advanced vmkfstools Options

Advanced options are tasks that you may perform infrequently. These tasks are not available through the management interface or are available in a limited form and are suggested for advanced users only.

Commits the Redo Log of the Specified File, Making the Associated Changes Permanent

```
-m -commit
```

If a virtual machine is in undoable or append mode, then the redo log is created automatically. The name of the redo log is derived by appending `.REDO` to the name of the file that contains the base disk image. You can commit the changes to the disk that are stored in the redo log by using the `commit` option or eliminate the changes by using the `rm` command to delete the redo log file.

Maps a Raw Disk or Partition to a File on a VMFS-2 Volume

```
-r --maprawdsk <raw-SCSI-device>
```

Once this mapping is established, you can access the raw disk just like a normal VMFS file. The file length of the mapping is the same as the size of the raw disk or partition. The mapping can be queried for the raw SCSI device name by using the `-P` option. By mapping a raw disk or partition to a file, you can manipulate this raw disk or partition as any other file. All VMFS-2 file-locking mechanisms apply to raw disks.

Displays Disk Geometry for a VMware Workstation or GSX Server Virtual Disk

```
-g -- geometry <virtual-disk>
```

The output is in the form: **Geometry information C/H/S is 1023/128/32**, where **C** represents the number of cylinders, **H** represents the number of heads, and **S** represents the number of sectors.

When importing VMware Workstation or VMware GSX virtual disks to VMware ESX Server, you may see a disk geometry mismatch error message. A disk geometry mismatch may also be the cause if you have problems loading a guest operating system or running a newly created virtual machine.

View the events log through the VMware Management Interface (Users and Events page for the virtual machine) or through the service console (the `vmware.log` file, found, by default, in the `<user>/vmware/<guest_operating_system>` directory). Look for **C/H/S** and compare this with the output of the `vmkfstools -g` command.

If the disk geometry information is different, then specify the correct information from the output of the `vmkfstools -g` command in the configuration file of the newly created virtual machine.

See the *ESX Server Administration Guide* for complete details on specifying the disk geometry in a virtual machine's configuration file.

Extends the Specified VMFS to the Specified Length

```
-X --extendfile # [gGmMkK]
```

Use this command to extend the size of a disk allocated to a virtual machine after the virtual machine has been created. The virtual machine that uses this disk file must be powered off when you enter this command. Also, the guest operating system must be able to recognize and use the new size of the disk, for example, by updating the file system on the disk to take advantage of the extra space. Specify the size in kilobytes, megabytes, or gigabytes by adding a suffix of **k** (kilobytes), **m** (megabytes), or **g** (gigabytes), respectively.

Manages SCSI Reservations of Physical Targets or LUNs

```
-L --lock [reserve|release|reset]
```

Caution: Be careful when using these commands. The `reserve`, `release`, and `reset` commands can interrupt the operations of other servers on a SAN, so use these commands with great caution.

The `-L reserve` command reserves the specified raw disk or the disk containing the specified VMFS volume. After the reservation, other servers get a SCSI reservation error if they attempt to access that disk, but the server that did the reservation can access the disk normally. The `-L release` command releases the reservation on the specified disk or disk containing the specified VMFS volume. Any other server can access the disk again. The `-L reset` command does a SCSI reset to the specified disk. Any reservation held by another server is released.

Recovers a VMFS

```
-R -recover
```

This command enables you to recover a VMFS accessible by multiple ESX servers when other `vmkfstools` commands indicate that the file system is locked by another ESX Server machine, but, in fact, no other server is currently accessing this file system. This situation may occur if the VMFS was being accessed by a server, for example, running a virtual machine, and that server crashed.

Note: You should only use this command if you are certain that no other ESX Server is still accessing the file system.

Scans the Specified vmhba Adapter for Devices and LUNs

```
-s --scan <FC_SCSI_adapter>
```

Note: We recommend that you use the `cos-rescan.sh` command rather than this option.

This option is useful for adapters connected to SANs, particularly if you are reconfiguring your SAN. If a new device or LUN becomes accessible through the adapter, then ESX Server registers this new virtual device for use by virtual machines. If an existing device or LUN is no longer used and appears to be gone, then it is removed from use by virtual machines.

Note: Only use this `-s` option for Fibre Channel adapters.

You can see the results of the scan by using `ls /vmfs` or looking at the contents of `/proc/vmware/scsi`.

Create or Resize a Swap File in a VMFS Volume of the Specified SCSI Device

```
-k --createswapfile # [gGmMkK]
```

The size is specified in bytes by default, but you can specify the size in kilobytes, megabytes, or gigabytes by adding a suffix of `k` (kilobytes), `m` (megabytes), or `g` (gigabytes), respectively.

Note: You must be logged in to the Service Console with root user permissions to create a swap file.

You can resize an existing swap file by specifying the new file size as an argument to the `-k` option:

1. Deactivate the swap file, if it is active, with `vmkfstools -y`.
2. Resize the swap file with the `-k` option.
3. Activate the swap file with `vmkfstools -w filename`.

If you try to resize an active swap file, ESX Server returns an error message. ESX Server does not automatically activate a swap file after it is created. Use `vmkfstools` with the `-w` option to activate a swap file. You can set a swap file to be activated automatically after a system reboot with the Activation Policy option of the Swap Management section in the Options tab of the Management Interface.

Imports the Contents of a Virtual Disk to the Specified File on a SCSI Device

```
vmkfstools -i ~/vms/nt4.vmdk
vmhba0:2:0:0:nt4.vmdk
```

This example illustrates importing the contents of a virtual disk that contains Windows NT 4.0 from the service console's file system to a file named `nt4.vmdk` on target 2 of SCSI adapter 0.

Glossary

Terminology

In other contexts, such as the VirtualCenter client console, raw device mapping may be described in terms such as "Mapping a VMFS volume into a datastore," "Mapping a system LUN," or "mapping a disk file to a physical disk volume." These terms all refer to raw device mapping.

The following terms are used in this document or related documentation:

Raw Disk — A disk volume accessed by a virtual machine as an alternative to a virtual disk file; it may or may not be accessed through a mapping file. This has been replaced with the introduction of RDM devices.

Raw Device — Any SCSI device accessed through a mapping file. For ESX Server 2.5, only disk devices are supported.

Raw LUN — A logical disk volume located in a storage area network (SAN).

Mapping File — A VMFS file containing metadata used to map and manage a raw device.

Mapping — An abbreviated term for a raw disk mapping.

Mapped Device — A raw device managed by a mapping file.

Metadata File — A mapping file.

Compatibility Mode — Type of SCSI device access virtualization (physical or virtual).

References

RDM

http://www.vmware.com/pdf/esx/esx25_rawdevicemapping.pdf

ESX Server Administration Guide, pp290 for `vmkfstools`

Backup Compatibility Guide



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
© 2005 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156 and 6,795,966; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective companies.

