


# Virtual Machine Backup Guide

ESX Server 3.0 and VirtualCenter 2.0



Virtual Machine Backup Guide

Revision: 20060615

Item: VI-ENG-Q206-216

You can find the most up-to-date technical documentation at:

<http://www.vmware.com/support/pubs>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

© 2006 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, 6,795,966, 6,880,022, 6,961,941, 6,961,806 and 6,944,699; 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.

All other marks and names mentioned herein may be trademarks of their respective companies.

**VMware, Inc.**

3145 Porter Drive  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

# Contents

Preface .....	vii
About This Book .....	viii
Intended Audience .....	viii
Document Feedback .....	viii
VMware Infrastructure Documentation .....	viii
Conventions and Abbreviations .....	ix
Technical Support and Education Resources .....	ix
Self-Service Support .....	ix
Online and Telephone Support .....	x
Support Offerings .....	x
VMware Education Services .....	x
Chapter 1 Introduction .....	1
Concepts .....	2
What to Back Up .....	2
Virtual Machine Contents .....	3
Virtual Machine Disks .....	3
Raw Device Mappings .....	4
Accessing and Managing Virtual Disk Files .....	4
Backup Components and Approaches .....	5
Backup Approaches .....	6
Using Traditional Backup Methods .....	6
Traditional Backup Considerations .....	7
Backing Up the Service Console .....	7
Backing Up Virtual Machines .....	8
Backup Client in a Virtual Machine .....	8
Backup Client in the Service Console .....	10
SAN Backups .....	11
NFS Backups .....	11
Using VMware Consolidated Backup .....	12
VMware Consolidated Backup Advantages .....	12

Chapter 2 VMware Consolidated Backup .....	13
Overview .....	14
How Consolidated Backup Works .....	14
Integration with Third-Party Backup Software .....	14
Requirements .....	15
VMware Virtual Machine Snapshot Technology .....	16
Considerations When Creating Snapshots .....	17
File-Level Backups .....	17
Full Virtual Machine Backups .....	18
Setting Up Consolidated Backup .....	18
VMware ESX Server and VirtualCenter Configuration .....	18
SAN Configuration .....	19
Configuring the Third-Party Software .....	20
Changing Backup Policies after ESX Server Upgrades .....	20
VCB Proxy Configuration .....	21
Hardware Requirements .....	21
Prerequisites .....	21
Configuring Windows on the VCB Proxy .....	22
Disabling Automatic Drive-Letter Assignment .....	22
Configuring the Consolidated Backup Framework .....	23
Installing a Backup Software Interoperability Module .....	25
Configuring Virtual Machines for Consolidated Backup .....	25
Consolidated Backup Restrictions .....	25
Using Consolidated Backup .....	26
Aliases .....	26
Configuring Backup Jobs .....	27
First-Time Backup .....	27
Advanced Configurations .....	27
Running Custom Quiescing Scripts .....	28
Canceling a Backup Job .....	29
Chapter 3 Restoration and Disaster Recovery .....	31
Restoring Your Files Using Consolidated Backup .....	32
Restore .....	32
Centralized Restore .....	32
Per-Group Restore .....	32
Self-Service Backup .....	33
Restoring Files Using the vcbRestore Utility .....	33

Data Recovery .....	33
Troubleshooting .....	33
<b>Chapter 4 Backup Scenarios .....</b>	<b>35</b>
Backup Usage Scenarios .....	36
A Typical Consolidated Backup Usage Scenario .....	36
<b>Appendix A Using Service Console to Back Up and Restore</b>	
<b>Virtual Machines .....</b>	<b>37</b>
General Configuration Settings for Consolidated Backup Utilities .....	38
Configuration File Settings .....	38
Backing Up Virtual Machines .....	40
Performing Backups .....	40
Identifying Virtual Machines .....	41
Identifying Virtual Machines by DNS Name or IP Address .....	41
Identifying Virtual Machines by BIOS UUID .....	41
Identifying Virtual Machines by MoRef .....	42
Identifying Groups of Virtual Machines .....	42
Specifying Backup Destinations .....	43
Backing Up to a Local Directory .....	43
Backing Up to a Remote Server .....	43
Restoring Virtual Machines .....	44
Restoring Virtual Machines to Original Locations .....	44
Restoring Virtual Machines to Alternative Locations .....	44
Copying a Catalog File .....	45
Editing a Catalog File .....	45
Restoring Virtual Machines Using an Alternate Catalog .....	47
Non-interactive Use of the vcbRestore Utility .....	48
<b>Appendix B Restoring Virtual Machines from</b>	
<b>ESX Server 2.5.x to ESX Server 3.0 .....</b>	<b>49</b>
Setting Configuration Parameters .....	50
Restoring ESX 2.5.x Server Virtual Machines .....	51
<b>Index .....</b>	<b>53</b>



# Preface

This preface describes the contents of the *Virtual Machine Backup Guide* and provides pointers to technical and educational resources.

This preface contains the following topics:

- [“About This Book”](#) on page viii
- [“Intended Audience”](#) on page viii
- [“Document Feedback”](#) on page viii
- [“VMware Infrastructure Documentation”](#) on page viii
- [“Conventions and Abbreviations”](#) on page ix
- [“Technical Support and Education Resources”](#) on page ix

## About This Book

This manual, the *Virtual Machine Backup Guide*, provides information on different methods you can use to perform backup and restore tasks. It also describes how to use VMware® Consolidated Backup, a new backup solution offered by ESX Server 3 and recommended to perform daily backups for virtual machines residing on a SAN.

## Intended Audience

The information presented in this manual is written for system administrators who are experienced Windows or Linux system administrators and who are familiar with virtual machine technology and datacenter operations.

## Document Feedback

If you have comments about this documentation, submit your feedback to:

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

## VMware Infrastructure Documentation

The VMware Infrastructure documentation consists of the combined VirtualCenter and ESX Server documentation set.

You can access the books in the VMware Infrastructure document set at:

<http://www.vmware.com/support/pubs>

## Conventions and Abbreviations

This manual uses the style conventions listed in [Table P-1](#).

**Table P-1.** Type Conventions

Style	Purpose
Monospace	Used for commands, filenames, directories, paths.
<b>Monospace bold</b>	Apply to indicate user input.
<b>Bold</b>	Use for these terms: Interface objects, keys, buttons Items of highlighted interest Glossary terms
<i>Italic</i>	Used for book titles.
< name >	Angle brackets indicate variable and parameter names.

## Technical Support and Education Resources

The following sections describe the technical support resources available to you:

- [“Self-Service Support”](#)
- [“Online and Telephone Support”](#)
- [“Support Offerings”](#)
- [“VMware Education Services”](#)

### Self-Service Support

Use the VMware Technology Network for self-help tools and technical information:

- Product Information – <http://www.vmware.com/products/>
- Technology Information – <http://www.vmware.com/vcommunity/technology>
- Documentation – <http://www.vmware.com/support/pubs>
- Knowledge Base – <http://www.vmware.com/support/kb>
- Discussion Forums – <http://www.vmware.com/community>
- User Groups – <http://www.vmware.com/vcommunity/usergroups.html>

For more information about the VMware Technology Network, go to <http://www.vmtn.net>.

## Online and Telephone Support

Use online support to submit technical support requests, view your product and contract information, and register your products. Go to <http://www.vmware.com/support>.

For customers with appropriate support contracts, use telephone support for the fastest response on priority 1 issues. Go to [http://www.vmware.com/support/phone\\_support.html](http://www.vmware.com/support/phone_support.html).

## Support Offerings

Find out how VMware's support offerings can help you meet your business needs. Go to <http://www.vmware.com/support/services>.

## VMware Education Services

VMware courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. For more information about VMware Education Services, go to <http://mylearn1.vmware.com/mgrreg/index.cfm>.

# CHAPTER 1 Introduction

Backup, restoration, and disaster recovery are among the most critical processes of datacenter management. VMware® ESX Server and VMware Infrastructure provide many different solutions, each suitable for a specific environment, to perform backup and restore tasks.

This introduction describes which resources should be backed up on an ESX Server system and explains options available for that backup.

This chapter includes the following sections:

- [“Concepts”](#) on page 2
- [“What to Back Up”](#) on page 2
- [“Backup Components and Approaches”](#) on page 5
- [“Using Traditional Backup Methods”](#) on page 6
- [“Using VMware Consolidated Backup”](#) on page 12

## Concepts

The following concepts are essential for your understanding of backup procedures:

- **Quiescing** – A process of bringing the on-disk data of a physical or virtual computer into a state suitable for backups. This process might include such operations such as flushing dirty buffers from the operating system's in-memory cache to disk, or other higher-level application-specific tasks.
- **Volume-level (image-level) backup** – A type of backup that backs up an entire storage volume.
- **File-level backup** – A type of backup that is defined at the level of files and folders.
- **Full backup** – Backs up all selected files.
- **Full virtual machine backup** – Backs up all files that comprise entire virtual machine. These files include disk images, .vmtx files, and so on.
- **Differential backup** – Backs up only those files that have changed since the last full backup.
- **Incremental backup** – Backs up only those files that have changed since the last backup, whether it is a full or incremental backup.
- **VCB proxy** – A physical machine running Microsoft Windows 2003, Consolidated Backup, and third-party backup software. Used to perform off-loaded file-level and full virtual machine backups.

## What to Back Up

Within the ESX Server environment, you need to back up the following major items regularly:

- **Virtual machine contents** – Because virtual machines are used frequently, critical information stored in their disk files constantly changes. As with physical machines, virtual machine data needs to be backed up periodically to prevent its corruption and loss due to human or technical errors.

The virtual machine data you back up includes virtual disks, Raw Device Mappings (RDM), configuration files, and so on. For more information, see [“Virtual Machine Contents”](#) on page 3.

- **ESX Server service console** – The service console, a customized version of Linux, is the ESX Server command-line management interface. It provides the ESX Server management tools and a command prompt for more direct management of ESX Server. It also keeps track of all the virtual machines on the server and their configurations.

---

**NOTE** In earlier releases, the service console was the main interface to the ESX Server host. With ESX Server 3 and later, the VI Client has priority, although you still might use the service console to perform some advanced administration operations.

---

During its lifetime, the service console doesn't experience any major changes other than periodic upgrades. In addition, in case of a failure, you can easily recover the state of your service console by reinstalling ESX Server. Therefore, although you might consider backing up the service console, it doesn't need to be backed up as frequently as the virtual machines and their data.

## Virtual Machine Contents

ESX Server uses **VMware File System (VMFS)** for its storage needs. VMFS is a simple, high-performance file system on physical SCSI disks and partitions. VMFS is used for storing large files, such as the virtual disk images for ESX Server virtual machines and the memory images of suspended virtual machines. ESX Server 3 supports VMFS-3. A VMFS-3 volume can span multiple disk partitions, also called extents.

### Virtual Machine Disks

In ESX Server 3, VMFS supports directories. Typically, there is one directory for each virtual machine on VMFS. This directory contains all the files that comprise the virtual machine, such as disk images, virtual machine configuration `.vmx` files, log files, and so on. The disk files are in a special format and generally use the `.vmdk` file extension.

---

**NOTE** All the information normally backed up in the enterprise infrastructure, including the operating system, applications, and data, is included in the virtual disks.

---

## Raw Device Mappings

Raw Device Mappings (RDM) lets you use all the features of VMware Virtual Infrastructure in conjunction with raw SAN LUNs. The mapping file is the file that is used to connect the raw LUN to the virtual machine and is referenced in the virtual machine's configuration.

Two modes exist for RDM: virtual compatibility and physical compatibility.

### Virtual Compatibility Mode

Virtual compatibility mode allows a mapping to act exactly as a virtual disk file does, including virtual machine snapshots.

In virtual compatibility mode, 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, and can have redo logs. To the virtual machine, the ESX Server presents the mapped device as a locally attached SCSI device.

### Physical Compatibility Mode

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 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 for the virtual machine that owns it. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical mode is useful when you need to run SAN management agents or other SCSI target-based software in the virtual machine. Physical mode is also used for virtual-to-physical clustering for cost-effective high availability.

## Accessing and Managing Virtual Disk Files

Typically, you use **Virtual Infrastructure (VI) Client** to perform a variety of operations on your virtual machines.

Direct manipulation of your virtual disk files on VMFS is possible through ESX Server service console and VMware SDKs, although using the VI Client is the preferred method.

From the service console, you can view and manipulate files in the `/vmfs/volumes` directory in mounted VMFS volumes with ordinary file commands, such as `ls` and `cp`. Although mounted VMFS volumes might appear similar to any other file system, such

as ext3, VMFS is primarily intended to store large files, such as disk images with the size of up to 2 TB. You can use ftp, scp, and cp commands for copying files to and from a VMFS volume as long as the host file system supports these large files.

---

**NOTE** In Linux, importing a large disk from a Common Internet File System (CIFS) mount hangs the ESX Server, which must then be rebooted. If Consolidated Backup is not used, large files need to be moved from the service console to tape backup. To do this effectively, use supported programs with no reported known issues. As a workaround, use `smbclient` to copy the large file onto a local directory on the service console and then import from there.

---

Additional file operations are enabled through the `vmkfstools` command. This command supports the creation of a VMFS on a SCSI disk and is used for the following:

- Creating, extending, and deleting disk images
- Importing, exporting, and renaming disk images
- Setting and querying properties of disk images
- Creating and extending a VMFS file system

For more information on the `vmkfstools` command, see the *Server Configuration Guide*.

## Backup Components and Approaches

When you perform a backup, the following three components of backup software are generally involved in the process:

- **Backup Client (Backup Agent)** – A program that scans virtual machine file systems and transfers data to be backed up to a backup server. During restore operations, the backup client writes the data into the file systems.
- **Backup Server** – A program that writes the data, pushed by the backup client, to a backup medium, such as a robotic tape library. During restore operation, the backup server reads the data from the backup medium and pushes it to the backup client.
- **Scheduler** – A program that allows you to schedule regular automatic backup jobs and coordinate their execution. Backups can be scheduled at periodic intervals, or individual files can be automatically backed up immediately after they have been updated.

## Backup Approaches

Each of the backup software components can be run in a virtual machine, on the service console, or on a physical machine, or a VCB proxy, running Microsoft Windows 2003. While the location of the scheduler isn't important, the locations of the backup server and backup client are important.

Depending on where you want to run each component, you choose one of the following approaches:

- Use traditional backup approach. With this approach, you deploy a backup client to every system that requires backup services. You can then regularly perform backups in an automated way.

With this approach, several methodologies exist. You can choose a specific method that better suites your needs and requirements.

For more information, see [“Using Traditional Backup Methods”](#) on page 6.

- Use VMware Consolidated Backup, which enables offloaded and impact-free backup for virtual machines running on ESX Server. This approach lets you use the virtual machine snapshot technology and SAN-based data transfer in conjunction with traditional file-based backup software.

When running Consolidated Backup, you can back up virtual machine contents from a centralized Microsoft Windows 2003 proxy server rather than directly from the ESX Server system. Utilizing a proxy server reduces the load on ESX Server allowing it to run more virtual machines.

For more information on Consolidated Backup, see [“VMware Consolidated Backup”](#) on page 13.

For more information on the snapshot technology that Consolidated Backup employs, see [“VMware Virtual Machine Snapshot Technology”](#) on page 16.

## Using Traditional Backup Methods

With the traditional backup methods, you deploy a backup agent on each host whose data needs to be secured. Backups are then conducted on a regular basis in an automated way.

The backup agent scans the file system for changes during periods of low system utilization and sends the changed information across the network to a backup server that writes the data to a backup medium, such as a robotic tape library.

Using traditional methods, you can back up your service console and virtual machines. For more information, see:

- [“Backing Up the Service Console”](#) on page 7
- [“Backing Up Virtual Machines”](#) on page 8

## Traditional Backup Considerations

When using traditional methods to back up your system, keep in mind the following:

- To be able to capture the data in its consistent state, perform backups at the times of the lowest activity on the network and when your computer resources are mostly idle. While performing backups, you might need to take critical applications off line.
- Make sure that network bandwidth between the server you are backing up and the backup server is sufficient.
- With a large number of servers, both physical and virtual, allocate enough resources to manage backup software on each host. Remember that managing agents in every virtual machine is very time consuming.

## Backing Up the Service Console

Because the service console doesn't experience any major changes during its lifetime and its state is easily recoverable in case of a failure, you might decide against backing it up. However, if you choose to back up the service console, you don't need to do it frequently.

Use the following methods when backing up service console:

- **File-Based** – The service console can be treated as a physical machine with a deployed backup agent. To restore the service console, first reinstall it, then reinstall the agent, and then restore the files that you backed up. This approach makes sense if management agents that are hard to set up have been deployed in the service console. Otherwise, this approach provides no advantage over not backing up the service console.
- **Image-Based** – Use third-party software to create a backup image that you can restore quickly. Use your boot CD or whatever the backup software created to restore the service console.

## Backing Up Virtual Machines

Depending on your requirements, you might choose one of the traditional methods for backing up your virtual machines. Traditional backup methods do not use Consolidated Backup.

Use the following table to compare available traditional methods.

**Table 1-1.** Recommended Traditional Backup Methods (No Consolidated Backup)

		Backup Server	
		Virtual Machine	Physical Machine
Backup Client	Virtual Machine	Method 1	Method 2
	Service Console	Method 3	Method 4

---

**NOTE** Running the backup server in the service console is not supported.

---

Traditional backup methods offer the following options:

- You can run backup clients from within a virtual machine performing file-level or system-level backups. As long as you are backing up over the network, no compatibility guide is needed.
- You can run backup clients from the ESX Server Service Console, backing up virtual machines in their entirety as `.dsk` and/or `.vmdk` files residing in the ESX Server host VMFS file system.
- You can back up virtual machine data by running a backup server within a virtual machine that is connected to a tape drive or other SCSI-based backup media attached to the physical system.

### Backup Client in a Virtual Machine

Method 1 and method 2 assume that you deploy your backup client in a virtual machine.

#### Method 1: Backup Server in a Virtual Machine

With this method, you deploy your backup client in one virtual machine while the backup server is in another virtual machine. Both virtual machines run on the same ESX Server system. Data between the two virtual machines moves through the virtual ethernet that connects these virtual machines.

---

**NOTE** Method 1 is not recommended in ESX Server except in a branch office scenario where no separate hardware for a VCB proxy or backup server is available.

---

When you use method 1, the backup agent performs quiescing of a virtual machine being backed up.

Method 1 is generally used for file-level backups of the data stored within the virtual machine's disk image.

#### Backup Server in a Virtual Machine

<b>Recommended:</b>	No
<b>File-level restore:</b>	Very easy
<b>Full virtual machine restore:</b>	No
<b>Quiescing:</b>	Excellent
<b>Load on ESX Server:</b>	Extremely high
<b>LAN-free backup:</b>	No
<b>Backup Impact:</b>	No
<b>Manageability:</b>	Very poor

#### Method 2: Backup Server in a Physical Machine

With this method, you deploy the backup client in a virtual machine while the backup server runs on a physical machine.

---

**NOTE** Instead of method 2, consider using Consolidated Backup.

---

Method 2 is used for file-level backups of the data stored within the virtual machine's disk image.

#### Backup Server in a Physical Machine

<b>Recommended:</b>	Yes
<b>File-level restore:</b>	Very Easy
<b>Full virtual machine restore:</b>	No
<b>Quiescing:</b>	Excellent
<b>Load on ESX Server:</b>	High
<b>LAN-free backup:</b>	No
<b>Backup Impact:</b>	No
<b>Manageability:</b>	Very poor

## Backup Client in the Service Console

Method 3 and method 4 assume that you deeply your backup client in the service console.

### Method 3: Backup Server in a Virtual Machine

With this method, you deploy the backup client in the service console while the backup server runs in the virtual machine.

---

**NOTE** Method 3 is not recommended in ESX Server except in a branch office scenario where no separate hardware for a VCB proxy or backup server is available.

---

Method 3 is used to perform image-level backups, or backups of entire virtual machines.

#### Backup Server in a Virtual Machine

<b>Recommended:</b>	No
<b>File-level restore:</b>	No
<b>Full virtual machine restore:</b>	Very easy
<b>Quiescing:</b>	Excellent
<b>Load on ESX Server:</b>	Extremely high
<b>LAN-free backup:</b>	No
<b>Backup Impact:</b>	No
<b>Manageability:</b>	Very poor

### Method 4: Backup Server in a Physical Machine

With this method, you deploy the backup client in the service console while the backup server runs on a physical machine.

---

**NOTE** Instead of method 4 consider using Consolidated Backup.

---

Method 4 is used to perform image-level backups.

#### Backup Server in a Physical Machine

<b>Recommended:</b>	Yes
<b>File-level restore:</b>	No
<b>Full virtual machine restore:</b>	Very Easy
<b>Quiescing:</b>	Excellent
<b>Load on ESX Server:</b>	High
<b>LAN-free backup:</b>	No
<b>Backup Impact:</b>	No
<b>Manageability/Scalability:</b>	Very poor

### SAN Backups

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.

If you decide to use SAN snapshots to back up your data, you must consider the following points:

- Some vendors support snapshots for both VMFS and RDMs. If both are supported, you can make either a snapshot of the whole virtual machine filesystem for a host, or snapshots for the individual virtual machines (one per disk).
- Some vendors support snapshots only for a setup using RDM. If only RDM is supported, you can make snapshots of individual virtual machines.

See your storage vendor's documentation for additional information. For more information on SAN, see the *SAN Configuration Guide*.

### NFS Backups

If your virtual machines are stored on external network-attached storage (NAS) systems using the network file system (NFS) protocol, you can perform image-level backups of the virtual machines.

See your storage vendor's documentation for additional information.

## Using VMware Consolidated Backup

In a modern datacenter environment, it has become increasingly difficult to apply the traditional approach to your backup processes. Using it might cause considerable overhead and a number of problems, some of which are described in [“Traditional Backup Considerations”](#) on page 7. To avoid many of the problems and issues, consider using VMware Consolidated Backup.

### VMware Consolidated Backup Advantages

VMware Consolidated Backup addresses most of the problems you encounter when performing traditional backups. Consolidated Backup helps you to:

- Reduce the load on your ESX Server systems by moving the backup tasks to one or more dedicated backup proxies.
- Avoid congesting and overloading the data center network infrastructure by enabling LAN-free backup.
- Eliminate the need for a backup window by moving to a snapshot-based backup approach.
- Simplify backup administration by making optional the deployment of backup agents in each virtual machine you back up.
- Back up virtual machines that are powered off.

## CHAPTER 2 **VMware Consolidated Backup**

VMware Consolidated Backup is a new backup solution offered by ESX Server 3 and is the recommended way to perform daily backups for virtual machines residing on a SAN. This method performs backups using a dedicated physical host (proxy client) rather than backing up from an ESX Server system itself.

You can use Consolidated Backup with a single ESX Server host or with a VirtualCenter Management Server.

This chapter includes the following information:

- [“Overview”](#) on page 14
- [“How Consolidated Backup Works”](#) on page 14
- [“Setting Up Consolidated Backup”](#) on page 18
- [“Configuring Virtual Machines for Consolidated Backup”](#) on page 25
- [“Consolidated Backup Restrictions”](#) on page 25
- [“Using Consolidated Backup”](#) on page 26
- [“Advanced Configurations”](#) on page 27

## Overview

VMware Consolidated Backup is a fast and efficient way of backing up data in virtual machines.

Consolidated Backup has the following characteristics:

- Offloads backup processes to a dedicated physical host
- Doesn't require backup agents in virtual machines
- Works with other industry-leading backup applications
- Doesn't restrict the use of Fibre Channel tapes
- Supports file-level backups for Windows virtual machines and full (image-level) backup for all guest operating systems

## How Consolidated Backup Works

Consolidated Backup enables offloaded and impact-free backup for virtual machines running on an ESX Server system by allowing traditional file-based backup software to leverage VMware virtual machine snapshot technology and efficient SAN-based data transfer.

### Integration with Third-Party Backup Software

In a typical Consolidated Backup scenario, backup scheduling is done by the third-party software. As a result, all the advanced scheduling and backup management functionality provided by the backup software is also available with Consolidated Backup.

Backup jobs are configured and launched through the third-party backup software. The basic backup workflow involves the following steps:

- 1 The third-party backup software schedules the backup job for a single virtual machine or a group of virtual machines. The job runs on the VCB proxy.
- 2 Consolidated Backup's pre-backup script runs on the VCB proxy, creating virtual machine snapshots (for Windows only) and mounting them on the proxy.
- 3 The third-party backup software performs the backup.
- 4 Consolidated Backup's post-backup script runs, removing the mount from the VCB proxy and removing the backup snapshots.

Even though the mounted disks appear to be writable, changes are cached on the proxy as “transient writes” and are discarded once the disk is unmounted. This feature lets you back up NT4 volumes that use an older version of NTFS. Windows 2003 running on the VCB proxy will fail any mount attempts for NT4 volumes if they are read-only.

For more information, see [“Advanced Configurations”](#) on page 27.

## Requirements

To be able to run Consolidated Backup, make sure that the following requirements are met:

- Microsoft Windows 2003 server is configured for Consolidated Backup. See [“Configuring the Consolidated Backup Framework”](#) on page 23.
- Windows virtual machines are in place. Required only for file-level backup. Full virtual machine backup (image-level) is supported for all guest operating systems.

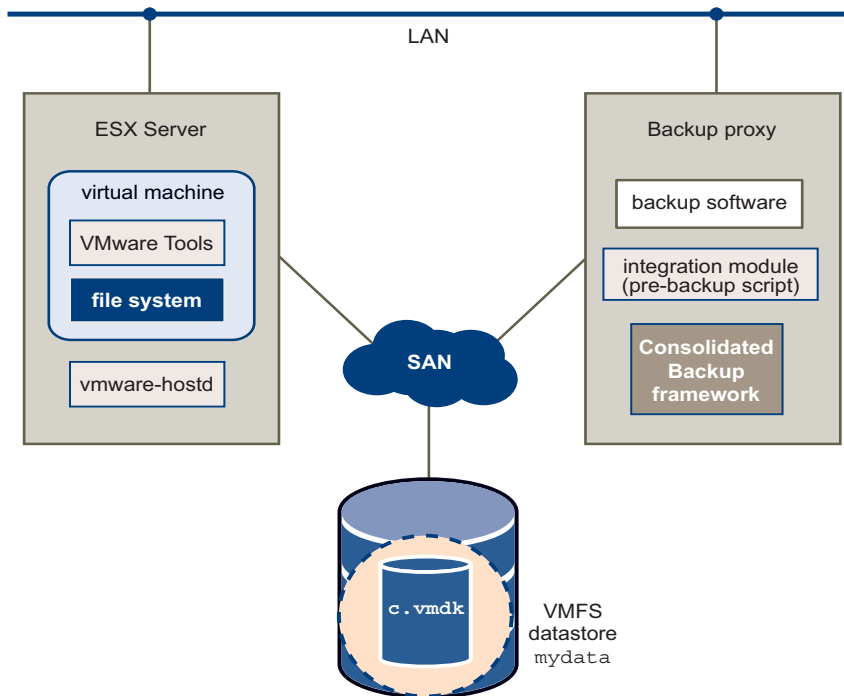
---

**CAUTION** File system consistency is guaranteed for Windows guest operating systems only.

---

- SAN is connected and LUN is masked.
- Backup administrator account exists on the VCB proxy with read access to the data. Because LUNs containing VMFS volumes are accessible by the proxy, to protect the VMFS volumes from accidental deletion, restrict access to the proxy to trained personnel only.

The following illustrates how different components of Consolidated Backup work together.



**Figure 2-1.** Consolidated Backup Components

## VMware Virtual Machine Snapshot Technology

ESX Server can create file-system-consistent snapshots of Windows virtual machines that have VMware Tools installed. No file system writes are pending at the time the snapshot of the virtual machine is taken. As a result, the snapshot represents a clean state of the virtual machine's disk images. Furthermore, system administrators can configure pre-freeze and post-thaw scripts in both Windows (application consistency) and Linux virtual machines to achieve even higher levels of data consistency.

After a snapshot of a virtual machine has been created, the frozen disk images can be accessed in different ways, as described in the remainder of this section.

## Considerations When Creating Snapshots

When creating snapshots, keep in mind the following:

- Taking a snapshot of a virtual machine freezes it for a very short period of time, usually considerably less than a second.
- Creating a memory snapshot through the UI can leave a virtual machine frozen for a couple of seconds. This might affect some very time-sensitive applications, such as DBHammer.
- Creating quiesced snapshots using the SYNC driver requires waiting for I/O to drain in the guest operating system. This might affect some very time-sensitive applications, such as DBHammer.
- If you choose not to install the SYNC driver when installing VMware Tools, you can avoid the delay caused by the I/O draining. However, it results in snapshots being only crash-consistent, unless you provide custom quiescing through pre/and post-scripts in the guest.

## File-Level Backups

For virtual machines running Windows, Consolidated Backup supports file-level backups. A file-level backup is ideal for preventing data loss due to errors, such as file corruption, caused by bugs in application programs, or user errors, such as accidental file deletion.

When running file-level backups, Consolidated Backup performs the following:

- 1 Analyzes volumes on virtual machine snapshots.
- 2 Mounts discovered volumes on the VCB proxy at predefined junction points.

Each junction point corresponds to a drive letter assigned to each partition in the virtual machine. For example:

```
C:\mnt\mytestvm.foo.com\letters\D
```

---

**CAUTION** Because the proxy can recognize only those partitions that have drive letters assigned to them, make sure that each virtual disk partition has a drive letter.

---

## Full Virtual Machine Backups

Consolidated Backup supports full virtual machine backups. A full virtual machine backup is operating-system neutral and can be performed regardless of the guest operating system.

This type of backup is suitable for restoring entire virtual machines in the event of a hardware failure or a system administrator error, such as the accidental deletion of an entire virtual machine or a disk image.

When running full virtual machine backups, the following steps are performed:

- 1 Consolidated Backup exports virtual machine disk images and configuration files to a local directory on the VCB proxy. For example:  
  
`C:\mnt\mytestvm.foo.com-fullVM`
- 2 Third-party backup software picks up the virtual machine files and moves them to the backup medium.

## Setting Up Consolidated Backup

Before you use Consolidated Backup, configure the following components:

- VMware ESX Server and VirtualCenter. See [“VMware ESX Server and VirtualCenter Configuration”](#) on page 18.
- The SAN fabric. See [“SAN Configuration”](#) on page 19.
- The third-party software. See [“Configuring the Third-Party Software”](#) on page 20.
- The backup proxy (VCB proxy). See [“VCB Proxy Configuration”](#) on page 21.
- Windows on the VCB proxy. See [“Configuring Windows on the VCB Proxy”](#) on page 22.

## VMware ESX Server and VirtualCenter Configuration

You should have an existing installation of an ESX Server and the Virtual Infrastructure (VI) Client or a multi-host configuration and VirtualCenter to manage them.

## SAN Configuration

Before installing Consolidated Backup, set up your ESX Servers to use VMware File System (VMFS) or virtual compatibility raw device mappings (RDM). You then need to configure the SAN fabric to which both the ESX Servers and the VCB proxy are connected.

For more information, see documentation provided by your SAN storage array and switch vendors, as well as the *SAN Configuration Guide*.

For Consolidated Backup, your SAN configuration must meet the following requirements:

- The VCB proxy must have access to:
  - All the SAN arrays containing VMFS volumes (datastores) with virtual disks on them
  - All the SAN arrays containing virtual compatibility RDMs that are supposed to be backed up using Consolidated Backup
- For every LUN containing VMFS or RDM data, the LUN ID on the proxy server must match the LUN ID as seen by the ESX Server.

---

**NOTE** There is no support for EMC AX100 managed with NaviExpress. Use full NaviSphere.

IBM ESS should be explicitly configured to have the consistency of LUN ID presentation.

---

Usually, the SAN configuration on the storage array involves adding the VCB proxy to all the host groups where the ESX Servers are located.

## Configuring the Third-Party Software

You must configure third-party backup software for use with Consolidated Backup. This involves enabling the use of the Consolidated Backup pre-scripts and post-scripts for your backup software. Also, you might have to turn on the cross junctions (mount points) option for your backup software.

When configuring third-party backup software, make sure that the backup client running on the VCB proxy doesn't use the following mechanisms. Both mechanisms require the backup software to alter the file system being backed up, which is not possible in a snapshot-based backup.

- Windows archive bit for incremental backups. The archive bit is an attribute of a file that some backup products use to determine whether the file should be backed up or not. Configure the backup client to use time stamps instead.
- Windows change journal.

For each supported third-party backup software, either the backup software vendor or VMware provides an Integration Module. This is a ZIP file containing all the required pre-backup and post-backup scripts.

The ZIP file contains a README.html file that describes how to use the Integration Module on the respective third-party backup software. Also, the README.html file outlines how to turn off the Windows archive bit and change journal functions in your third party backup software, if required.

### Changing Backup Policies after ESX Server Upgrades

If you have upgraded your ESX Server software 2.x to 3.0, you need to modify all VMFS volume paths that you configured in the backup software. A path format that ESX Server 3 uses is different from the ESX Server 2.x format and follows this standard:

- VMFS volume  
/vmfs/volumes/<file\_system\_UUID> or /vmfs/volumes/<file\_system\_label>
- VMFS file  
/vmfs/volumes/<file system label|file system UUID>/[dir]/myDisk.vmdk

## VCB Proxy Configuration

To use Consolidated Backup, you need to configure a backup proxy.

### Hardware Requirements

The VCB proxy must be able to run Microsoft Windows 2003. In addition, the proxy requires the following hardware components:

- Network adapter (NIC)
- Fibre Channel host bus adapter (HBA)

### Prerequisites

To be able to install Consolidated Backup on the VCB proxy, make sure that the following requirements are met:

- The proxy is running Microsoft Windows 2003. Consolidated Backup doesn't support any other versions of Windows on the proxy.
- Networking on the backup proxy is configured so that the proxy can establish a connection to VirtualCenter or its ESX Server peer.

If there is a firewall between the backup proxy and the VirtualCenter or ESX Server peer, the firewall must permit TCP/IP connections to VirtualCenter or the ESX Server peer. By default, VirtualCenter expects incoming connections at TCP/IP port 902.

- The third-party backup software to be used with Consolidated Backup is installed and correctly configured.

Verify the configuration of the third-party backup software at this time by running a backup and restoration job on a local directory on the VCB proxy.

## Configuring Windows on the VCB Proxy

After setting up the VCB proxy, you need to configure Windows that runs on the proxy. Configuring Windows involves the following:

- [“Disabling Automatic Drive-Letter Assignment”](#) on page 22
- [“Configuring the Consolidated Backup Framework”](#) on page 23
- [“Installing a Backup Software Interoperability Module”](#) on page 25

### Disabling Automatic Drive-Letter Assignment

All versions of Windows, except Windows 2003 Enterprise Edition and Windows 2003 Datacenter Edition, automatically assign drive letters to each visible new technology file system (NTFS) and file allocation table (FAT) volume.

For Consolidated Backup, you need to change this default behavior so that volumes are not automatically mounted on the proxy.

---

**CAUTION** If you do not perform this configuration step, data corruption for virtual machines using RDM can occur.

---

### To prevent Windows from automatically assigning drive letters to RDM

- 1 Shut down the Windows proxy.
- 2 Disconnect the Windows proxy from the SAN or mask all the LUNs containing VMFS volumes or RDM for virtual machines.
- 3 Boot the proxy and log into an account with administrator privileges.
- 4 Open a command-line interface.
- 5 Run the diskpart utility by typing:
 

```
diskpart
```

The diskpart utility starts up and prints its own command prompt.
- 6 Disable automatic drive-letter assignment to newly seen volumes by typing at the diskpart command prompt:
 

```
automount disable
```
- 7 Clean out entries of previously mounted volumes in the registry by typing at the diskpart command prompt:
 

```
automount scrub
```

- 8 Exit the diskpart utility by typing:  
exit
- 9 Shut down Windows.
- 10 Reconnect the Windows proxy to the SAN, or unmask all previously masked LUNs containing either VMFS volumes or RDM.
- 11 Boot the proxy.

### **To install the basic Consolidated Backup framework**

- 1 Log into the backup proxy using an account with administrator privileges.
- 2 Install the Consolidated Backup framework by running `setup.exe` from your CD-ROM or electronic distribution.
- 3 During the installation, pick an installation directory for Consolidated Backup or accept the default one.

### **Configuring the Consolidated Backup Framework**

Essential configuration for Consolidated Backup is stored in a configuration file called `config.js`. It is located in a subdirectory named `config` within the install directory for Consolidated Backup.

The following table provides an overview of all the configuration settings in this file.

**Table 2-1.** Configuration Settings

Option	Default	Description
BACKUPROOT	C:\mnt	<p>Directory in which all the virtual machine backup jobs are supposed to reside.</p> <p>For each backup job, a directory with a unique name derived from the backup type and the virtual machine name is created here.</p> <p>Make sure this directory exists before attempting any virtual machine backups.</p> <p>For full virtual machine backups, the volume containing this mount point must be large enough to hold the exported disk images of the largest virtual machine to be handled.</p>
HOST	(no default)	Host name/port of the VirtualCenter server or the ESX Server peer used by the VCB proxy.
PORT	902	Port number to connect to on the VirtualCenter or ESX Server peer.
USERNAME	(no default)	User ID to be used for logging into the VirtualCenter host or ESX Server peer.
PASSWORD	(no default)	Password to be used for logging into the VirtualCenter host or ESX Server peer.
SNAPSHOT_POLICY	auto-matic	<p>Valid options:</p> <p><b>automatic:</b> The Consolidated Backup framework creates and deletes backup snapshots for virtual machines on demand. This is the default used most of the time.</p> <p><b>manual:</b> The Consolidated Backup framework does not create or delete any snapshots but assumes that a backup snapshot named <code>_VCB_BACKUP_</code> already exists and uses this snapshot for backup purposes. This option is useful for creative scripting.</p> <p><b>createonly:</b> The Consolidated Backup framework creates a backup snapshot when the pre-backup script is being run, but it does not remove the snapshot after backup. This option is used if you need to run a verification job. Your verification script would then be responsible for tearing down the mount.</p> <p><b>deleteonly:</b> The Consolidated Backup framework assumes that a backup snapshot named <code>_VCB_BACKUP_</code> already exists and does not attempt to create one. However, the snapshot is deleted by the post-backup script. This option is useful for creative scripting.</p>

## Installing a Backup Software Interoperability Module

Finally, you must install a Consolidated Backup Interoperability module that matches your third-party backup software.

For directions on how to install a particular interoperability module, refer to the documentation provided with the interoperability module.

## Configuring Virtual Machines for Consolidated Backup

In general, no particular configuration is required within the virtual machine to support Consolidated Backup.

However, you must install a new version of VMware Tools corresponding to ESX Server 3 in each protected virtual machine. Without installing VMware Tools, the snapshots that Consolidated Backup creates for backup will be crash-consistent only. That is, no file system synchronization will be performed.

## Consolidated Backup Restrictions

In general, Consolidated Backup can back up virtual machines on a backup proxy, thus offloading backup from the ESX Server and the protected virtual machine.

In some situations, Consolidated Backup cannot be used for backing up data in a virtual machine. If you encounter any of these situations, you should deploy a backup agent in the virtual machine and perform backup from within the virtual machine. The backup agent should be supplied by your third-party backup software vendor.

For details on installation of the backup agent, refer to the documentation provided with the interoperability module matching your backup software.

You cannot use Consolidated Backup to do any of the following:

- Back up virtual machines with disk images stored on a storage device that the proxy cannot access
- Back up virtual machines with virtual disks that are physical compatibility RDMs
- Back up virtual machines that do not have an Internet protocol (IP) address or a domain name server (DNS) name associated with them
- Perform a file-level backup of virtual machines running operating systems other than Windows NT 4.0, Windows 2000, Windows XP, Windows XP Professional, or Windows 2003
- Back up virtual machines that reside on NAS/NFS or iSCSI storage devices

- Perform a file-level backup simultaneously with a full backup for the same virtual machine

---

**CAUTION** Running two backup jobs on the same VCB proxy, one that performs a full backup for one virtual machine while another performs a file-level backup for another virtual machine, can trigger a failure of one of these operations. If this happens, restart the failed operation.

---

## Using Consolidated Backup

If you have multiple virtual machines to back up, group these virtual machines together and manage that group as a single entity in your backup software by configuring DNS aliases for the proxy.

---

**NOTE** Consolidated Backup supports a maximum of 60 concurrently mounted virtual machines. For example, you can concurrently mount 60 virtual machines that have a C: drive, or 30 virtual machines that have a C: and a D: each.

---

After you have associated a group of virtual machines with one host name, you can set up a backup job for each alias using the alias as the client name for the job.

## Aliases

You can set up different aliases for a group of virtual machines, all pointing to the same IP address of the VCB proxy. This lets you associate different permissions with the group. For example, restore permissions are tied to these aliases, determining who can restore and which virtual machines can be restored from.

If your department grows and you add another physical machine, aliases make it easy to move a group to a different proxy. You can add new proxies as the datacenter grows and then move jobs for the group by pointing the alias to a new proxy.

## Configuring Backup Jobs

Consolidated Backup is able to make backups of virtual machines on the backup proxy.

When configuring backup jobs, make sure that:

- All jobs are assigned to the proxy.
- All jobs are specified by one of the following directories:
  - file-level: C:\mnt\mytestvm.foo.com\letters\D
  - image-level: C:\mnt\mytestvm.foo.com-fullVM
- Each job is scheduled to run at specific time.

---

**NOTE** If you plan to run multiple backup jobs on the same proxy at the same time, remember that backup products might have limitations on a number of jobs you can run in parallel.

---

For more details on how to configure backup jobs for virtual machines, refer to the documentation provided with your third-party backup software.

## First-Time Backup

When you perform a first backup for a particular virtual machine, the virtual machine has to be powered on, or the backup fails on ESX Server.

After you have completed the first backup of the virtual machine, Consolidated Backup can perform backups of the virtual machine regardless of its power state at backup time.

## Advanced Configurations

On certain occasions, you might need to perform some advanced configurations when using Consolidated Backup. For example, you can run custom scripts to create a quiescent snapshot of your virtual machine. Also, you might need to run a post-backup command for your virtual machines to cancel backup jobs.

## Running Custom Quiescing Scripts

When you use Consolidated Backup, your virtual machines are automatically quiesced when you start the backup process.

You can also run custom pre-freeze and post-thaw quiescing scripts to create a quiescent snapshot of your virtual machine. You deploy and run the custom quiescing scripts inside the protected virtual machine.

When running the scripts, you can use the SYNC driver, an optional feature that holds incoming I/O and flushes all dirty data to a disk, thus making file systems consistent.

Because the SYNC driver is optional, you might decide not to install it when installing VMware Tools. The SYNC driver is currently not supported on the following operating systems:

- 64-bit guest operating systems
- Any operating system other than Windows

Running the scripts involves the following steps.

### Step 1 Running a pre-freeze script

Consolidated Backup runs the following pre-freeze script within the virtual machine being backed up:

- For Windows:  
`C:\Windows\pre-freeze-script.bat`
- For all other operating systems:  
`/usr/sbin/pre-freeze-script`

If the pre-freeze script returns a non-zero exit code, the snapshot create operation fails.

### Step 2 Engaging the SYNC driver (optional)

Engage the SYNC driver to hold incoming I/O and flush all dirty data to a disk. This helps to make the file systems consistent.

If this step fails, proceed to step 5 and fail the snapshot creation.

### Step 3 Creating a snapshot

In this step, an actual quiescent snapshot of your virtual machine is created. If this step takes too long and times out, the snapshot create operation fails and the snapshot is deleted. Proceed to step 5.

### Step 4 Disengaging the SYNC driver

Disengage the SYNC driver to allow I/O again. This step can fail if the snapshot creation in step 3 took too long and timed out.

### Step 5 Running post-thaw script

Consolidated Backup runs the following post-thaw script within the virtual machine:

- For Windows:  
C:\Windows\post-thaw-script.bat
- For all other operating systems:  
/usr/sbin/post-thaw-script

Even though VMware Tools does not currently check the exit status of this script, it should return 0 if successful.

## Canceling a Backup Job

If a backup operation is canceled from your third-party software while the backup is in process, the virtual machine might not be unmounted from the proxy server, and the snapshot might not be deleted. This is to be expected because the cleanup script was unable to run. To fix the problem, you must manually run the post-backup command for each virtual machine.

### To run the post-backup command

- 1 Check the folders in the C:\mnt directory to determine the virtual machine host names.
- 2 For each virtual machine host name, run the following command. Run this command from the generic subdirectory in the Consolidated Backup installation directory.

```
cscript /nologo <VCB default installation path ["C:\program
files\VMware\VMware Consolidated Backup Framework"]> post-
command.wsf <virtual_machine_hostname>
```



## CHAPTER 3    **Restoration and Disaster Recovery**

This chapter describes how to restore your data or recover from a disaster. You need to find a balance between the number of agents that you want to use and the ease with which you can restore your data.

This chapter includes the following information:

- [“Restoring Your Files Using Consolidated Backup”](#) on page 32
- [“Restoring Files Using the vcbRestore Utility”](#) on page 33
- [“Data Recovery”](#) on page 33

## Restoring Your Files Using Consolidated Backup

Consolidated Backup helps you perform file-based restores of your virtual machines.

### Restore

For information on restoring workflows, consult the documentation that comes with the integration module for your third-party backup software.

In general, three different restoration workflows are supported:

- No backup software in virtual machine. Restoration is done by the administrator on a backup proxy network share that is accessible by the protected virtual machine.
- Backup software in dedicated virtual machines and data moved to target virtual machines.
- Backup software deployed in every protected virtual machine. Restoration is done directly by the system administrator or the user.

### Centralized Restore

When performing a centralized restore, you have a group of virtual machines on ESX Server, a proxy, and a backup agent on the proxy in a dedicated virtual machine that you are planning to use to restore your data. In this case, you use the backup software to get the data to the proxy that is running the agent. After the administrator restores the data to the central server, you can copy it back to the virtual machine using the Common Internet File System (CIFS) remote-access, file-sharing protocol.

Pros: There are a minimum number of agents to maintain.

Cons: Data restoration is now centralized, and the administrator is involved in file-level restoration.

### Per-Group Restore

When performing a per-group restoration, one virtual machine has a backup agent for each group, such as accounting, engineering, and marketing. The group administrator restores workflows to a per-group restore host. Files are copied to a target virtual machine using CIFS file share.

Pros:

- Restorations can be delegated.
- This is a good compromise between the number of agents and ease of restoration.

Cons: This is not a complete self-service restoration.

## Self-Service Backup

Backup agents are deployed in every virtual machine. The user can use the agent to back up data to tape and restore the same way. The backup agent in the virtual machine is used to restore the data.

Pros: This is self-service restoration.

Cons: Agents are required in each virtual machine.

## Restoring Files Using the vcbRestore Utility

The vcbRestore utility is a command-line utility that you use to restore data that has been backed up using image-based backup.

For information on how to use this utility, see [“Using Service Console to Back Up and Restore Virtual Machines”](#) on page 37.

## Data Recovery

The following guidelines can aid you in recovering your data:

- Make sure you have full virtual machine image-based backups.
- Back up your VirtualCenter Database.
- Make sure you have your license keys.
- Make sure you have enough servers to run all the virtual machines you plan to restore.

Enabling migration with VMotion or using DRS greatly enhances your disaster recovery capabilities.

## Troubleshooting

In ESX 3.0, VMFS volumes are identified by unique identifiers. The name of the directory under which each VMFS volume is mounted in `/vmfs/volumes` corresponds to this unique identifier. The unique identifier is assigned to the volume automatically during formatting, and you cannot change it.

VMFS volumes can have user-friendly labels. These labels show up as symbolic links in `/vmfs/volumes` that point to the corresponding directory. For example, for a VMFS volume with the unique identifier `43a0552e-ae6093b2-47a1-00145e0a7ec0` and the label `storage1`, the following entries are created under `/vmfs/volumes`:

- A directory named `43a0552e-ae6093b2-47a1-00145e0a7ec0`, under which the file system is mounted
- A symbolic link named `storage1`, pointing to the directory `43a0552e-ae6093b2-47a1-00145e0a7ec0`

Graphical user interfaces (GUIs) for backup software that allow you to select files for backup show only the directory (the unique ID) in their Browse Directory pane.

### **To identify a VMFS volume by its label**

- 1 Browse the `/vmfs/volumes` directory in your backup software GUI.  
The symbolic links pointing to the VMFS volume mount points show up in the file selection pane.
- 2 Use these entries to find the unique ID for the file system label you need.
- 3 Select the directory corresponding to this unique ID in the directory pane for browsing.

When performing file-based backups, the backup application uses paths referencing the unique identifier, so the backed-up files show up as the following:

```
/vmfs/volumes/43a0552e-ae6093b2-47a1-00145e0a7ec0/vm01/vm01.vmdk
```

When restoring files from the backup application, you might need to perform a reverse mapping to identify the correct VMFS volume label (in this example, `storage1`) corresponding to this unique identifier. To be able to do this, consider backing up the symbolic link itself while performing backups.

## CHAPTER 4 **Backup Scenarios**

This chapter describes real-world scenarios that can help you plan your backup strategies.

This chapter includes the following information:

- [“Backup Usage Scenarios”](#) on page 36
- [“A Typical Consolidated Backup Usage Scenario”](#) on page 36

## Backup Usage Scenarios

The following are the most recommended use cases:

- **Datacenter** – Consolidated Backup at the file level performed every night.
- **Datacenter** – Consolidated Backup at the image level performed periodically for Windows or nightly for Linux. This is a disaster recovery scenario.
- **Agents in virtual machines** – Incremental backup for Linux.
- **Backup server in a virtual machine** – For branch offices, agents in the virtual machine as well.

## A Typical Consolidated Backup Usage Scenario

This is an example of how you can use Consolidated Backup to protect data in virtual machines:

- 1 The system administrator configures backup schedules and policies in the third-party backup software.  
  
For example, the system administrator might instruct the backup software to back up `D:\Data` on `vm37.company.com` daily at 3:05 a.m.
- 2 The backup software schedules this backup job automatically.
- 3 When the backup software launches this job, it calls into the Consolidated Backup framework by using a pre-backup script. Consolidated Backup performs the following:
  - a Contacts a VirtualCenter instance or an ESX Server peer, and requests it to create a snapshot of the virtual machine to be backed up.
  - b Makes this snapshot available (mounted) on the backup proxy. This makes the data that needs to be backed up visible to the third-party backup software.
- 4 The third-party backup software performs the backup procedure of copying changed data to the backup media.
- 5 At the end of the backup job, the third-party backup software calls into the Consolidated Backup framework, using a post-backup script in which Consolidated Backup does the following:
  - a Detaches (unmounts) the snapshot from the backup proxy.
  - b Asks VirtualCenter or its ESX Server peer to remove the virtual machine snapshot.

# APPENDIX A **Using Service Console to Back Up and Restore Virtual Machines**

This appendix describes how to back up and restore virtual machines using the service console. The appendix walks you through the process of configuring the Consolidated Backup command-line utilities and provides examples on how to use these utilities.

This appendix includes the following sections:

- [“General Configuration Settings for Consolidated Backup Utilities”](#) on page 38
- [“Backing Up Virtual Machines”](#) on page 40
- [“Restoring Virtual Machines”](#) on page 44

## General Configuration Settings for Consolidated Backup Utilities

Before using service console Consolidated Backup utilities, edit the `/etc/vmware/backuptools.conf` configuration file to set the most common parameters for these tools.

Because this configuration file is parsed as a Bourne shell script, you should follow general syntax conventions of the Bourne shell when editing the file:

- Use the `#` character to indicate a comment.
- Do not use spaces when entering variables. For example, `F00="bar"` should have no spaces around the equals sign.
- Use a backslash before entering any special characters, such as `$`. For example, `\$server`.

Administrators familiar with Bourne shell script programming can use all the standard Bourne shell mechanisms, such as command execution, for example ``foo``, or use environment variables.

### Configuration File Settings

Use the `/etc/vmware/backuptools.conf` configuration file to set up the following options.

#### VCHOST

Specifies the URL of the Virtual Center instance that manages the ESX Server host being backed up or restored. `VCHOST` should point to the Virtual Center instance managing the host.

If you perform the backup or restore operations on a standalone host, you can use `localhost` as the host name.

---

**NOTE** You can use the `-h` command-line option for any Consolidated Backup command-line utility to override this setting.

---

#### USERNAME

Specifies the user name to log into the VirtualCenter instance defined by `VCHOST`. The user must have privileges to be able to register or create virtual machines.

---

**NOTE** You can use the `-u` command-line option for any Consolidated Backup command-line utility to override this setting.

---

**PASSWORD**

Specifies the password corresponding to USERNAME. This option allows you to perform virtual machine backups in a non-interactive way.

---

**NOTE** Because specifying a password in a configuration file can present a security risk, make sure that the Service Console is not used by anyone except an ESX Server administrator.

---

---

**NOTE** You can use the `-p` command-line option for any Consolidated Backup command-line utility to override this setting.

---

**VMNAMECACHE**

The most common way of identifying a virtual machine for backup purposes is by its DNS name or by its IP address. However, when you back up a virtual machine from a standalone ESX Server host, the ESX Server host can recognize the IP address only when the virtual machine is powered on and running VMware Tools.

To be able to perform backups of the virtual machine on the standalone ESX Server host even when the virtual machine is powered off, you should maintain a cache file. The cache file records the IP address of the virtual machine each time the virtual machine is being backed up. This allows you to perform the future backups of this virtual machine regardless of its power state.

VMware recommends that you do not change the default setting.

---

**NOTE** You can use the `-c` command-line option for `vcbMounter` to override this setting. The `vcbRestore` command does not use this setting.

---

**TEMPDIR**

If you are using the secure copy capabilities of the Consolidated Backup command-line utilities, you can use this option to specify a temporary holding space for your virtual machine data.

This holding space must have enough free storage to hold the largest of your virtual machines.

---

**NOTE** This setting cannot be overridden from the command line.

---

## Backing Up Virtual Machines

You can use `vcbMounter` to back up an entire virtual machine in the service console. The `vcbMounter` utility creates a quiesced snapshot of the virtual machine and exports the snapshot into a set of files, which can be later used to restore the virtual machine. To back up the set of files, you can use any file-based third-party backup software.

Before backing up a virtual machine using `vcbMounter`, determine the following:

- Which virtual machine to back up.  
For information on identifying virtual machines, see [“Identifying Virtual Machines”](#) on page 41.
- Where to store the backup data.  
Consolidated Backup service console supports different transport plug-ins to either back up the virtual machine to a local directory or back it up to a remote directory using `scp`. For more information, see [“Specifying Backup Destinations”](#) on page 43.

## Performing Backups

After setting up configuration options as described in [“Configuration File Settings”](#) on page 38, enter the following command in the command line:

```
vcbMounter -a <virtual_machine_identifier> -r <backup_destination>,
```

where

- `<virtual_machine_identifier>` is a unique identifier of the virtual machine you’re backing up. For information on identifying virtual machines, see [“Identifying Virtual Machines”](#) on page 41.
- `<backup_destination>` specifies the location for backup data. For information on how to specify a backup destination, see [“Specifying Backup Destinations”](#) on page 43.

---

**NOTE** When backing up a group of virtual machines, use the `vcbSnapAll` command instead of `vcbMounter`. For information on how to identify the group you want to back up, see [“Identifying Groups of Virtual Machines”](#) on page 42.

---

Follow these examples when backing up virtual machines:

- Backing up the virtual machine `vm37.company.com` to the local directory `/home/VMs/vm37`:

```
vcbMounter -a ipaddr:vm37.company.com -r /home/VMs/vm37
```

- Backing up the virtual machine `vm37.company.com` to the directory `/backups/VMs/vm37`. The directory is located on the remote server `backups.company.com` with the user ID `vmware`. The `backups.company.com` host is running a secure shell (ssh) server. You can use Consolidated Backup's secure copy (scp) plug-in to transfer the virtual machine to `backups.company.com`.

```
vcbMounter -a ipaddr:vm37.company.com -r scp://vmware@backups.company.com:/backups/
VMs/vm37
```

- Backing up a virtual machine on a standalone ESX Server host. To identify the virtual machine, use the virtual machine's name displayed in the VI Client. The virtual machine is backed up to the local directory `/home/VMs/vm37`.

```
vcbMounter -a name:"Virtual Machine 37" -r /home/VMs/vm37
```

---

**NOTE** The virtual machine name argument is case sensitive.

---

## Identifying Virtual Machines

You can use different standards to specify the virtual machine you want to back up.

### Identifying Virtual Machines by DNS Name or IP Address

The most common way of identifying virtual machines is to use their DNS name or IP address. To identify the virtual machine, use the following specification:

```
ipaddr:<DNS name or IP address>
```

For example, to refer to the virtual machine `vm37.company.com` with the IP address `10.17.5.12`, use one of the following search specifiers:

- `ipaddr:vm37.company.com`
- `ipaddr:10.17.5.12`

### Identifying Virtual Machines by BIOS UUID

You can identify a virtual machine by its universally unique identifier (UUID). Use the following search specifier:

```
uuid:<uuid>
```

For example:

```
uuid:564d78a1-8c1c-59b4-fa02-be14138797be
```

## Identifying Virtual Machines by MoRef

Internally, VirtualCenter and ESX Server refer to objects by Managed Object References (MoRef). To identify a virtual machine by MoRef, follow these examples:

- `moref:vm-00027` – Use this format when accessing VirtualCenter Server.
- `moref:248` – Use this format when accessing the ESX Server host.

Because MoRefs change every time the VirtualCenter server or the host instance that Consolidated Backup connects to gets restarted, you should not use MoRefs to identify virtual machines. However, when running a shell script to back virtual machines, you can use MoRefs to identify the virtual machines.

For example, you can write a script that uses `vcbVmName` with the `any:` search specifier to get a list of all virtual machines first, and then performs custom filtering to produce a list of only those virtual machines you want to back up. Virtual machines on this list can use MoRefs as their identifiers. Another part of your script can then call `vcbMounter` on each of these MoRefs to perform the backup operations.

In a case like this, using MoRefs rather than other identifiers, such as UUID, causes less search overhead because the entire list of all virtual machines doesn't need to be parsed each time the identifier is used.

## Identifying Groups of Virtual Machines

When you need to back up a group of virtual machines, you use the `vcbSnapAll` command instead of `vcbMounter`. You identify a specific group by using one of the following search specifiers:

- `powerstate:on|off|suspended` – Finds all virtual machines with the specified power state.
- `any:` – Finds all virtual machines.
- `none:` – Doesn't find any virtual machines. You can use this option for testing purposes.

## Displaying Virtual Machine Information

To search for a particular virtual machine and get information about it, use `vcbVmName`.

Follow these examples:

- `vcbVmName -s powerstate:on` – Lists all powered-on virtual machines.
- `vcbVmName -s any:` – Lists all known virtual machines.
- `vcbVmName -s ipaddr:vm37.company.com` – Displays information about the virtual machine with the specified address.

The following is the sample output you get after using `vcbVmName`:

```
bash #vcbVmName -s name:vm37.company.com Found VM:
moref:192
name:Virtual Machine 37
uuid:564d78a1-8c1c-59b4-fa02-be14138797be
ipaddr:10.17.5.31
```

## Specifying Backup Destinations

You can back up a virtual machine to a local directory or to a remote server using `scp`.

### Backing Up to a Local Directory

When backing up a virtual machine to a local directory, you can specify the path to the directory or use the file transport plug-in descriptor.

For example, to back up a virtual machine to the local directory `/home/VMs/vm37`, you can use one of the following specifiers:

- `/home/VMs/vm37`
- `file://home/VMs/vm37`

---

**NOTE** You don't need to create the destination subdirectory, such as `/home/VMs/vm37`, in advance because the backup operation will create it. However, the directory that lists your destination subdirectory, for example `/home/VMs`, must exist before you start a backup process.

---

### Backing Up to a Remote Server

When backing up a virtual machine to a remote server, you can use a corresponding `scp` plug-in. Use the following syntax:

```
scp://<user>@<host>:<path>
```

To perform the `scp` backup in an automated way, use RSA key-based authentication. In this case, `scp` will not prompt you for a password during backup.

For example, you need to back up a virtual machine to the directory `/backup/VMs/vm37` located on the remote server `backups.company.com` that uses the `vmware` user ID. Enter the following:

```
scp://vmware@backups.company.com:/backups/VMs/vm37
```

---

**NOTE** Prior to backup, make sure that the `/backups/VMs` directory already exists on the remote server. However, you do not need to create the `/backups/VMs/vm37` directory because it will be created during the backup operation.

---

## Restoring Virtual Machines

You can restore a virtual machine to its original location or to another location of your choice.

### Restoring Virtual Machines to Original Locations

If you set up all configuration options as described in [“Configuration File Settings”](#) on page 38, the following is the only command you need to pass to `vcbRestore` to restore a virtual machine:

```
vcbRestore -s <backup_directory>
```

For information on how to specify a backup directory, see [“Specifying Backup Destinations”](#) on page 43.

Follow these examples when restoring your virtual machines:

- Restoring a virtual machine from a local backup directory named `/home/VMs/vm37`:

```
vcbRestore -s /home/VMs/vm37
```

- Restoring a virtual machine from the remote server `backup.company.com`, directory `/backups/VMs/vm37`, and user ID `vmware`:

```
vcbRestore -s scp://vmware@backup.company.com:/backups/VMs/vm37
```

### Restoring Virtual Machines to Alternative Locations

When you need to restore a virtual machine to a location other than its original location, or to a different ESX Server host, you use the virtual machine’s `catalog` file. `vcbMounter` creates this file for each virtual machine it backs up. The `catalog` file contains summary information about the virtual machine, such as its display name, its power state at the time of backup, and so on.

## To restore a virtual machine to an alternative location

- 1 Make a copy of the virtual machine's catalog file.  
See ["Copying a Catalog File"](#) on page 45.
- 2 In the copy of the catalog file, specify the new settings for datastores, folder path, and resource pool.  
See ["Editing a Catalog File"](#) on page 45.
- 3 Restore the virtual machine using vcbRestore.  
See ["Restoring Virtual Machines Using an Alternate Catalog"](#) on page 47

## Copying a Catalog File

When restoring a virtual machine to a location other than the original, you start by making a copy of the virtual machine's catalog file.

For example, you need to make a copy of the catalog file of the /home/VMs/vm37 virtual machine. Enter the following:

```
cp /home/VMs/vm37/catalog /tmp/catalog-vm37
```

## Editing a Catalog File

In the copy of the catalog file you made, you need to change the following settings:

- **Datastore** – The datastore identifies where to store all the files that comprise a virtual machine.
- **Folder path** – The virtual machine's folder path defines where the virtual machine will be placed in the VirtualCenter folder hierarchy.
- **Resource pool** – This host-specific configuration item determines the virtual machine's behavior with respect to DRS (Distributed Resource Scheduling). When you use multiple ESX Servers managed by VirtualCenter, this item also specifies the host that will run the virtual machine.

---

**NOTE** If you change the name of the virtual machine in the catalog file, vcbRestore doesn't pick up the new name from the file, but instead uses the original virtual machine name specified in the .vmx file.

You can change the name of the virtual machine later using the VI Client.

---

## Changing Datastore Paths

The datastore path in the catalog file identifies where to store all the files that comprise a virtual machine. Change datastore paths in the following entries:

- `disk.scsi*.diskname` – Names and locations of all disks associated with this virtual machine.
- `config.vmx` – Location for the virtual machine's main configuration file.
- `config.suspenddir` – Location for the memory snapshots taken when the virtual machine gets suspended.
- `config.logdir` – Location for the virtual machine's log files.

By default, all these entries use the same path, which points to the same directory on the same datastore. It is highly recommended that you follow this standard when changing the path.

The datastore paths have the following syntax:

```
[<datastore_name>] <path_on_datastore>
```

You can obtain a list of valid datastore names from the datastore browser in your VirtualCenter client, or by looking at the file system labels of your VMFS volumes in the service console under `/vmfs/volumes`.

## Changing Folder Paths

The virtual machine's folder path in the catalog file specifies the folder within the VirtualCenter folder hierarchy where the restored virtual machine will be placed.

### To change the folder path for the virtual machine

- 1 Identify the folder, which will store the virtual machine, by running the following command in the service console:

```
vcbUtil -c vmfolders
```

Running this command assumes that you have set up appropriate configuration options as described in [“Configuration File Settings”](#) on page 38.

- 2 In the catalog file, set the folder path to one of the folder paths printed out by the command above.

## Changing Resource Pools

The resource pools entry in the catalog file determines the virtual machine's behavior with respect to DRS (Distributed Resource Scheduling). When you use multiple ESX Servers managed by VirtualCenter, this item also specifies the host that will run the virtual machine.

### To change the resource pool setting for the virtual machine

- 1 Identify the resource pool, which the virtual machine will use, by running the following command:

```
vcbUtil -c resourcepools
```

Running this command assumes that you have set up appropriate configuration options as described in [“Configuration File Settings”](#) on page 38.

- 2 In the catalog file, set the resource pool to one of the options provided by the command above.

## Restoring Virtual Machines Using an Alternate Catalog

After modifying the settings in the virtual machine's alternate catalog, use this file to restore the virtual machine.

To restore the virtual machine, use the `-a` entry to specify the alternate catalog.

For example, to restore a virtual machine backed up under `/home/VMs/vm37` by using the alternate catalog file `/tmp/catalog-vm37`, enter:

```
vcbRestore -s /home/VMs/vm37 -a /tmp/catalog-vm37
```

## Non-interactive Use of the vcbRestore Utility

By default, vcbRestore prompts you what to do when the restore operation detects a file that already exists or a virtual machine already known to VirtualCenter.

If vcbRestore is used by a script in a non-interactive way, use the `-b` command-line entry to specify the default behavior. The following options are available:

### prompt

Prompts a user what to do before overwriting files or configurations of virtual machines already known to VirtualCenter.

### overwrite

Overwrites any existing files and virtual machine configurations known to VirtualCenter during restore.

### keep

Preserves existing files and configurations of virtual machines known to VirtualCenter without replacing them.

### abort

Terminates the restore operation after detecting an existing file or a virtual machine configuration already known to VirtualCenter.

# APPENDIX B **Restoring Virtual Machines from ESX Server 2.5.x to ESX Server 3.0**

This appendix describes how to restore virtual machines, which were created and backed up on ESX Server 2.5.x, in ESX Server 3.0 using the service console.

This appendix includes the following sections:

- [“Setting Configuration Parameters”](#) on page 50
- [“Restoring ESX 2.5.x Server Virtual Machines”](#) on page 51

## Setting Configuration Parameters

To restore virtual machines from ESX Server 2.5.x to ESX Server 3.0, you first need to set up configuration parameters in the `/etc/vmware/backuptools.conf` file.

To set up general parameters, follow recommendations in [“General Configuration Settings for Consolidated Backup Utilities”](#) on page 38.

In addition, define the following specific parameters.

### DSPATH

Specifies the path to a datastore where your restored virtual machine will reside. To avoid setting up this option individually for each virtual machine you restore, use the `%VMNAME%` entry. During the restore process, the base name of the virtual machine's `.vmx` configuration file substitutes this entry.

For example, to restore virtual machines to the `oldvms` datastore using the base name of the virtual machine's `.vmx` file, enter the following:

```
DSPATH="[oldvms] %VMNAME%"
```

This entry restores the virtual machine with the `myvm.vmx` file into `[oldvms] /myvm`.

### VMHOST

Specifies the host for the virtual machine you restore. The virtual machine will be powered on from this host.

### RESOURCEPOOL

Specifies the resource pool for the virtual machine you restore. For more information on this parameter, see [“Changing Resource Pools”](#) on page 47.

---

**NOTE** Make sure to select a valid resource pool on the host you specified in `VMHOST`. Typically, the resource pool name contains the name of the corresponding ESX Server host. You can also use the `%VMHOST%` entry, which will be replaced by the value you assigned to `VMHOST`.

---

### FOLDERPATH

Specifies the folder within the VirtualCenter hierarchy, in which the restored virtual machine will be placed. For more information on this parameter, see [“Changing Folder Paths”](#) on page 46.

## Restoring ESX 2.5.x Server Virtual Machines

After defining all necessary settings in the `/etc/vmware/backuptools.conf` file, you can restore virtual machines that were backed up on ESX Server 2.5.x.

The restore process is the same as for the ESX Server 3.0 virtual machines.

For more information, see [“Restoring Virtual Machines”](#) on page 44.



# Index

## A

archive bit **20**

## B

backup destinations **43**

BIOS UUID **41**

## C

catalog files **45**

change journal **20**

configuration file **38**

configuration files, exporting **18**

## D

directory, /vmfs/volumes **4**

disk images, exporting **18**

DNS name **41**

## F

file allocation table **22**

## I

IP address **41**

## K

Knowledge base **ix**

## M

metadata **4**

MoRef **42**

## N

new technology file system **22**

## P

physical compatibility **4**

physical SCSI disks **3**

## R

Report LUNs command **4**

## U

User groups **ix**

utilities **38**

## V

VCB utilities **38**

virtual compatibility **4**

virtual-to-physical clustering **4**

VMFS **3**

vmkfstools command **5**

VMware community forums **ix**

