

# HP Data Protector 8.00

## Integration Guide for Virtualization Environments

HP Part Number: N/A  
Published: June 2013  
Edition: Second



© Copyright 2013 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft®, Windows®, Windows XP®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Adobe and Acrobat are trademarks of Adobe Systems Incorporated.

Java is a registered trademark of Oracle and/or its affiliates.

Oracle® is a registered US trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of The Open Group.

LiveVault® is a registered trademark of Autonomy Corporation plc.

---

# Contents

Publication history.....	9
About this guide.....	10
Intended audience.....	10
Documentation set.....	10
Help.....	10
Guides.....	10
Documentation map.....	13
Abbreviations.....	13
Map.....	14
Integrations.....	14
Document conventions and symbols.....	15
Data Protector graphical user interface.....	16
General information.....	16
HP technical support.....	16
Subscription service.....	16
HP websites.....	17
Documentation feedback.....	17
I VMware.....	18
1 Data Protector Virtual Environment integration.....	20
Introduction.....	20
Integration concepts.....	21
Supported environments.....	21
vCloud Director environment.....	21
vCenter environment.....	21
Migration of virtual machines.....	22
vCenter Server system in a cluster.....	23
Standalone ESX/ESXi Server environment.....	23
Data Protector components.....	24
Data Protector Cell Manager.....	24
Data Protector Virtual Environment Integration component.....	24
Data Protector Disk Agent component.....	25
Data Protector Media Agents.....	25
Backup concepts.....	25
What is backed up?.....	25
Virtual machines.....	26
Virtual machine templates.....	26
vStorage Image backup method.....	26
Snapshot management.....	27
vCD vStorage Image backup method.....	28
Backup types.....	28
Snapshot handling mode.....	29
Changed block tracking.....	35
Quiescence.....	36
Disk space requirements.....	37
Free space required option.....	37
Backup disk buffer.....	38
Backup parallelism.....	38
Backup considerations.....	38
Restore concepts.....	39
Restore of VMware objects backed up with vStorage Image method.....	39

Restore to a datacenter.....	39
Restore to a directory.....	39
Restore of VMware objects backed up with vCD vStorage Image method.....	40
Restore to an organization.....	40
Restore to a directory.....	40
Restore chain.....	40
Restore considerations.....	40
Performance.....	41
Configuring the integration.....	41
Prerequisites.....	42
Before you begin.....	42
Importing and configuring VMware clients.....	42
Changing the configuration of VMware clients.....	46
Using the Data Protector GUI.....	46
Using the Data Protector CLI.....	47
Checking the configuration of VMware clients.....	48
Using the Data Protector GUI.....	48
Using the Data Protector CLI.....	48
Configuring virtual machines.....	49
Using the Data Protector GUI.....	49
Using the Data Protector CLI.....	54
Customizing the Data Protector behavior with omnirc options.....	56
Backup.....	56
Backup limitations.....	56
Creating backup specifications.....	56
Modifying backup specifications.....	61
Scheduling backup sessions.....	64
Scheduling example.....	64
Previewing backup sessions.....	65
Using the Data Protector GUI.....	65
Using the Data Protector CLI.....	65
What happens during the preview?.....	65
Starting backup sessions.....	65
Using the Data Protector GUI.....	66
Using the Data Protector CLI.....	66
Preparing for disaster recovery.....	67
Restore.....	67
Restore limitations.....	67
Finding information for restore.....	68
Using the Data Protector GUI.....	68
Using the Data Protector CLI.....	69
Restoring using the Data Protector GUI.....	69
Restoring using the Data Protector CLI.....	76
Recovering virtual machines manually.....	77
Recovering virtual machines after restore to a directory.....	78
Recovering with the VM configuration file in the VMX format.....	78
Recovering with the VM configuration file in the XML format.....	83
Recovering virtual machines after restore to a datacenter.....	83
Recovering virtual machines after restore to an organization.....	83
Restoring using another device.....	83
Cleaning up a datastore after a failed restore.....	84
Disaster recovery.....	84
Monitoring sessions.....	84
Troubleshooting.....	84
Before you begin.....	85

Checks and verifications.....	85
Problems.....	85
<b>2 Data Protector VMware (Legacy) integration.....</b>	<b>89</b>
Introduction.....	89
Integration concepts.....	89
Supported environments.....	90
VirtualCenter environment.....	90
Migration of virtual machines.....	91
VirtualCenter Server system in a cluster.....	92
Standalone ESX Server environment.....	92
ESXi Server systems.....	92
Backup.....	92
Restore.....	92
Limitations in VMware vSphere environments.....	93
Data Protector installation components.....	93
Data Protector Cell Manager.....	93
Data Protector VMware Integration (Legacy) component.....	93
Data Protector Media Agents.....	93
Backup concepts.....	93
What is backed up?.....	94
Virtual machines.....	94
Filesystems.....	94
Backup methods.....	95
Snapshot method.....	95
Suspend method.....	101
VCBimage method.....	102
VCBfile method.....	103
Backup types.....	105
Considerations.....	105
Restore concepts.....	106
Restore of virtual machines.....	106
Restore to datacenter.....	106
Restore to default directory.....	107
Restore of filesystems.....	107
Restore considerations.....	107
Configuring the integration.....	108
Prerequisites.....	108
Before you begin.....	108
Configuring VMware users.....	109
Configuring clusters.....	109
ESX Server systems in a cluster.....	109
Configuring backup proxy systems.....	109
Configuring VMware management clients.....	109
Using the Data Protector GUI.....	110
Using the Data Protector CLI.....	112
Configuring virtual machines.....	112
Using the Data Protector GUI.....	112
Using the Data Protector CLI.....	114
Checking the configuration of VMware clients.....	115
Using the Data Protector GUI.....	115
Using the Data Protector CLI.....	115
Backup.....	116
Creating backup specifications.....	116
Modifying backup specifications.....	119

Scheduling backup sessions.....	119
Scheduling example.....	119
Previewing backup sessions.....	120
Using the Data Protector GUI.....	120
Using the Data Protector CLI.....	120
What happens during the preview?.....	120
Starting backup sessions.....	121
Using the Data Protector GUI.....	121
Using the Data Protector CLI.....	121
Preparing for disaster recovery.....	122
Restore.....	122
Finding information for restore.....	122
Using the Data Protector GUI.....	122
Using the Data Protector CLI.....	123
Restoring using the Data Protector GUI.....	123
Restoring using the Data Protector CLI.....	128
Recovering virtual machines manually.....	130
Recovery after restoring outside a datacenter.....	130
Recovery after restoring to a datacenter.....	135
Linking virtual machine disk files.....	135
Consolidating virtual machine disks.....	137
Importing virtual machine disks.....	138
Restoring using another device.....	140
Disaster recovery.....	140
Monitoring sessions.....	140
Troubleshooting.....	141
Before you begin.....	141
Checks and verifications.....	141
Problems.....	141
<b>II Microsoft Hyper-V.....</b>	<b>144</b>
<b>3 Data Protector Virtual Environment integration.....</b>	<b>146</b>
Introduction.....	146
Integration concepts.....	147
Supported environments.....	147
Standalone environments.....	147
Clustered environments.....	147
Migration of virtual machines.....	150
Cluster Shared Volumes.....	151
Hyper-V Replica.....	151
Virtual machines on Windows file shares.....	151
Data Protector installation components.....	151
Data Protector Cell Manager.....	151
Data Protector Virtual Environment Integration component.....	151
Data Protector Disk Agent component.....	151
Data Protector MS Volume Shadow Copy Integration component.....	152
Data Protector Media Agent component.....	152
Backup concepts.....	152
Hyper-V Image backup method.....	152
Backup types.....	153
Microsoft Hyper-V backup types.....	154
VSS backup types.....	154
Quiescence.....	154
Restore chain protection.....	155
Backup considerations.....	155

Object copy considerations.....	159
Restore concepts.....	159
Restore of virtual machines.....	159
Restore to the default location.....	159
Restore to a different location.....	160
Restore to a directory .....	160
Restore chain validation.....	160
Restore considerations.....	161
Configuring the integration.....	163
Prerequisites.....	163
Limitations.....	164
Before you begin.....	164
Enabling automatic mounting of new volumes on Microsoft Hyper-V systems.....	164
Configuring Microsoft Hyper-V clusters.....	164
Importing and configuring Microsoft Hyper-V systems.....	164
Changing the configuration of Microsoft Hyper-V systems.....	166
Using the Data Protector GUI.....	166
Using the Data Protector CLI.....	166
Customizing the Data Protector behavior with omnirc options.....	167
Backup.....	167
Creating backup specifications.....	167
Modifying backup specifications.....	171
Scheduling backup sessions.....	172
Scheduling example.....	172
Starting backup sessions.....	173
Using the Data Protector GUI.....	173
Using the Data Protector CLI.....	173
Restore.....	174
Limitations.....	174
Finding information for restore.....	174
Using the Data Protector GUI.....	174
Using the Data Protector CLI.....	175
Restoring using the Data Protector GUI.....	175
Restoring using the Data Protector CLI.....	177
Merging virtual machine snapshots manually.....	179
Restore of cluster virtual machines.....	179
Restoring a replicated virtual machine.....	180
Re-enabling the replication.....	180
Reverting a restored replica VM to an application-consistent recovery point.....	180
Restoring using another device.....	180
Monitoring sessions.....	180
Troubleshooting.....	180
Before you begin.....	181
Checks and verifications.....	181
Problems.....	181
III Citrix XenServer.....	185
4 Data Protector Citrix XenServer script solution.....	186
Introduction.....	186
Integration concepts.....	186
Types of backup.....	186
Online backup.....	186
Offline backup.....	186
Disaster recovery.....	186
Backup processes.....	187

Online backup.....	187
Offline backup.....	187
Types of restore.....	187
Restore processes.....	187
Restore from online backup.....	187
Restore from offline backup.....	188
Restore considerations.....	188
Main integration components.....	189
Installation of the integration.....	189
Prerequisites.....	189
Installation of the integration scripts.....	190
Integration script functions.....	191
Backup using the integration.....	192
Updating configuration script DPxen_config.py for backup.....	192
Example backup configurations.....	193
Creating a backup specification.....	194
Restore using the integration.....	195
Updating configuration script DPxen_config.py for restore.....	195
Example restore configurations.....	196
Specifying a restore.....	197
Notes on restore.....	199
Special considerations.....	199
Further information.....	199
Glossary.....	200
Index.....	229



---

## Publication history

Guide updates may be issued between editions to correct errors or document product changes. To ensure that you receive updated or new editions, subscribe to the appropriate product support service. See your HP sales representative for details.

**Table 1 Edition history**

Part number	Guide edition	Product
N/A	June 2013	Data Protector release 8.00
N/A	June 2013 (second edition)	Data Protector release 8.00

---

# About this guide

This guide describes how to configure and use Data Protector to protect your virtualizations environments.

## Intended audience

This guide is intended for backup administrators responsible for planning, setting up, and maintaining network backups. It assumes you are familiar with:

- Basic Data Protector functionality
- Administration of the respective virtual environment

Conceptual information can be found in the *HP Data Protector Concepts Guide*, which is recommended to fully understand the fundamentals and the model of Data Protector.

## Documentation set

The Help and other guides provide related information.

---

**NOTE:** The documentation set available at the HP support website at <http://support.openview.hp.com/selfsolve/manuals> contains the latest updates and corrections.

---

## Help

Data Protector provides Help topics and context-sensitive (F1) Help for Windows and UNIX platforms. Install the Help during the Data Protector setup procedure by selecting the installation component English Documentation (Guides, Help) (Windows systems) or OB2-DOCS (on UNIX systems). Once installed, the Help resides in the following directory:

**Windows systems:** `Data_Protector_home\help\enu`

**UNIX systems:** `/opt/omni/help/C/help_topics`

You can access the Help from the top-level directory of any installation DVD-ROM without installing Data Protector:

**Windows systems:** Open `DP_help.chm`.

**UNIX systems:** Unpack the zipped tar file `DP_help.tar.gz` and open `DP_help.htm`.

## Guides

Data Protector guides are available in the electronic PDF format. Install the PDF files during the Data Protector setup procedure by selecting the installation component English Documentation (Guides, Help) (on Windows systems) or OB2-DOCS (on UNIX systems). Once installed, the guides reside in the following directory:

**Windows systems:** `Data_Protector_home\docs`

**UNIX systems:** `/opt/omni/doc/C`

You can also access the guides:

- From the **Help** menu of the Data Protector graphical user interface
- From the HP support website at <http://support.openview.hp.com/selfsolve/manuals> (where the most up-to-date guide versions are available)

Data Protector guides are:

- *HP Data Protector Getting Started Guide*  
This guide contains information to get you started with using Data Protector. It lists installation prerequisites, provides instructions on installing and configuring a basic backup environment and procedures for performing backup and restore. It also lists resources for further information.
- *HP Data Protector Concepts Guide*  
This guide describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented Help.
- *HP Data Protector Installation and Licensing Guide*  
This guide describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This guide also gives details on how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.
- *HP Data Protector Troubleshooting Guide*  
This guide describes how to troubleshoot problems you may encounter when using Data Protector.
- *HP Data Protector Disaster Recovery Guide*  
This guide describes how to plan, prepare for, test, and perform a disaster recovery.
- *HP Data Protector Command Line Interface Reference*  
This guide describes the Data Protector command-line interface, command options, and their usage as well as provides some basic command-line examples. It is located in the following directory:  
**Windows systems:** `Data_Protector_home\docs\MAN`  
**UNIX systems:** `/opt/omni/doc/C/`  
On UNIX systems, you can use the `omniintro` man page to display a list of the available Data Protector commands. You can then execute the `man CommandName` command to retrieve information about each Data Protector command.
- *HP Data Protector Product Announcements, Software Notes, and References*  
This guide gives a description of new features of HP Data Protector 8.00. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.
- *HP Data Protector Integration Guides*  
These guides describe how to configure and use Data Protector to back up and restore various databases and applications. They are intended for backup administrators and operators. There are six guides:
  - *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*  
This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server.
  - *HP Data Protector Integration Guide for Oracle and SAP*  
This guide describes the integrations of Data Protector with Oracle Server, SAP R/3, and SAP MaxDB.

- *HP Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*  
This guide describes the integrations of Data Protector with the following IBM applications: Informix Server, IBM DB2 UDB, and Lotus Notes/Domino Server.
- *HP Data Protector Integration Guide for Sybase and Network Data Management Protocol Server*  
This guide describes the integrations of Data Protector with Sybase Server and Network Data Management Protocol Server.
- *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*  
This guide describes the integration of Data Protector with the Microsoft Volume Shadow Copy Service. This guide also documents application writer specifics.
- *HP Data Protector Integration Guide for Virtualization Environments*  
This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, and Citrix XenServer.
- *HP Data Protector Zero Downtime Backup Concepts Guide*  
This guide describes Data Protector zero downtime backup and instant recovery concepts and provides background information on how Data Protector works in a zero downtime backup environment. It is intended to be used with the task-oriented *HP Data Protector Zero Downtime Backup Administrator's Guide* and the *HP Data Protector Zero Downtime Backup Integration Guide*.
- *HP Data Protector Zero Downtime Backup Administrator's Guide*  
This guide describes how to configure and use the integration of Data Protector with HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP 3PAR StoreServ Storage, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.
- *HP Data Protector Zero Downtime Backup Integration Guide*  
This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases.
- *HP Data Protector Granular Recovery Extension User Guide for Microsoft Exchange Server*  
This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft Exchange Server. Graphical user interface of the Data Protector Granular Recovery Extension for Microsoft Exchange Server is integrated into the Microsoft Management Console. This guide is intended for Microsoft Exchange Server administrators and Data Protector backup administrators.
- *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*  
This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft SharePoint Server. The Data Protector Granular Recovery Extension is integrated into Microsoft SharePoint Server Central Administration and enables you to recover individual items. This guide is intended for Microsoft SharePoint Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for VMware vSphere*  
This guide describes how to configure and use the Data Protector Granular Recovery Extension for VMware vSphere. The Data Protector Granular Recovery Extension is integrated into VMware vCenter Server and enables you to recover individual items. This guide is intended for VMware vCenter Server users and Data Protector backup administrators.
- *HP Data Protector Deduplication*  
This technical white paper describes basic data deduplication concepts, principles of Data Protector integration with Backup to Disk devices and its use of deduplication. It also provides instructions how to configure and use deduplication in Data Protector backup environments.
- *HP Data Protector Integration with Autonomy IDOL Server*  
This technical white paper describes all aspects of integrating Data Protector with Autonomy IDOL Server: integration concepts, installation and configuration, Data Protector backup image indexing, full content search-based restore, and troubleshooting.
- *HP Data Protector Integration with Autonomy LiveVault*  
This technical white paper all aspects of integrating Data Protector with Autonomy LiveVault: integration concepts, installation and configuration, backup policy management, cloud backup, cloud restore, and troubleshooting.

## Documentation map

### Abbreviations

Abbreviations in the documentation map that follows are explained below. The documentation item titles are all preceded by the words "HP Data Protector".

Abbreviation	Documentation item
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE Exchange	Granular Recovery Extension User Guide for Microsoft Exchange Server
GRE SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Help
Install	Installation and Licensing Guide
IG IBM	Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino
IG MS	Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server
IG VSS	Integration Guide for Microsoft Volume Shadow Copy Service
IG O/S	Integration Guide for Oracle and SAP
IG Var	Integration Guide for Sybase and Network Data Management Protocol Server
IG VirtEnv	Integration Guide for Virtualization Environments
IG IDOL	Integration with Autonomy IDOL Server
IG LV	Integration with Autonomy LiveVault

Abbreviation	Documentation item
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concepts	ZDB Concepts Guide
ZDB IG	ZDB Integration Guide

## Map

The following table shows where to find information of different kinds. Shaded squares are a good place to look first.

									Integr. guides						ZDB			GRE		
	Help	GS	Concepts	Install	Trouble	DR	CLI	PA	MS	O/S	IBM	Var	VSS	VirtEnv	Concepts	Admin	IG	Exchange	SPS	VMware
Backup	X	X	X						X	X	X	X	X	X	X	X	X			
CLI							X													
Concepts, techniques	X		X						X	X	X	X	X	X	X	X	X	X	X	X
Disaster recovery	X		X			X														
Installation, upgrade	X	X		X				X												
Instant recovery	X		X												X	X	X			
Licensing	X			X				X												
Limitations	X				X			X	X	X	X	X	X	X			X			
New features	X							X												
Planning strategy	X		X												X					
Procedures, tasks	X			X	X	X			X	X	X	X	X	X		X	X	X	X	X
Recommendations			X					X							X					
Requirements				X				X	X	X	X	X	X	X						
Restore	X	X	X						X	X	X	X	X	X		X	X	X	X	X
Supported configurations															X					
Troubleshooting	X			X	X				X	X	X	X	X	X		X	X	X	X	X

## Integrations

Look in these guides for details of the integrations with the following software applications:

Software application	Guides
Autonomy IDOL Server	IG IDOL
Autonomy LiveVault	IG LV
IBM DB2 UDB	IG IBM
Informix Server	IG IBM
Lotus Notes/Domino Server	IG IBM
Microsoft Exchange Server	IG MS, ZDB IG, GRE Exchange

Software application	Guides
Microsoft Hyper-V	IG VirtEnv
Microsoft SharePoint Server	IG MS, ZDB IG, GRE SPS
Microsoft SQL Server	IG MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG VSS
Network Data Management Protocol (NDMP) Server	IG Var
Oracle Server	IG O/S, ZDB IG
SAP MaxDB	IG O/S
SAP R/3	IG O/S, ZDB IG
Sybase Server	IG Var
VMware vCloud Director	IG VirtEnv
VMware vSphere	IG VirtEnv, GRE VMware

Look in these guides for details of the integrations with the following families of disk array systems:

Disk array family	Guides
EMC Symmetrix	all ZDB
HP P4000 SAN Solutions	ZDB Concepts, ZDB Admin, IG VSS
HP P6000 EVA Disk Array Family	all ZDB, IG VSS
HP P9000 XP Disk Array Family	all ZDB, IG VSS
HP 3PAR StoreServ Storage	ZDB Concepts, ZDB Admin, IG VSS

## Document conventions and symbols

**Table 2 Document conventions**

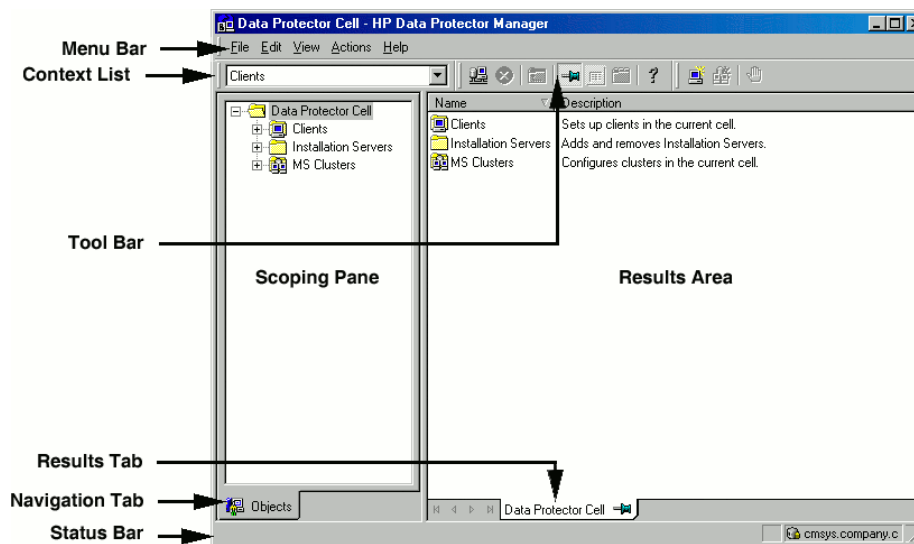
Convention	Element
Blue text: “ <a href="#">Document conventions</a> ” (page 15)	Cross-reference links and e-mail addresses
Blue, underlined text: <a href="http://www.hp.com">http://www.hp.com</a>	Website addresses
<b>Bold</b> text	<ul style="list-style-type: none"> <li>Keys that are pressed</li> <li>Text typed into a GUI element, such as a box</li> <li>GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes</li> </ul>
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none"> <li>File and directory names</li> <li>System output</li> <li>Code</li> <li>Commands, their arguments, and argument values</li> </ul>
<i>Monospace, italic</i> text	<ul style="list-style-type: none"> <li>Code variables</li> <li>Command variables</li> </ul>
<b>Monospace, bold</b> text	Emphasized monospace text

- 
- CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.
- 
- IMPORTANT:** Provides clarifying information or specific instructions.
- 
- NOTE:** Provides additional information.
- 
- TIP:** Provides helpful hints and shortcuts.
- 

## Data Protector graphical user interface

Data Protector provides a graphical user interface for Microsoft Windows operating systems. For information about it, see the *HP Data Protector Help*.

**Figure 1 Data Protector graphical user interface**



## General information

General information about Data Protector can be found at <http://www.hp.com/go/dataprotector>.

## HP technical support

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

## Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/e-updates>



After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

## HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/software>
- <http://support.openview.hp.com/selfsolve/manuals>
- <http://www.hp.com/support/downloads>

## Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message with the subject line Feedback on Data Protector documentation to [AutonomyTPFeedback@hp.com](mailto:AutonomyTPFeedback@hp.com). All submissions become the property of HP.

---

## Part I VMware

Data Protector offers different ways to back up VMware virtual machines online. Choose the appropriate backup and restore solution depending on your VMware solution and the desired functionality.

### VMware Virtual Infrastructure

To back up virtual machines from a VMware Virtual Infrastructure environment, you can use the following solution:

- **Data Protector VMware (Legacy) integration**

This integration enables you to back up and restore the following VMware logical objects:

- Datacenters
- Virtual machines
- Virtual machine disks

The integration uses the VI SDK, a web-service API, to communicate with VMware Virtual Infrastructure. See [“Data Protector VMware \(Legacy\) integration” \(page 89\)](#).

### VMware vSphere

To back up virtual machines from a VMware vSphere environment, you can use the following solutions:

- **Data Protector Virtual Environment integration**

This integration enables you to back up and restore the following VMware logical objects:

- Datacenters
- Virtual machines
- Virtual machine disks
- Virtual machine templates

The integration uses the VMware vStorage API to communicate with VMware vSphere. See [“Data Protector Virtual Environment integration” \(page 20\)](#).

- **Data Protector VMware (Legacy) integration**

This integration enables you to back up and restore the following VMware logical objects:

- Datacenters
- Virtual machines
- Virtual machine disks

The integration uses the VI SDK, a web-service API, to communicate with VMware vSphere. See [“Data Protector VMware \(Legacy\) integration” \(page 89\)](#).

## VMware vCloud Director

To back up virtual machines from a VMware vCloud Director environment, you can use the following solution:

- **Data Protector Virtual Environment integration**

This integration enables you to back up and restore the following VMware logical objects:

- Virtual applications
- Virtual machines

The integration uses the VMware vStorage API to communicate with VMware vCloud Director.

See “Data Protector Virtual Environment integration” (page 20).

**NOTE:** You can also back up VMware virtual machines using common Data Protector filesystem backup functionality, which operates on the file level. The smallest object that you can back up or restore this way is a file. To ensure data consistency, you must shut the virtual machines down before starting a backup session.

**Table 3 Data Protector backup solutions for VMware**

Feature		Disk Agent (filesystem backup)	VMware (Legacy) integration	Virtual Environment integration
Online backup			✓ <sup>1</sup>	✓
Crash-consistent backup		✓	✓	✓
Application-consistent backup				✓
Granularity		File	VM disk	VM disk
Backup types	Full	✓	✓	✓
	Incremental	✓	✓	✓
	Differential	✓	✓	✓
Are zero downtime backup and instant recovery supported?		✓		
Is VCB software required?			✓ <sup>2</sup>	
Is changed block tracking supported?				✓
Where does the Data Protector integration component need to be installed?		ESX(i) Server system	<ul style="list-style-type: none"><li>• ESX(i) Server system</li><li>• vCenter Server system</li><li>• Backup proxy</li></ul>	Backup host <sup>3</sup>
Extra licenses needed		<ul style="list-style-type: none"><li>• Zero Downtime Backup (optional)</li><li>• Instant Recovery (optional)</li></ul>	<ul style="list-style-type: none"><li>• On-line Extension</li></ul>	<ul style="list-style-type: none"><li>• On-line Extension</li></ul>

<sup>1</sup> In a VMware vSphere environment, online backup of virtual machines using the Data Protector VMware (Legacy) integration “Snapshot method” (page 95) is not supported.

<sup>2</sup> The VCB software is needed to back up virtual machines that reside on an ESXi Server system.

<sup>3</sup> The integration agent can be installed on a dedicated system, from which Data Protector remotely contacts vCenter Server or ESX(i) Server system, increasing the range of support for VMware products compared with the Data Protector VMware (Legacy) integration.

---

# 1 Data Protector Virtual Environment integration

## Introduction

This chapter explains how to configure and use the Data Protector Virtual Environment integration for VMware vSphere and VMware vCloud Director. Data Protector integrates with VMware vSphere and VMware vCloud Director, including ESX Server, ESXi Server systems, and vCenter Server systems, to back up and restore the following **VMware objects**:

VMware vSphere:

- Virtual machines
- Virtual machine disks
- Virtual machine templates

VMware vCloud Director:

- Virtual applications
- Virtual machines

## Backup

The following backup methods are available:

- vStorage Image
- vCD vStorage Image

These are snapshot-based methods that can be used to back up virtual machines while powered off (**offline backup**) or actively used (**online backup**).

For details on the backup methods, see [“vStorage Image backup method” \(page 26\)](#).

Data Protector offers interactive and scheduled backups of the following types:

- Full
- Incremental
- Differential

For details on the backup types see, [“Backup types” \(page 28\)](#).

## Restore

Virtual machines can be restored:

- **To a datacenter** (integration with VMware vSphere)  
This datacenter can be a standalone ESX(i) Server system or any of the datacenters created and managed by a vCenter Server system.
- **To an organization** (integration with VMware vCloud Director)  
This organization can be any of the organizations created and managed by a VMware vCloud Director.
- **To a directory**  
This can be a directory on any client with the Data Protector Virtual Environment component installed. After such a restore, you need to manually move the restored virtual machine images to an ESX Server or ESXi Server system, using the VMware Converter.

This chapter provides information specific to the Data Protector Virtual Environment integration for VMware vSphere and VMware vCloud Director. For other limitations and recommendations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HP Data Protector Help*.

# Integration concepts

## Supported environments

Data Protector supports environments where ESX Server and/or ESXi Server systems are managed through a vCloud Director (**vCloud Director environments**), through a vCenter Server (**vCenter environments**) as well as environments with standalone ESX and ESXi Server systems (**standalone ESX(i) Server environments**). Mixed environments, in which some of the ESX Server and/or ESXi Server systems are managed through a vCenter Server system and some are standalone, are also supported. You can even have multiple vCenter Server systems in your environment, each managing its own set of ESX Server and/or ESXi Server systems.

### vCloud Director environment

VMware vCloud Director is software for managing resources from several vCenter Server systems. It enables you to create and manage virtual datacenters, each of which can be used by a different organization. Data Protector supports backup and restore of virtual machines from such virtual datacenters.

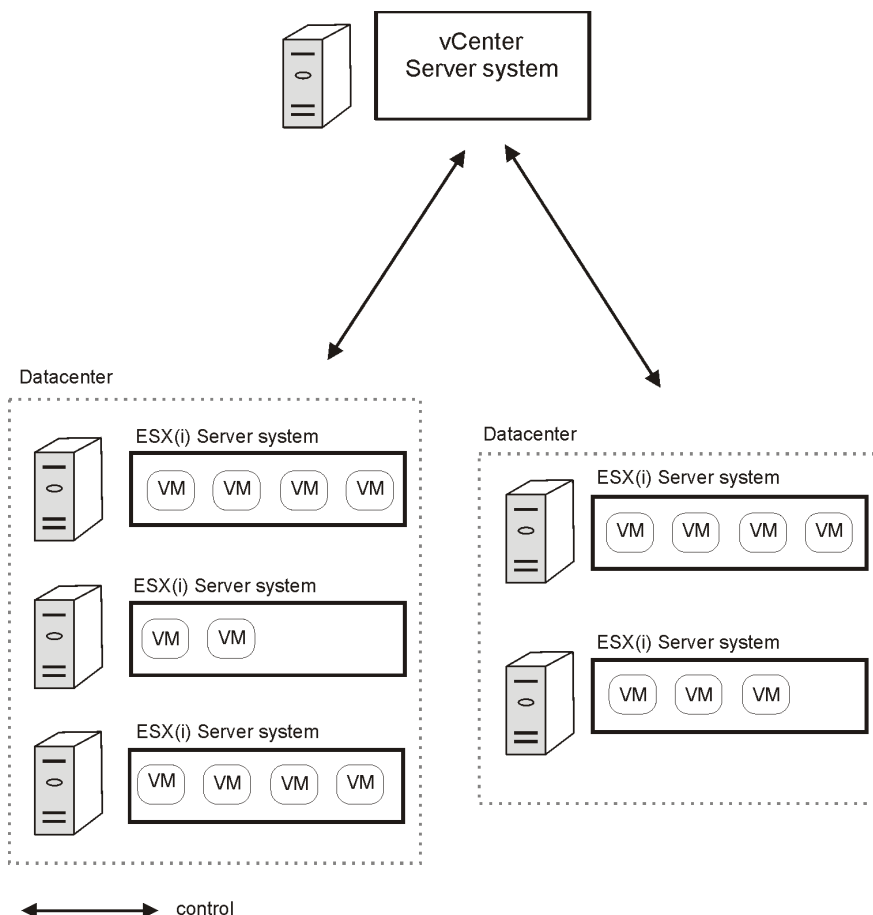
**NOTE:** To be able to back up virtual machines from virtual datacenters using the vCD vStorage Image backup method, make sure that you import all vCenter Server systems used by VMware vCloud Director in the Data Protector Cell as **VMware vCenter** clients.

### vCenter environment

In a vCenter environment, Data Protector communicates with the VMware vSphere through the vCenter Server system. All backup and restore requests are sent there.

In one session, you can back up virtual machines from one or multiple datacenters.

**Figure 2 vCenter environment**



ESX Server or ESXi Server system	VMware platform capable of hosting multiple virtual machines.
VM	Virtual machine. Virtualized x86 or x64 PC environment, in which a guest operating system and associated application software can run.
Datacenter	An organizational unit that consists of one or more ESX Server and/or ESXi Server systems and the related storage for virtual machines (datastores). Datastores can reside on local disks/RAID, iSCSI or SAN storage.

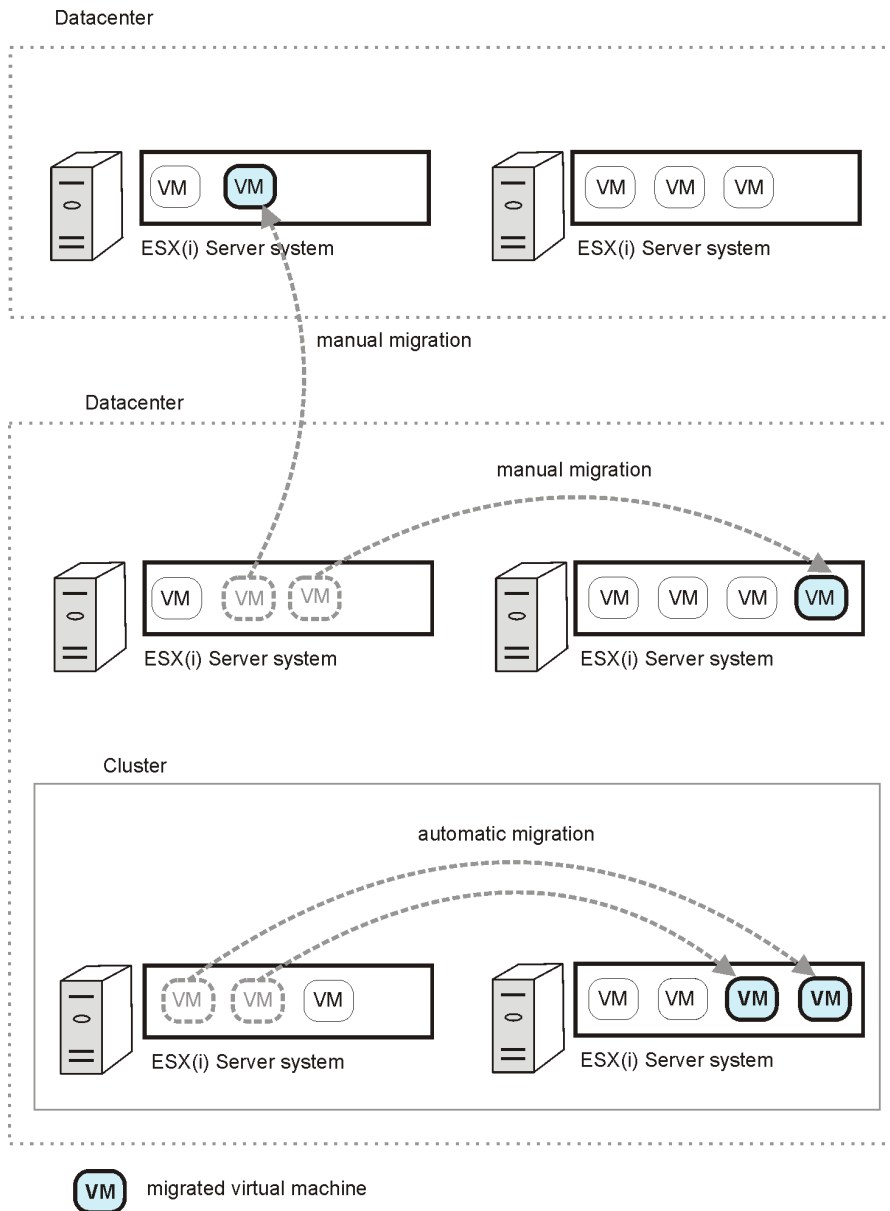
### Migration of virtual machines

In a vCenter environment, Data Protector supports migration (using VMotion and Storage VMotion) of virtual machines between ESX(i) Server systems in the same datacenter and, for supported versions of VMware vSphere, in different datacenters.

Virtual machines migrate from one ESX(i) Server system to another for various reasons:

- If ESX(i) Server systems are configured in a **VMware high availability cluster**, virtual machines automatically migrate when the original ESX(i) Server system fails.
- If ESX(i) Server systems are configured in a **VMware load balancing cluster**, virtual machines automatically migrate to ESX(i) Server systems with less workload.
- You can start a migration of a virtual machine manually, using the VMware vSphere client.

**Figure 3 Migration of virtual machines**



Whatever the reason for the migration, you do not need to create a new backup specification afterwards. Data Protector will automatically find the migrated virtual machines and back them up.

### vCenter Server system in a cluster

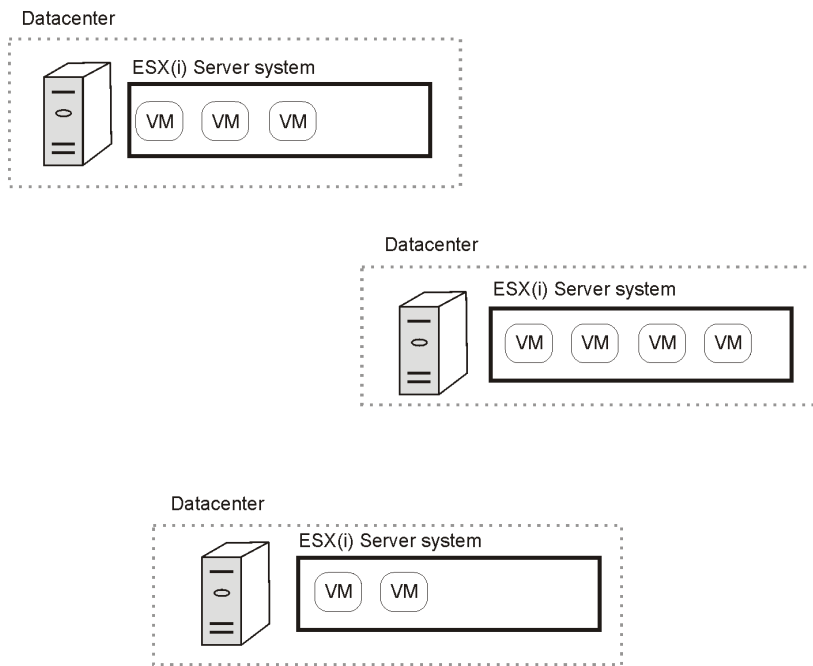
Data Protector also supports environments with a vCenter Server system in a Microsoft Cluster Service cluster. After a failover in such a cluster, you do not need to change the backup specification. However, if the failover occurs during a backup or restore session, the session fails and has to be restarted.

### Standalone ESX/ESXi Server environment

In a standalone ESX(i) Server environment, Data Protector communicates with VMware vSphere through the ESX(i) Server system. All backup and restore requests are sent there.

In a single session, you can back up virtual machines from only one ESX(i) Server system.

**Figure 4 Standalone ESX/ESXi Server environment**



## Data Protector components

### Data Protector Cell Manager

The Data Protector Cell Manager can be installed on a virtual machine, vCenter Server system, or a separate system outside the virtualization environment.

### Data Protector Virtual Environment Integration component

The Data Protector Virtual Environment Integration component (**VEAgent**) must be installed on at least one Data Protector client in the cell. This client will be called the **backup host** and can be:

- A virtual machine
- The vCenter Server system
- The Data Protector Cell Manager
- A dedicated physical backup host

The component includes the following major parts:

- `vepa_bar.exe`, activated during backup and restore operations on virtual environments.
- `vepa_util.exe`, activated during browsing and query operations on virtual environments.
- `vepalib_vmware.dll`, a dynamic link library for VMware vSphere specific backup, restore, query, and browsing tasks.
- `vepalib_vcd.dll`, a dynamic link library for VMware vCloud Director specific backup, restore, query, and browsing tasks.
- `vepalib_hyperv.dll`, a dynamic link library for Microsoft Hyper-V specific backup, restore, query, and browsing tasks.<sup>1</sup>

---

**NOTE:** VEPA stands for Virtual Environment Protection Agent.

---

1. The Data Protector Virtual Environment integration can also be used to back up virtual machines from Microsoft Hyper-V virtualization environments. For details, see the Microsoft Hyper-V part.



## Data Protector Disk Agent component

The Data Protector `Disk Agent` component must be installed on the backup host to use the browse directory button (this button is used when you restore to a directory on a backup host).

## Data Protector Media Agents

The Data Protector `Media Agent` component must be installed on clients that will transfer data to backup devices.

---

**NOTE:** The `Virtual Environment Integration` component does not enable the backup host to write to backup devices. You still need a client with the `Media Agent` component installed. However, note that the `Cell Manager`, `Media Agent`, and `Virtual Environment Integration` components can all be installed together on the same system (physical or virtual).

---

## Backup concepts

### What is backed up?

Using the Data Protector `Virtual Environment` integration, you can back up the following VMware objects:

VMware vSphere:

- Virtual machines
- Virtual machine disks
- Virtual machine templates

VMware vCloud Director:

- Virtual applications
- Virtual machines

Data Protector identifies datacenters and virtual machines by their VMware vSphere inventory path. A standalone ESX Server system has only one datacenter `/ha-datacenter` and two folders: `/host` and `/vm`. Virtual machines are stored in the folder `/host`.

*Example:*

Datacenter: `/ha-datacenter`

Virtual machine: `/vm/myvm1`

In a vCenter environment, you can organize virtual machines and datacenters within folders that you create yourself. If you subsequently move a virtual machine, you do not need to create a new backup specification because Data Protector will find a virtual machine using its UUID.

*Example:*

Virtual machine: `/vm/myfolder1/myfolder2/.../myvm2`

Datacenter: `/myfolder/mydatacenter`

In a vCloud Director environment, you can organize virtual machines within vApps, vDatacenters, and organizations that you create yourself.

*Example:*

Virtual machine: `/ORG22/vDCOrg22/vAppORG22/vm1Org22`

Organization: `/vCD1/Mngmt/ORG22`

## Virtual machines

When you back up a virtual machine, you actually back up virtual machine files of the following types:

- .vmx
- .vmdk
- .vmsn

## Memory file

You can also specify that when a virtual machine is online during the backup, its memory is saved to a file and backed up together with other virtual machine files.

---

**NOTE:** The Backup memory file option can only be set at backup specification level (which means it applies to all virtual machines selected in the backup specification).

---

## Virtual machine disks

Data Protector supports backup of individual virtual machine disks when using a VStorage Image backup method. In this case, all virtual machine files are backed up, except for virtual machine disks that are not specified. You can run full, incremental, and differential backups.

---

**NOTE:** After you add a new disk to a virtual machine, make sure you run a full backup session for the updated virtual machine.

---

## Virtual machine templates

You can also back up virtual machine templates when using a VStorage Image backup method. When you create a backup specification, expand the **vm** folder and select the desired virtual machine templates.

## vStorage Image backup method

The vStorage Image backup method provided by the Data Protector Virtual Environment integration is based on the VMware vStorage technology. For this method, a single central **backup host** is used to back up all virtual machines hosted by ESX(i) Server systems in a Data Protector cell. This backup host can be a dedicated physical host, a virtual machine, or the Cell Manager. The important thing is that it has the Data Protector Virtual Environment Integration component (**VEAgent**) installed.

During a vStorage Image backup, VEAgent first establishes connections between the backup host and the virtualization host (an ESX(i) Server system). This connection can either be through a vCenter Server system (in the case of a vCenter environment), or direct (in the case of a standalone ESX(i) Server environment). It then requests a snapshot of the virtual machine that is to be backed up, via the vStorage API for Data Protection (VADP). This snapshot is used during the period of the backup in order to keep the virtual machine in a consistent state.

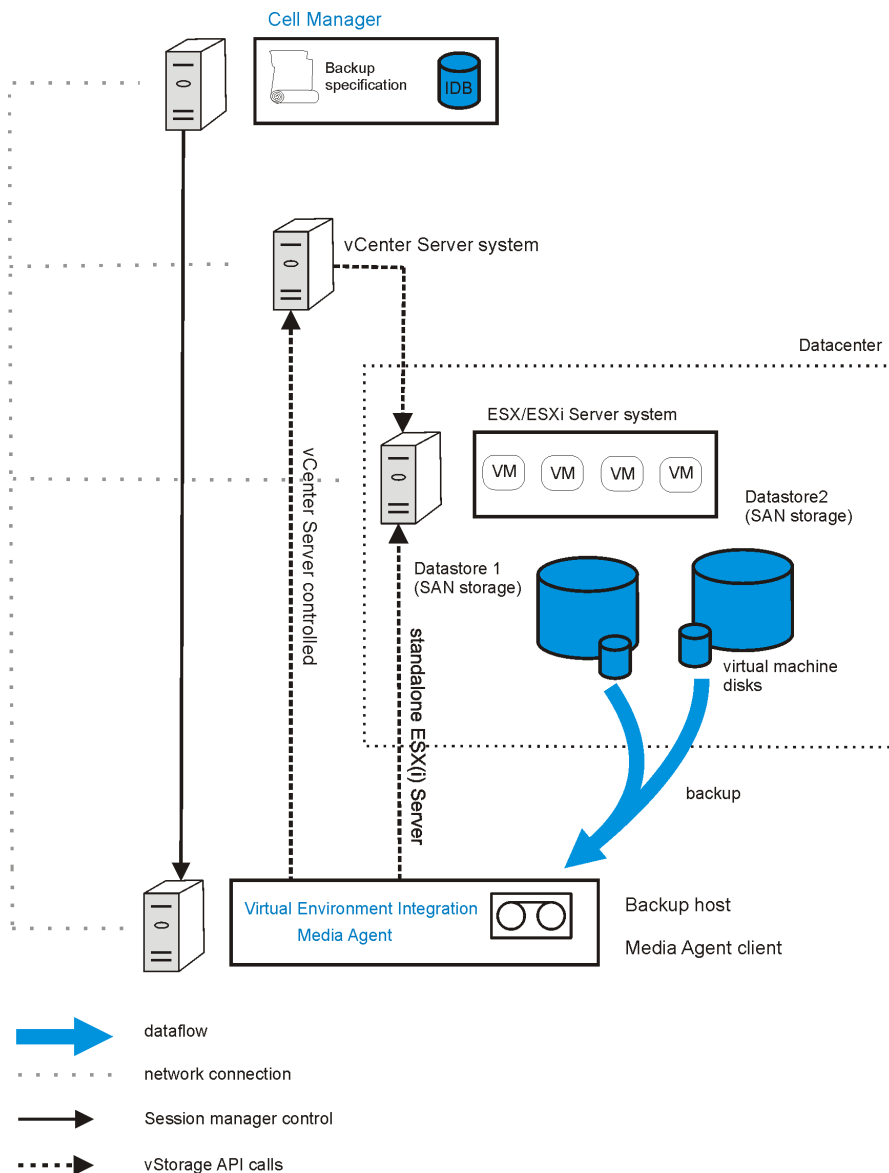
VEAgent then opens the virtual machine disks across LAN or SAN, initializes the Media Agent client and controls the transfer to it of the virtual machine and all its associated data.

---

**NOTE:** You can use the `OB2_VEAGENT_OPEN_DISK_TIMEOUT` omnirc option to specify the time interval between opening two different disks. By default, a new disk is opened every 2 seconds.

---

**Figure 5 vStorage Image method**



In “vStorage Image method” (page 27), the backup host is also a Media Agent client (it has the Media Agent component installed and a device connected to it).

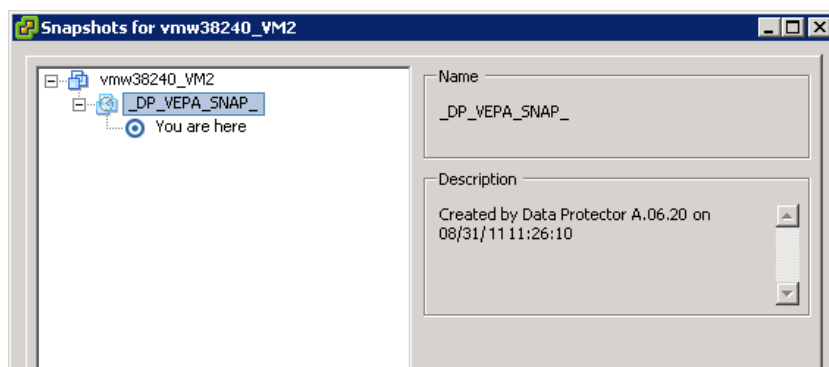
### Snapshot management

The vStorage Image backup method relies on being able to create virtual machine snapshots. A virtual machine snapshot is an operation that puts the virtual machine into a consistent state. All subsequent changes made to the virtual machine disks are recorded to separate files. Note that snapshot operation is not supported by all virtual machine disks. For example, snapshots of independent disks are not supported; consequently, this type of virtual disks cannot be backed up using the Data Protector Virtual Environment integration. For details, see the VMware documentation.

During a vStorage Image backup, Data Protector creates a snapshot and copies data from the consistent state to Data Protector media. Then Data Protector deletes the snapshot. Note that snapshots created by Data Protector (**DP snapshots**) are distinguished from other snapshots by the label `_DP_VEPA_SNAP_` containing the product name, a description, and a timestamp.

**NOTE:** It does not matter if a user-created snapshot exists in a virtual machine snapshot tree at the time of backup; the user-created snapshot is backed up together with other virtual machine files. However, if Data Protector detects a user-created snapshot at the time of restore, which is only possible if the virtual machine to be restored still exists, the virtual machine is not restored. To restore such a virtual machine, you first need to manually remove all existing user-created snapshots. Also note that although the user-created snapshots of the virtual machine were backed up, when the virtual machine is restored, only the latest state is restored; the snapshots are not.

**Figure 6 Virtual machine snapshot tree**



Existing virtual machine snapshots reduce the overall performance of a virtual machine. For this reason, Data Protector automatically removes DP snapshots once they are no longer needed.

❗ **IMPORTANT:** Do not use the label `_DP_VEPA_SNAP_` for snapshots that you create for other purposes, otherwise they will be deleted by Data Protector.

The number of Data Protector snapshots that remain in the snapshot tree after a backup depends on three factors:

- Whether changed block tracking has been specified
- The Data Protector snapshot handling mode selected
- The backup type specified

The snapshots that remain in a snapshot tree play a great role for incremental and differential virtual machine backups. For details, see [“Backup types” \(page 28\)](#).

Since the snapshot handling mode can affect virtual machine performance, it is specified per virtual machine.

## vCD vStorage Image backup method

The vCD vStorage Image backup method provided by the Data Protector Virtual Environment integration is a variant of the vStorage Image backup method. It is used exclusively to back up virtual machines running on VMware ESX(i) Server systems from several vCenter Servers systems managed by VMware vCloud Director. Conceptually, both backup methods are identical.

For details on the vCD vStorage Image backup method concepts, see [“vStorage Image backup method” \(page 26\)](#).

## Backup types

The type of backup to be performed is specified at backup specification level, either in the Scheduler page, or in the Start Backup dialog box for an interactive backup. However, the type of backup that can actually be performed depends on the snapshot handling mode specified and on whether changed block tracking is used or not.

Using the vStorage Image or vCD vStorage Image backup methods, you can perform the following backup types:

**Table 4 Backup types**

Full	Backs up the complete virtual machine (disk), including the virtual machine memory file (if specified).
Incremental	Backs up the changes made to a virtual machine (disk) since the last full, incremental, or differential backup.
Differential	Backs up the changes made to a virtual machine (disk) since the last full backup.

For incremental and differential backup sessions, you must also specify how Data Protector should identify changes:

- At file level
- At disk block level

To identify changes at file level (changed block tracking is not enabled), Data Protector uses the VMware snapshot functionality. For details, see [“Snapshot handling mode” \(page 29\)](#).

To identify changes at disk block level, Data Protector uses the VMware changed block tracking functionality. For details, see [“Changed block tracking” \(page 35\)](#).

The possible combinations of snapshot handling mode and changed block tracking for VMware backups are summarized in [Table 5 \(page 29\)](#).

**Table 5 Snapshot management**

Changed block tracking	VM snapshot handling mode	Number of snapshots remaining after backup
Disabled	Disabled	0
	Single	1
	Mixed	up to 2
Enabled	Disabled	0
	Single	0
	Mixed	0

**NOTE:** When changed block tracking functionality is enabled, the number of snapshots remaining after backup is always 0.

### Snapshot handling mode

The snapshot handling mode enables you to control the number of Data Protector snapshots left in a virtual machine's snapshot tree after a backup. You can set it per virtual machine, or specify a common setting which can be overridden for individual virtual machines. The types of backup that can be performed depend on this setting. The selected snapshot handling mode must therefore be suitable for the type of backup chain anticipated. The following snapshot handling modes are available:

- **Disabled:** Supports only full backups. The snapshot made at the beginning of a backup session is used only to create a consistent state. After the data transfer completes, the snapshot is removed. For details, see [“Snapshot mode: disabled” \(page 30\)](#).
- **Single:** Supports full, differential, and incremental backups in the following backup chains:  
*Full, differential, differential, differential ...*  
*Full, incremental, incremental, incremental ...*

You cannot mix incremental and differential backups within the same backup chain. The snapshot that is made at the beginning of a backup session is used to create a consistent state. After the backup completes, one DP snapshot remains in the snapshot tree. It is needed to

track changes made since the last full or incremental backup. For details, see [“Snapshot mode: single” \(page 31\)](#).

- **Mixed:** Supports full, differential, and incremental backups in all possible backup chain forms. For example:

*Full, incremental, incremental, differential, incremental, differential ...*

The snapshot that is made at the beginning of a backup session is used to create a consistent state. After the backup completes, up to two DP snapshots remain in the snapshot tree. One is needed to track changes made since the last full backup and the other to track changes made since the last backup (incremental or differential). For details, see [“Snapshot mode: mixed” \(page 33\)](#).

When performing snapshot operations on a backed up virtual machine, you must be careful not to break your backup chains.

- 
- ❗ **IMPORTANT:** A VMware object’s snapshot which was not created with Data Protector cannot be used to set up a Data Protector backup chain (restore chain) for that object.

A backup chain gets broken if you perform any of the following operations:

- Delete a snapshot
- Revert to a snapshot
- Create a snapshot without involving Data Protector
- Change snapshot handling mode
- Add a new virtual machine disk or rename an existing one
- Restore the virtual machine
- Enable changed block tracking

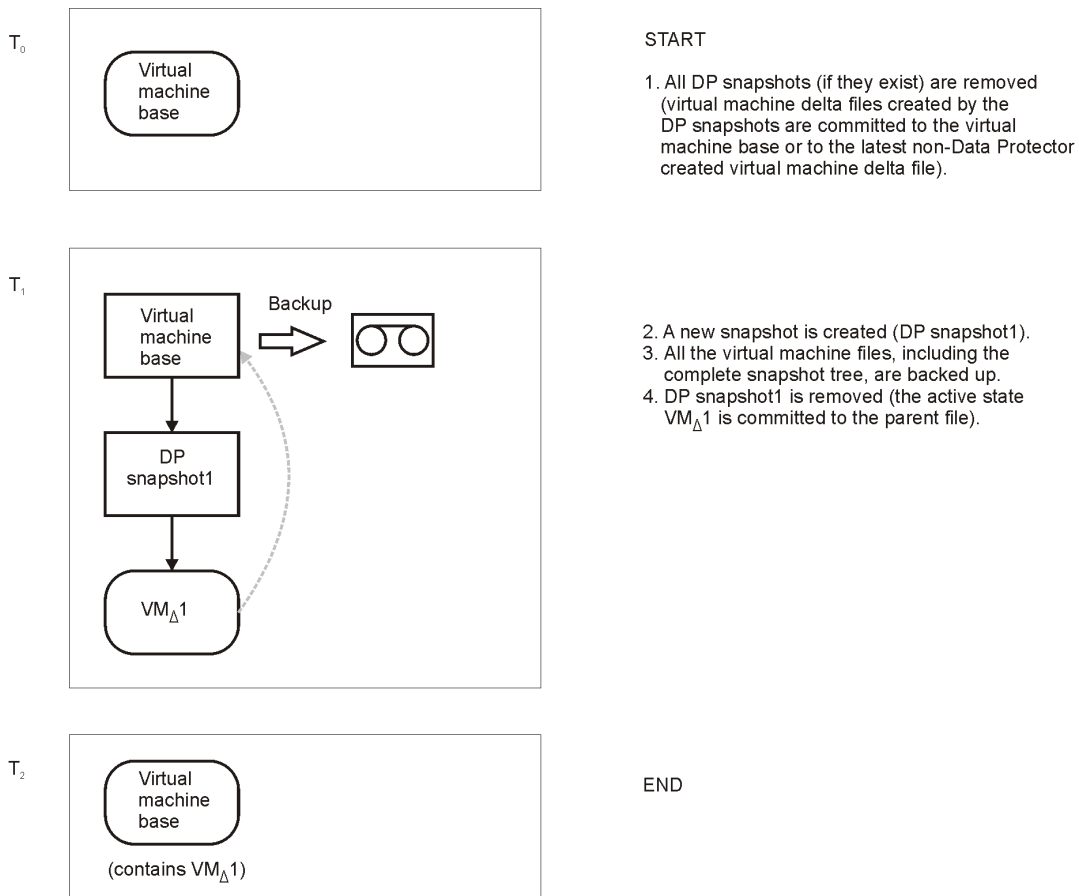
After completing any of the above operations, you must first run a full backup to start a new backup chain. If you run a session for an incremental backup or a differential backup instead, Data Protector switches the effective backup type in the session to full.

---

### **Snapshot mode: disabled**

A full backup in **Disabled** mode progresses as shown in the following figure.

**Figure 7 Full backup (disabled mode)**



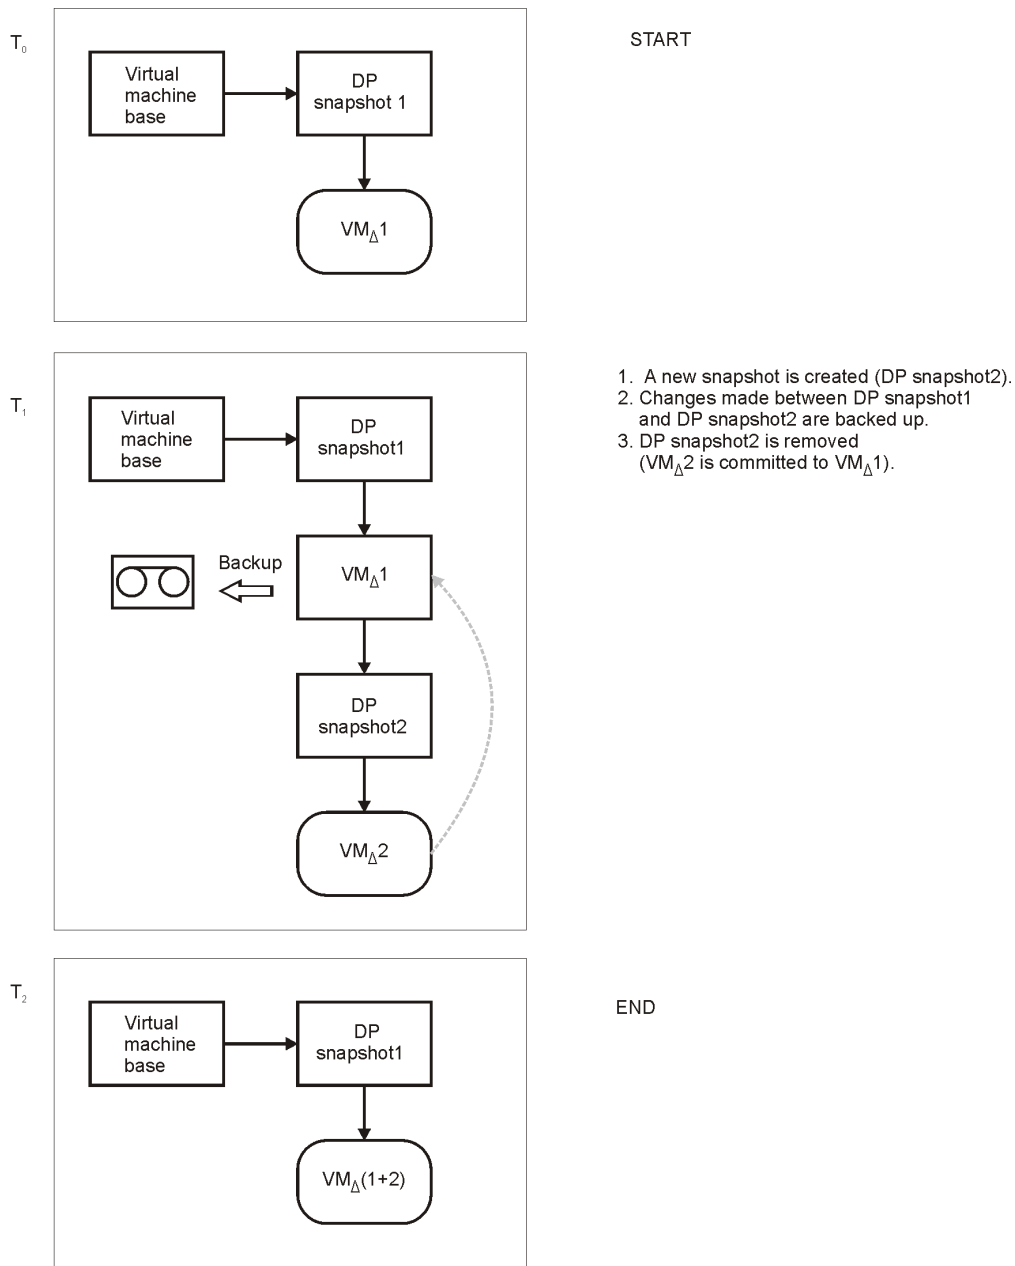
**Table 6 Legend**

T <sub>n</sub>	Boxes denoted by T <sub>i</sub> show how the virtual machine snapshot tree changes in time.
Virtual machine base	The rectangle denoted by Virtual machine base represents the virtual machine base or the last non-Data Protector created virtual machine delta file on the active branch.
VM <sub>Δ</sub>	A rectangle denoted by VM <sub>Δ</sub> represents a virtual machine delta file created by a Data Protector snapshot.
DP snapshot	A rectangle denoted by DP snapshot represents a process (snapshot operation triggered by Data Protector). This process closes the current active state to become a read-only file. At the same time, it creates a new delta file, which becomes the active state. The active state is denoted by round corners. The snapshot created by Data Protector is named <code>_DP_VEPA_SNAP_</code> .

### Snapshot mode: single

A full backup in the **Single** mode progresses in the same way as a full backup in the **Disabled** mode, with the exception that the DP snapshot is not removed at the end (you end up with one DP snapshot). A subsequent differential backup is shown in the following figure.

**Figure 8 Differential backup (single mode)**

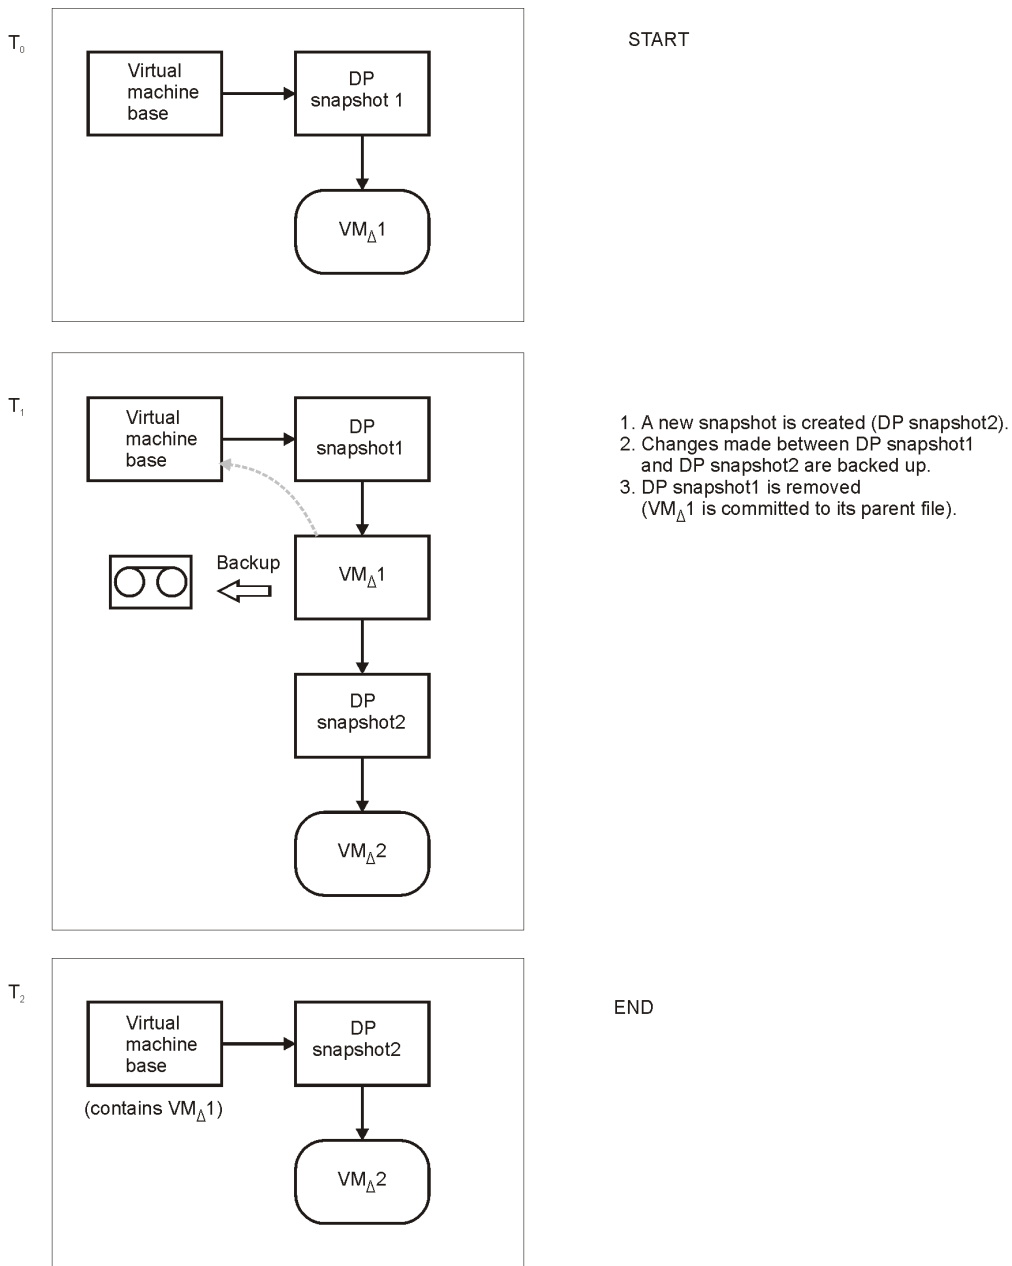


DP snapshot1 remains in the snapshot tree to track changes made since the last full backup.

A backup chain consisting of a full backup that is followed by incremental sessions progresses in the same way, with the exception that, at the end of an incremental session, DP snapshot1 is removed instead of DP Snapshot2 (see [“Incremental backup \(single mode\)”](#) (page 33)).



**Figure 9 Incremental backup (single mode)**



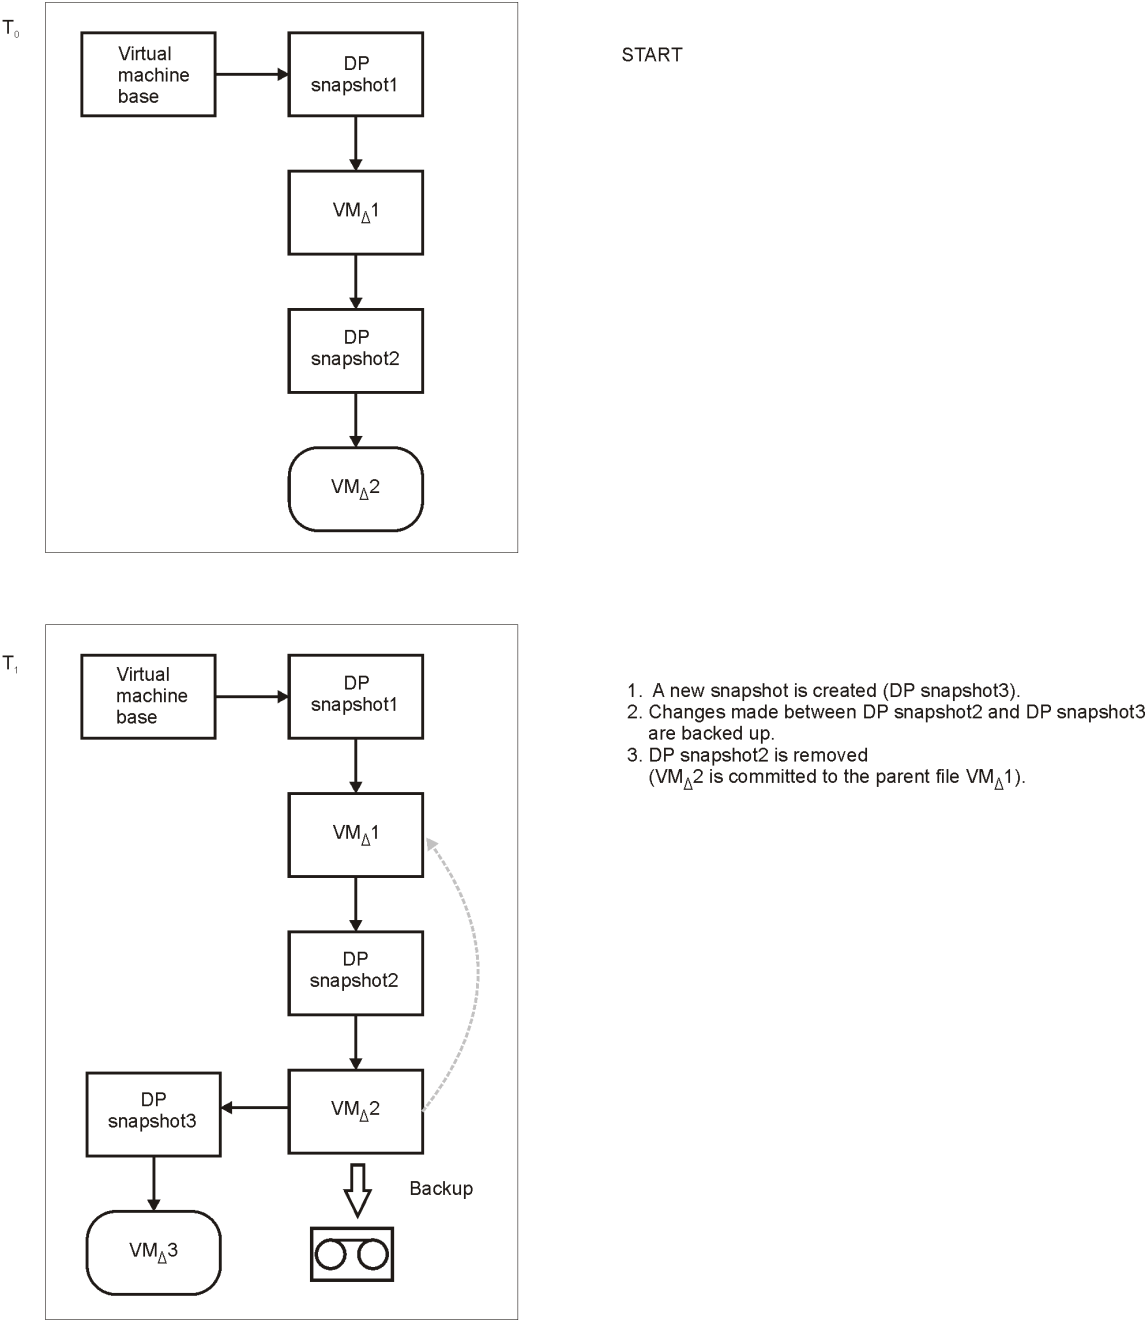
DP snapshot2 remains in the snapshot tree to track changes made since the last incremental backup.

#### Snapshot mode: mixed

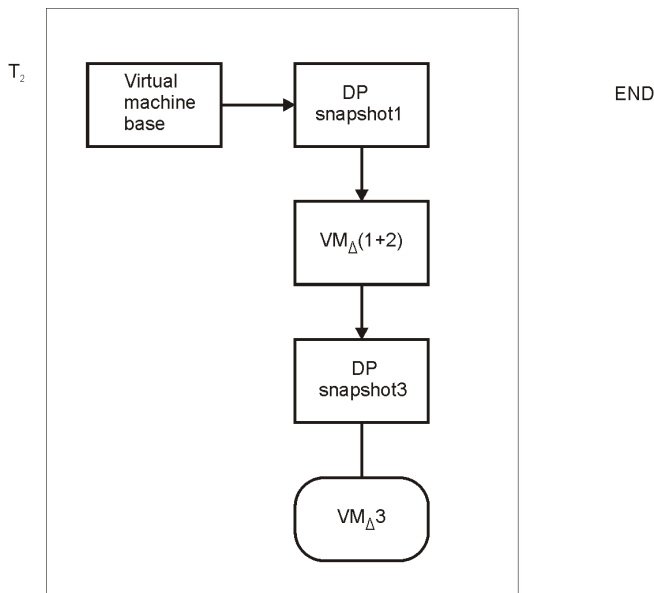
A full backup in **Mixed** mode progresses in the same way as a full backup in **Single** mode (you end up with one DP snapshot). A subsequent differential or incremental backup progresses in the same way as a differential or incremental backup in **Single** mode, with the exception that no DP snapshot is removed at the end (you end up with two DP snapshots).

The progress of a subsequent incremental backup is shown below.

Figure 10 Incremental backup (single mode)



**Figure 11 Incremental backup (mixed mode)**



DP snapshot1 and DP snapshot3 remain in the snapshot tree to track changes made since the last full and the last backup respectively.

A subsequent differential backup progresses in the same way as an incremental backup session described in [“Incremental backup \(mixed mode\)” \(page 35\)](#), with the following exceptions:

- DP snapshot2 is removed before DP snapshot3 is created.
- Instead of changes made between DP snapshot2 and DP snapshot3, changes made between DP snapshot1 and DP snapshot3 are backed up.

### Changed block tracking

Changed block tracking (CBT) is a feature of later versions of VMware that can be used to improve the efficiency and speed of your backups.

For CBT, change IDs are used. A change ID is an identifier for the state of a virtual disk at a specific point in time. It is saved by the virtual disk logic whenever a snapshot is taken of the disk.

The main advantage of using changed block tracking is most noticeable on incremental or differential backups, because:

- It is not necessary to keep virtual machine snapshots on the system until the next backup, greatly reducing the system overhead.
- The changes to be backed up are calculated more easily, by obtaining change information from the kernel, rather than calculating it from snapshots.

During full backups, only active blocks on the disk are backed up, and unallocated blocks are ignored. This makes the backups space-efficient and faster.

When changed block tracking is enabled, a virtual machine's performance is slightly impacted, but this is small relative to what you gain. If it is enabled within VMware vSphere, Data Protector will use it. You can also enable it using the Data Protector GUI, if required.

When changed block tracking is used, Data Protector snapshots are still used to put a virtual machine in a consistent state. However, when the backup completes, they are deleted. Only changed block tracking log files change IDs are kept.

---

**NOTE:**

- Not all types of virtual disk support changed block tracking. In case of an unsupported disk, a virtual machine backup fails.
  - When changed block tracking is first enabled, the next backup for a virtual machine will always be a full backup to provide a reference point for the tracking. In other words, a new backup chain is started.
  - A CBT backup chain (Full, Differential, Incr,...) gets broken when you perform a restore session. When the restore session completes, run a full backup again to start a new backup chain, otherwise subsequent incremental and differential backup sessions will fail.
- 

**Backup flow**

1. Data Protector triggers a snapshot.
2. A change ID is recorded for the current backup.

If this is the first snapshot taken after changed block tracking was enabled, all active blocks are identified and change ID 0 is recorded.

In case of a full backup, this change ID becomes the start reference point for a new backup chain.
3. This step depends on the backup type selected:
  - *Full*: The blocks changed since change ID 0 are identified.
  - *Incremental*: The blocks changed since the change ID for the previous backup (full, incremental, or differential) are identified.
  - *Differential*: The blocks changed since the change ID for the previous full backup are identified.
4. The identified blocks are backed up.
5. The snapshot is deleted.

**Table 7 Example of a backup chain with changed block tracking**

Snapshot	Change ID	Blocks identified	Blocks backed up
1st after CBT enabled	ID 0	All active blocks	—
Full backup	ID $n$	Changed since ID 0	All active blocks from ID 0 + changed blocks since ID 0
Incremental backup	ID $n+m$	Changed since ID $n$	Changed block since ID $n$
Incremental backup	ID $n+p$	Changed since ID $n+m$	Changed blocks since ID $n+m$
Differential backup	ID $n+q$	Changed since ID $n$	Changed blocks since ID $n$
Full backup	ID $r$	Changed since ID 0	All active blocks from ID 0 + Changed blocks since ID 0

**Quiescence**

On Windows guest operating systems, it is possible to use the VSS framework to freeze, or quiesce, the states of the applications running within a virtual machine before a snapshot is created. If quiescence is selected for a Windows guest operating system, the snapshot process quiesces all system writers and registered application writers before taking the snapshot of the virtual machine. During a backup with quiescence enabled, a .zip file is created that contains all the BCD and writer manifests located on the datastores concerned. Data Protector backs up this file. During a

restore, the file is restored to the same location and then post-restore VSS steps like roll-forward can be performed manually.

**NOTE:** Specifying Quiescence can slow down the speed of backup sessions considerably.

## Disk space requirements

A virtual machine backup requires sufficient disk space on the datastore(s) where the virtual machine disks reside.

### Free space required option

You can use the Data Protector **Free space required %** option to make sure that a virtual machine is backed up only if there is enough free space.

The required free space is calculated based on the size of virtual machine disks just before the snapshot is created. Data Protector checks all datastores where the virtual machine disks reside. If one of the datastores does not meet the specified percentage of free space, no snapshot is created and the backup of the virtual machine fails with an error.

When backing up more than one virtual machine, the check is applied to each virtual machine separately. The virtual machines which pass the check are backed up and the ones which do not pass it are not.

If you specify 0%, the check is omitted.

### Examples

The following examples illustrate how the **Free space required %** option works:

1. Backup of a single virtual machine "test1" with the disk "disk1" residing on the datastore "datastore1":  
If you specify 30% in the **Free space required %** option and the datastore has a size of 100 GB, the backup succeeds if there is at least 30 GB of free space on the datastore.
2. Backup of a single virtual machine "test1" with two disks, "disk1" and "disk2" residing on two datastores, "datastore1" and "datastore2":  
If you specify that 30% of free space is required, the backup succeeds if there is at least 30 GB of free space on each datastore.
3. Backup of two virtual machines, the virtual machine "test1" with the disk "disk1" on the datastore "datastore1" and the virtual machine "test2" with the disk "disk2" on the datastore "datastore2":  
If you require 30% of free space, the backup of both virtual machines succeeds if there is at least 30 GB of free space available on each datastore. If for example, the datastore "datastore1" has less than 30% of free space and the datastore "datastore2" has at least 30% of free space, the backup of the virtual machine "test1" fails and the backup of the virtual machine "test2" succeeds. If both datastores have less than 30% of free space, the backup of both virtual machines fails.

**Table 8 Disk space requirements**

Backup methods	Required disk space on datastores	Explanation
vStorage Image and vCD vStorage Image	The sum of the sizes of all the virtual machine disks, plus: <ul style="list-style-type: none"><li>• The size of any quiescence zip files, if quiescence is specified.</li><li>• The size of the virtual machine memory file, if specified.</li><li>• The size of the virtual machine snapshots if they are present (this</li></ul>	When a virtual machine snapshot is taken, changes made to the virtual machine disks are recorded to separate files (one delta file is created for each virtual machine disk). A delta file can grow up to the original virtual disk size.

**Table 8 Disk space requirements** *(continued)*

Backup methods	Required disk space on datastores	Explanation
	applies to Single and Mixed snapshot handling modes with CBT disabled).	

### Backup disk buffer

You can specify a disk buffer for your backup using the omnirc option

OB2\_VEAGENT\_BACKUP\_DISK\_BUFFER\_SIZE.

The SAN and the HotAdd backups support disk buffer sizes from 1 MB to 256 MB. By default, their disk buffer size is 8 MB. However, network backups, such as NBD and NBD (SSL), are always performed using the default disk buffer size of 1 MB.

### NOTE:

- If there is not enough memory available for the specified disk buffer size, a fallback to 1 MB disk buffer size will happen to keep the backup running, and a warning message will be displayed.
- Using bigger disk buffer sizes improves the backup performance, but it also increases the memory consumption. At the certain level the backup performance does not improve anymore due to the limits of your backup host.

### Backup parallelism

By default, virtual machines are backed up in parallel. In rare cases this may lead to problems. For example, backup sessions may end unexpectedly. In such cases, you may set the Data Protector OB2\_VEAGENT\_THREADED\_BACKUP omnirc option on the backup host to 0 to disable parallel backups.

Note that virtual machine disks are backed up sequentially in both cases.

For details on how to use Data Protector omnirc options, see the *HP Data Protector Help* index: "omnirc options".

### Backup considerations

- **Concurrent backup sessions**  
Backup sessions that use the same devices cannot run in parallel.  
You cannot back up virtual machines and an ESX(i) Server systems or a datacenters where the virtual machines reside in parallel.
- **Transportation modes**  
You can use various transportation modes for backup. For details, see ["Advanced virtual machine settings"](#) (page 53).
- **Thick and thin disks**  
Data Protector cannot detect whether virtual machine disks are thick or thin. In both cases, complete disks are backed up (that is, the complete space allocated at a disk's creation time is backed up, even if the space is still empty). To create space-efficient backups, enable changed block tracking. Note that not all datastores support changed block tracking.
- **Migration and vCD vStorage Image backup**  
Migration of virtual machines or vApps through the VMware vCloud Director during the vCD vStorage Image backup is not supported.

## Restore concepts

You can restore VMware objects backed up with either of vStorage Image backup methods in different ways.

### Restore of VMware objects backed up with vStorage Image method

Virtual machines, virtual machine disks, and virtual machines templates backed up with the **vStorage Image** method can be restored:

- To a datacenter
- To a directory on a backup host

#### Restore to a datacenter

By default, virtual machines are restored to the original datacenter and the original datastore, but you can select a different datacenter if you want.

By default, Data Protector deletes a virtual machine (if it exists) before it is restored, even if it resides in a different datacenter from the datacenter you restore to.

---

**NOTE:** If you select an ESX(i) Server client in the Restore client option (destination client) in the restore wizard, the migrated virtual machine will not be deleted, as the ESX(i) Server client is not able to detect virtual machines that are located on different ESX(i) Server clients (only a vCenter client can do that). Consequently, you end up with two virtual machines having the same UUID.

---

Alternatively, you can choose to restore virtual machines only if they do not exist, leaving existing virtual machines intact.

For the restore, you can also specify the following:

- Whether the memory state of the virtual machine should be restored (if a memory file was included in the backup)
- Whether the restored virtual machines should be registered in the datacenter
- Whether the restored virtual machine snapshots should be consolidated when the restore completes
- Whether the restored virtual machines should be powered on

The restore options provided are, by default, set to restore virtual machines to the original datacenter.

#### Restore of individual virtual machine disks

To be able to restore individual virtual machine disks to a datacenter, the original virtual machine must still exist. Otherwise, the restore fails.

Here is the progress of a restore session:

1. The virtual machine is powered off.
2. If the disks to be restored still exist, they are removed.
3. The disks are restored from the backup.

---

**NOTE:** After restore, virtual disks that are part of a dynamic disk set or virtual disks from different points in time may require additional user action (for example, mounting, resignaturing, or recovery) in the guest operating system and/or applications running inside.

---

#### Restore to a directory

When restoring to a directory (restore outside a datacenter) all the files of virtual machines are restored to a directory of your choice (for example, `C:\tmp`) on a backup host.

In the directory you specified, subdirectories are created with names corresponding to those of the datastores where the virtual machines (their virtual disks) resided at the time of backup. The files related to the virtual disks are restored to the respective subdirectories.

After such a restore, the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX(i) Server system, using the VMware Converter as described in [“Recovering virtual machines after restore to a directory”](#) (page 78).

## Restore of VMware objects backed up with vCD vStorage Image method

VMware vCloud Director vApps and VMs backed up with the **vCD vStorage Image** method can be restored:

- To an organization
- To a directory on a backup host

### Restore to an organization

By default, VMs are restored to the original organization, and consequently to the original vDatacenter, vApp, and vCenter, but you can select a different organization, vDatacenter, vApp, and vCenter if you want.

For the restore, you can also specify the following:

- Whether the memory state of the virtual machine should be restored (if a memory file was included in the backup)
- Whether the restored virtual machine snapshots should be consolidated when the restore completes
- Whether the restored virtual machines should be powered on

The restore options provided are, by default, set to restore virtual machines to the original destination.

### Restore to a directory

When restoring to a directory (restore outside a datacenter or an organization) all the files of virtual machines are restored to a directory of your choice (for example, `C:\tmp`) on a backup host.

After such a restore, the virtual machines are not functional. You need to manually move them to an organization as described in [“Recovering virtual machines after restore to an organization”](#) (page 83).

## Restore chain

When you restore a virtual machine from a backup created in an incremental or differential session, Data Protector automatically restores the complete backup chain, starting with the last full backup, which is then followed by the last differential and all subsequent incremental backups (if they exist) up to the selected session.

## Restore considerations

- **Concurrent sessions**

Restore sessions that use the same devices cannot run concurrently.

- **Failed restore sessions**

Sometimes, when a virtual machine restore fails, Data Protector creates extra files on the datastore which you need to clean up manually when the session completes. Otherwise, corrupt virtual machine backups may be created in subsequent sessions and restore from a such a backup also fails. For details, see [“Cleaning up a datastore after a failed restore”](#) (page 84).

When restoring a virtual machine to a non-original datastore whose block size is not compatible with the virtual machine disks' sizes (that is, a `.vmdk` file size is not a multiple of the datastore block size), the restore fails.



- **Virtual machines in vApp**

When you restore a virtual machine that resided in a vApp container at the time of backup, the virtual machine is not restored back to the vApp container, but to the ESX(i) Server root level. If the virtual machine in the vApp container still exists, it is deleted or the restore is skipped, depending on what you select in the Existing virtual machine handling option.

- **Partial restore from a vStorage Image backup**

When performing a partial restore from a vStorage Image backup (for example, when restoring only some out of many backed up VM disks), the default option **Delete after restore** is ignored and the **Delete before restore** option is used instead.

- **Transportation modes**

The following recommendations for specific virtual machine transportation modes apply:

- *SAN transportation mode*: To use the SAN transportation mode for restore:
  - Select a physical backup host for a restore session.
  - Ensure that the storage volumes that are presented to both the backup host and ESX(i) Server systems are not read-only. For details on how to check storage volume properties, see [“A restore session using SAN transportation mode fails” \(page 86\)](#).
  - Ensure that the storage volume size is a multiple of the underlying VMFS block size. Otherwise, the Write operation to the remainder fails. For example, if the storage volume size is 16.3 MB and the block size 1 MB, writing to the remaining 0.3 MB fails. For details, see the VMware Knowledge Base at:  
<http://kb.vmware.com/selfservice/microsites/searchEntry.do>.  
Search for “Best practices when using SAN transport for backup and restore”.

The SAN transportation mode automatically switches to NBD if:

- The backup host is a virtual machine.
- You select an incremental or differential backup session for restore.
- *Hotadd transportation mode*: Hotadd transportation mode is available for restore, however VMware does not support multiple disks. Therefore, in a HotAdd environment use the omnirc option `OB2_VEAGENT_RESTORE_TRANSPORT_METHOD` to set the restore transport mode to NBD.

## Performance

You can improve backup and restore performance by setting the device block size to the maximum value of 1024 kB.

## Configuring the integration

Configure the integration as follows:

- Import VMware clients into the Data Protector Cell.
- Configure virtual machines you want to back up.

## Prerequisites

- Make sure that you have a correctly installed and configured VMware vSphere or VMware vCloud Director environment.  
For supported versions, platforms, devices, and other information, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.
- Make sure you have the necessary VMware vSphere privileges granted to the user account that is used to connect to the vCenter Server. For details, see [Step 4](#).
- You have to be either a member of the VMware vCloud Director System Administrator user group or have granted the system administration user rights to perform backup and restore operations for VMware vCloud Director client.
- Make sure that you have correctly installed Data Protector. For information on installing Data Protector in various architectures, see the *HP Data Protector Installation and Licensing Guide*.  
Make sure that you have at least one client with the Virtual Environment Integration component installed (**backup host**) in your environment. No special configuration is required for the backup host after installation.

- 
- ① **IMPORTANT:** The client you intend to use as the backup host should not have the VMware Consolidated Backup (VCB) software installed.
- 

If you intend to restore virtual machine files to a directory on a backup host, also install the Disk Agent component on the backup host. Otherwise, you will not be able to use the **Browse** button to specify the target directory (however, you will still be able to type the directory by yourself).

## Before you begin

- Configure devices and media for use with Data Protector.

## Importing and configuring VMware clients

With the Data Protector Virtual Environment integration, it is not necessary to install any Data Protector components on the VMware clients (vCenter Server systems, ESX(i) Server systems, vCloud Director). To make them Data Protector clients, the VMware clients must be properly imported into the Data Protector cell and configured.

- 
- ① **IMPORTANT:** A client cannot be imported and configured as VMware vCenter client, VMware vCloud Director client, and Hyper-V client at the same time.
- 

### Procedure

To import a client into a Data Protector cell:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell**, right-click **Clients**, and select **Import Client**.
3. In the Import client page, enter the client name in the **Name** option, select the appropriate client type (**VMware ESX(i)**, **VMware vCenter**, or **VMware vCloud Director**) from the **Type** drop-down list and click **Next**.

Figure 12 Importing a VMware vCenter Server client (Name and Type)

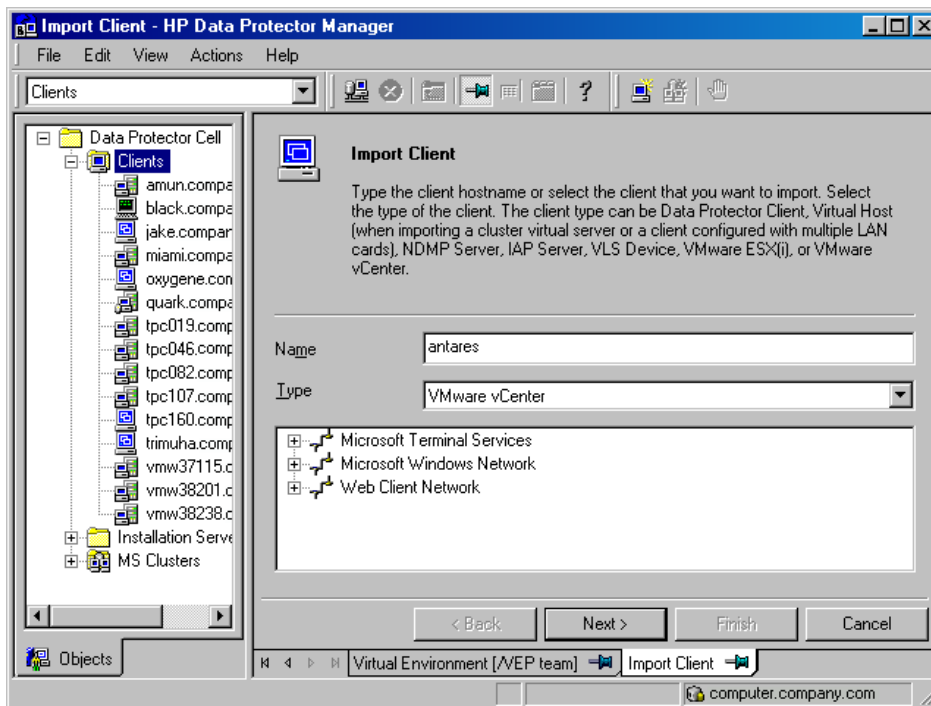
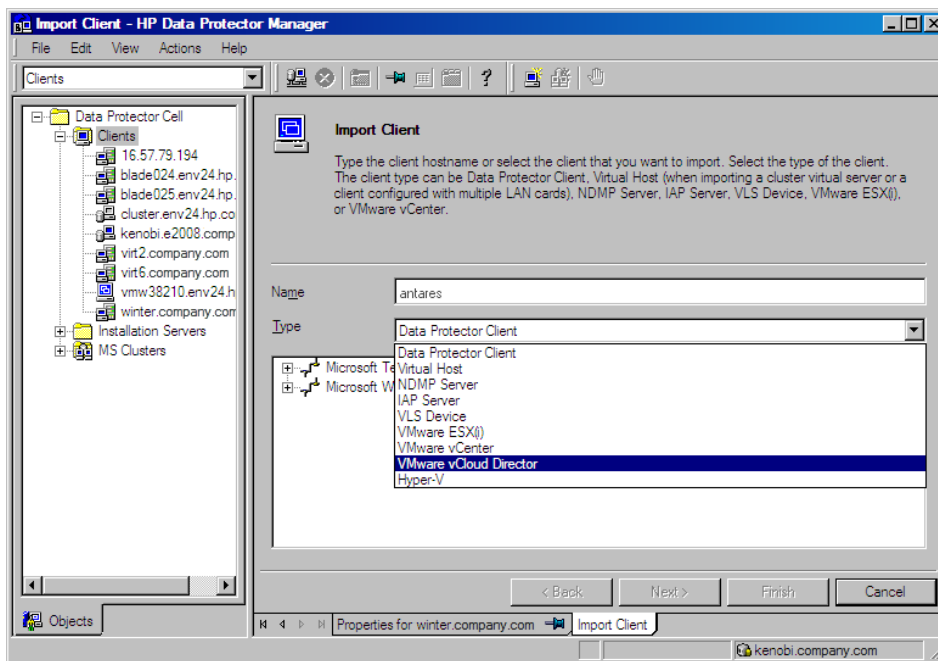


Figure 13 Importing a VMware vCloud Director client (Name and Type)



4. If you select **Standard security**, you need to manually specify which login credentials Data Protector should use to connect to the VMware client:

**Port:** Specify the port that VMware vSphere is using. By default, VMware uses the port 443.

**Username and Password:** Specify an operating system user account that has the following VMware vSphere privileges:

Datastore -> Allocate space  
Datastore -> Browse datastore  
Datastore -> Low level file operations  
Datastore -> Remove file  
Datastore -> Rename datastore  
Folder -> Delete folder  
Folder -> Rename folder  
Global -> Disable methods  
Global -> Enable methods  
Global -> Licenses  
Host -> Configuration -> Maintenance  
Host -> Inventory -> Add standalone host  
Network -> Assign network  
Resource -> Assign virtual machine to resource pool  
Resource -> Remove resource pool  
Resource -> Rename resource pool  
Sessions -> Validate session  
vApp -> Delete  
vApp -> Rename  
vApp -> Add virtual machine  
Virtual machine -> State -> Revert to snapshot  
Virtual machine -> Configuration \*  
Virtual machine -> Interaction -> Answer question  
Virtual machine -> Interaction -> Power Off  
Virtual machine -> Interaction -> Power On  
Virtual machine -> Inventory -> Create new  
Virtual machine -> Inventory -> Register  
Virtual machine -> Inventory -> Remove  
Virtual machine -> Inventory -> Unregister  
Virtual machine -> Provisioning \*  
Virtual machine -> State -> Create snapshot  
Virtual machine -> State -> Remove snapshot

- **Web service:** Optionally, change the web service entry point URI. Default: /sdk

If you select **Integrated security**, which is only available for VMware vCenter Server systems, provided that both the application client and the backup host are Windows systems, Data Protector connects to the VMware vCenter Server system with the user account under which the Data Protector Inet service on the backup host is running. Ensure that this user account has appropriate VMware vSphere rights to connect to the VMware vCenter Server system and the Data Protector Inet service on the backup host is configured for user impersonation.

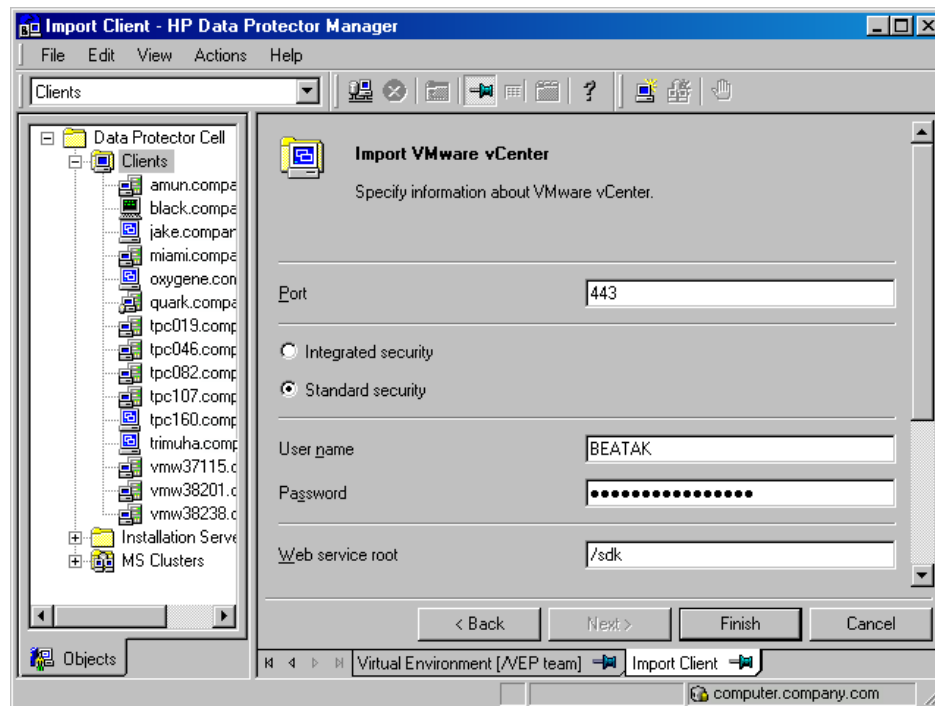
For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".

For the **Port** and **Web service root** options, Data Protector uses the values that are currently specified for the standard security. Integrated security is based on Security Support Provider Interface (SSPI).

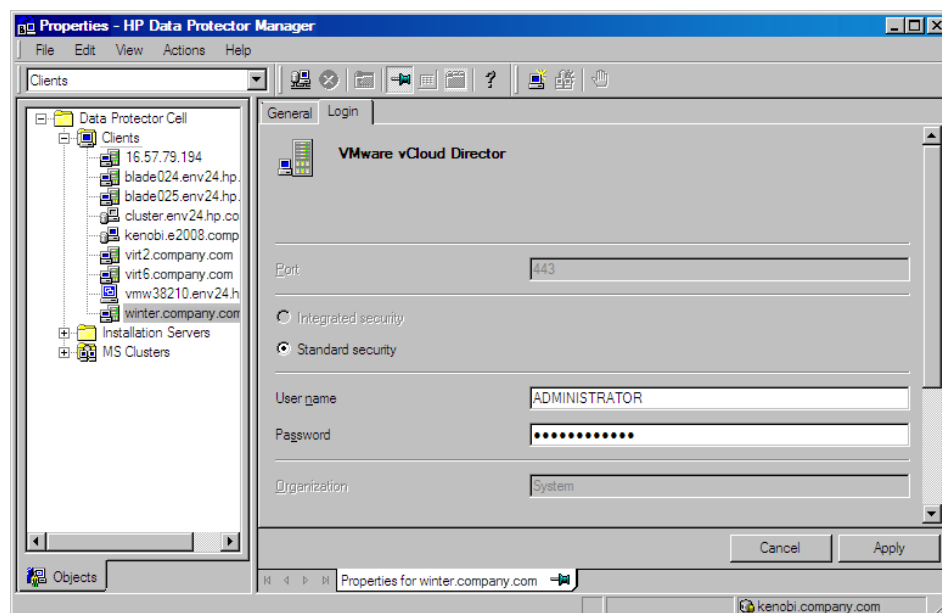
**Organization:** The organization in the VMware vCloud Director that you connect to. This option cannot be modified. By default, the *System* organization, which is only accessible to VMware vCloud Director system administrators, is used.

Click **Finish**.

**Figure 14 Importing a VMware vCenter Server client (login credentials)**



**Figure 15 Importing a VMware vCloud Director client (login credentials)**



**NOTE:** For details on how to change or check the parameters later, see “Changing the configuration of VMware clients” (page 46) and “Checking the configuration of VMware clients” (page 48).

## Changing the configuration of VMware clients

When you update the credentials for connecting to a VMware client (vCenter Server, ESX(i) Server, or vCloud Director client), you actually update the `cell_info` file which resides on the Data Protector Cell Manager. Therefore, you can only change the login credentials if you have the Data Protector `Clients configuration` user right. For details on the Data Protector user rights, see the *HP Data Protector Help* index: "user groups".

To update the credentials, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

You can update the credentials in two different places: in the Clients or in the Backup context.

#### Clients context

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand Clients, and then select the client for which you want to change the login credentials.
3. In the Results Area, click the **Login** tab.
4. Update the credentials and click **Apply**.

#### Backup context

It is assumed that a backup specification for the VMware client for which you want to change the login credentials already exists.

1. In the Context List, click **Backup**.
2. Open a backup specification for the VMware client for which you want to change the login credentials.
3. In the Source page, right-click the client at the top, and select **Configure**.

4. In the Configure Virtual Environment dialog box, update the values and click **OK**.

**Figure 16 Changing the configuration of a VMware vCenter Server or VMware ESX(i) Server client**

The screenshot shows the 'Configure Virtual Environment' dialog box with the 'Connection' tab selected. The title bar reads 'Configure Virtual Environment'. Below the tab, there is a folder icon and the text 'Virtual Environment configuration' and 'Specify security information for the connection.' The 'Client:' field is set to 'amun.company.com'. There are two radio buttons: 'Integrated security' (unselected) and 'Standard security' (selected). Below these are five text input fields: 'Username' with 'administrator1', 'Password' with masked characters, 'Web service root' with '/sdk', and 'Port' with '443'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

**Figure 17 Changing the configuration of a VMware vCloud Director client**

The screenshot shows the 'Configure Virtual Environment' dialog box with the 'Connection' tab selected. The title bar reads 'Configure Virtual Environment'. Below the tab, there is a folder icon and the text 'Virtual Environment configuration' and 'Specify security information for the connection.' The 'Client:' field is set to 'winter.company.com'. There are two radio buttons: 'Integrated security' (unselected) and 'Standard security' (selected). Below these are four text input fields: 'Username' with 'administrator', 'Password' with masked characters, 'Organization' (empty), and an unlabeled empty field. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

## Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.

## 2. Execute:

### For **Integrated security**:

```
vepa_util.exe
command
--config
--virtual-environment vmware
--host VMwareClient
--security-model 1
```

### For **Standard security**:

VMware vCenter Server or VMware ESX(i) Server client

```
vepa_util.exe
command
--config
--virtual-environment vmware
--host VMwareClient
--security-model 0
--username Username
{--password Password | --encoded-password Password}
[--webroot WebServiceRoot]
  [--port WebServicePort]
```

VMware vCloud Director client

```
vepa_util.exe
command
--config
--virtual-environment vCD
--host vCDClient
--security-model 0
--username Username
{--password Password | --encoded-password Password}
```

The message \*RETVAL\*0 indicates successful configuration.

For descriptions of the options, see the `vepa_util.exe` man page or the *HP Data Protector Command Line Interface Reference*.

## Checking the configuration of VMware clients

During the configuration check, Data Protector tries to connect to a VMware client using the login credentials from the `cell_info` file on the Data Protector Cell Manager.

To verify the connection, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

You can verify the connection to a VMware client after you have created at least one backup specification for this client.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Virtual Environment**. Click the backup specification for the VMware client to be checked.
3. In the Source page, right-click the VMware client and select **Check configuration**.

### Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.



## 2. Execute:

VMware vCenter Server or VMware ESX(i) Server client

```
vepa_util.exe  
command  
--check-config  
--virtual-environment vmware  
--host VMwareClient
```

VMware vCloud Director client

```
vepa_util.exe  
command  
--check-config  
--virtual-environment vCD  
--host vCDClient
```

The message `*RETVAL*0` indicates successful configuration.

For option description, see the `vepa_util.exe` man page or the *HP Data Protector Command Line Interface Reference*.

## Configuring virtual machines

To configure a virtual machine means to specify how the virtual machine should be backed up.

You can specify the following:

- Whether a virtual machine's disks should be defragmented and shrunk before they are backed up.
- (Windows virtual machines only) Whether a quiescence snapshot should be taken to make applications running inside a virtual machine consistent for backup.
- Which transportation mode should be used during backups.
- Which VMware functionality should be used to detect changes for incremental and differential virtual machine backups: the VMware snapshot functionality or the changed block tracking functionality (**CBT**).

For each datacenter, you can specify:

- *Common settings* that apply to all virtual machines in the datacenter.
- *Virtual machine-specific settings* that override the common settings. If there are no virtual machine-specific settings, the common settings are used for that particular virtual machine.

All these settings are saved in a datacenter-specific file `VMwareClient%DatacenterPath` on the Cell Manager. The file is used for all backup sessions that use any of the backup specifications for this datacenter.

Similarly, backup sessions using any of the backup specifications for *All datacenters*, use the settings from the file `VMwareClient%AllDatacenters`.

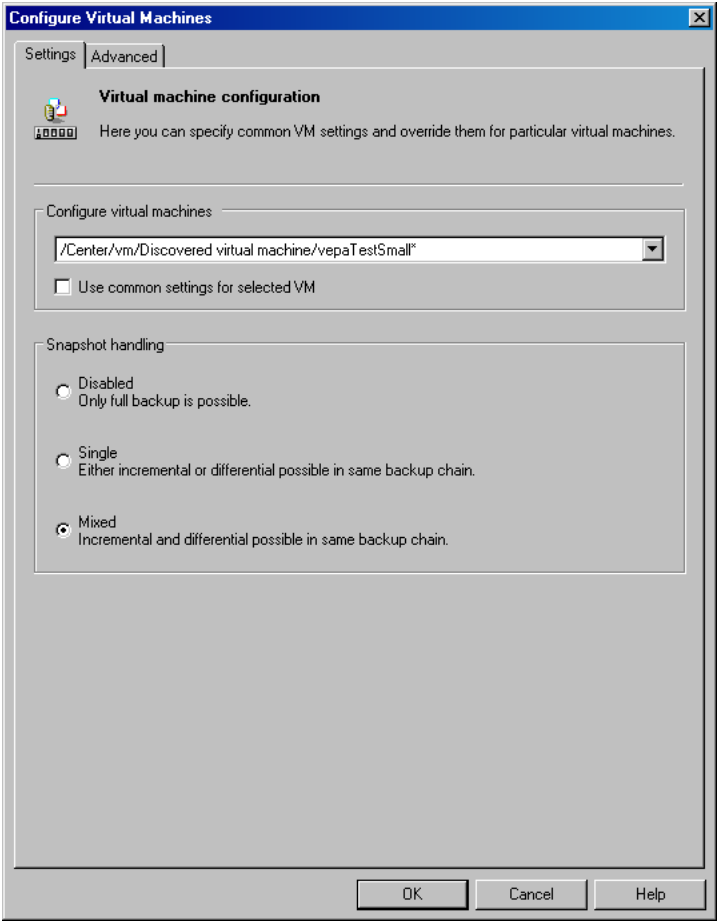
The files `VMwareClient%DatacenterPath` and `VMwareClient%AllDatacenters` are created or updated when you create or update a backup specification for a specific datacenter or all datacenters respectively.

To configure virtual machines, use the Data Protector GUI or CLI.

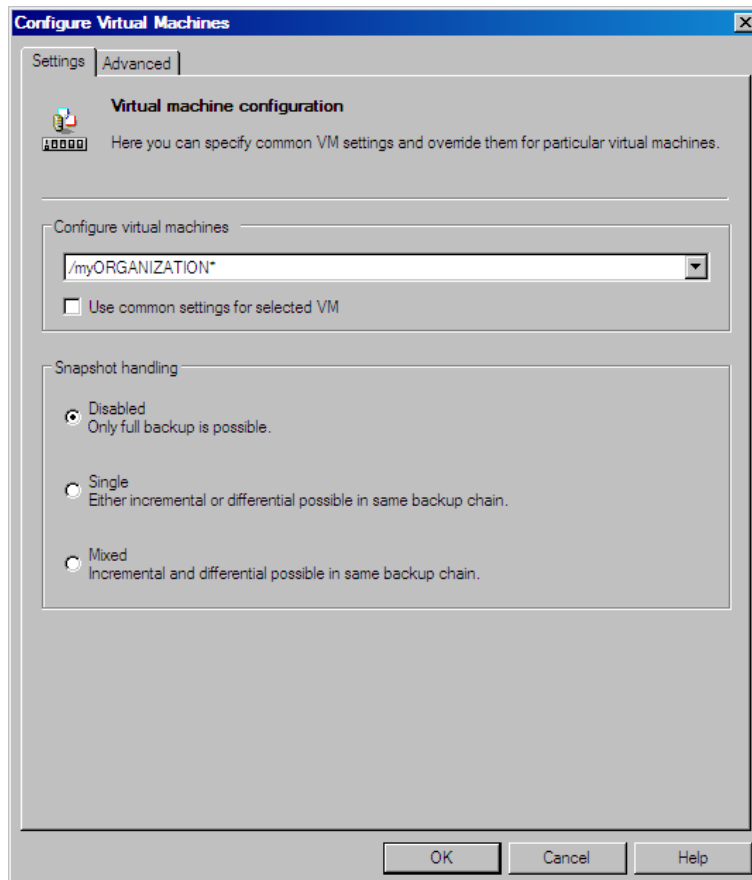
## Using the Data Protector GUI

You can configure virtual machines when you create or modify a backup specification. In the Source page of a backup specification, right-click the client system at the top or any of the virtual machines listed below, and select **Configure Virtual Machines**.

Figure 18 Configure virtual machines (settings for VMware vCenter Server client)



**Figure 19 Configure virtual machines (settings for VMware vCloud Director client)**



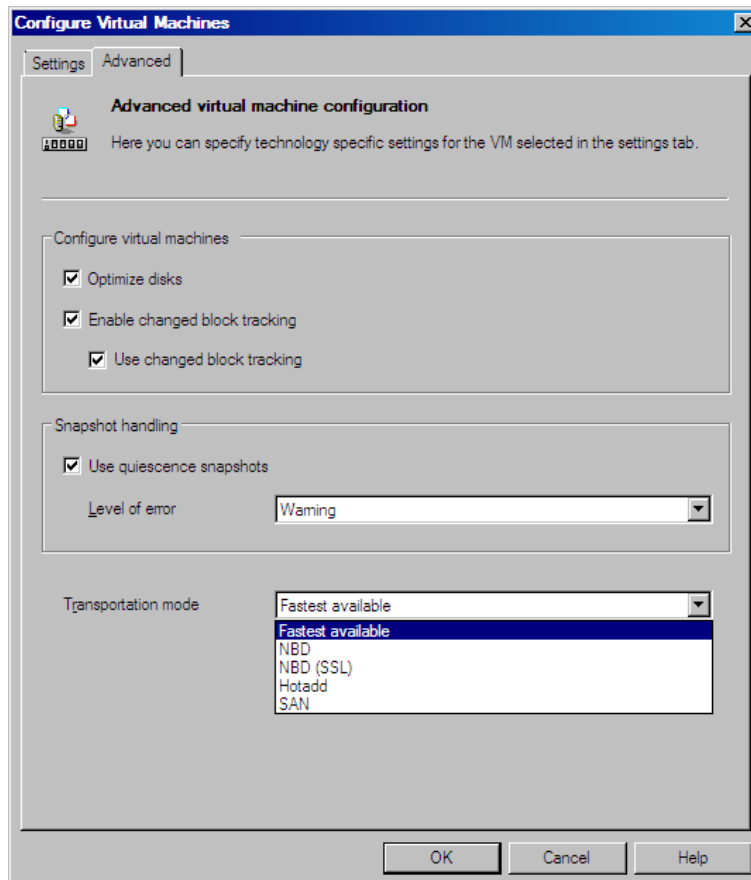
In the Configure Virtual Machines dialog box **Settings** page, specify the following settings:

**Table 9 Virtual machine settings**

<b>Configure virtual machine</b>	Select whether you want to specify common virtual machine settings ( <b>Common VM Settings</b> ) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.	
	<b>Use common settings for selected VM</b>	Available only if a virtual machine is selected. Select this option if you want the common settings to apply to the selected virtual machine. Default: selected
	<b>Use default settings</b>	Available only if <b>Common VM Settings</b> is selected. Select this option to set default values for the common virtual machine settings. Default: selected
<b>Snapshot handling</b>	<b>Disabled</b> (default)	This mode supports only full backups. The virtual machine snapshot that is created during backup is removed at the end of the session. For details, see <a href="#">“Snapshot mode: disabled” (page 30)</a> .
	<b>Single</b>	This mode supports full, incremental, and differential backups. However, you cannot mix incremental and differential backups within the same backup chain.  Data Protector keeps one DP snapshot for backup purposes. For details, see <a href="#">“Snapshot mode: single” (page 31)</a> .
	<b>Mixed</b>	This mode supports full, incremental, and differential backups, in all combinations.  Data Protector keeps up to two DP snapshots for backup purposes. For details, see <a href="#">“Snapshot mode: mixed” (page 33)</a> .

In the Configure Virtual Machines dialog box **Advanced** page, specify the following settings:

**Figure 20 Configure virtual machines (Advanced)**



**Table 10 Advanced virtual machine settings**

Configure virtual machine	Specify advanced backup options to be used for the virtual machine backup. The options specified in this page apply to the virtual machines that are currently selected in the Settings page.	
	<b>Optimize disks</b>	<p>This option is applicable only for thin virtual machine disks.</p> <p>Select this option to defragment and shrink virtual machine disks (.vmdk files) before they are backed up. Shrinking a virtual machine disk reclaims unused space and so reduces the amount of space the disk occupies on the ESX(i) Server system drive. Consequently, this reduces the size of backup data. However, note that such a backup needs more time to complete.</p> <p>Note that defragmentation results in disk block changes (even if there are no changes made to a virtual machine). Consider this if you are using changed block tracking for incremental and differential backups.</p> <p>Default: not selected.</p>
	<b>Enable changed block tracking</b>	<p>Select this option to enable the VMware changed block tracking functionality for the selected virtual machines.</p> <p>If the functionality has already been enabled using the VMware vSphere client, you cannot disable it by clearing this option.</p> <p>Default: not selected.</p>
	<b>Use changed block tracking</b>	<p>Available only if <b>Enable changed block tracking</b> is selected.</p> <p>Select this option to use the VMware changed block tracking functionality for incremental and differential backup sessions. For details, see <a href="#">“Changed block tracking”</a> (page 35).</p> <p>Not all datastores support changed block tracking. If this option is selected and the datastore does not support this functionality, incremental and differential backup sessions will fail.</p> <p>Default: not selected.</p>

**Table 10 Advanced virtual machine settings** *(continued)*

<b>Snapshot handling</b>	<b>Use quiescence snapshots</b>	<p>Applicable for Windows virtual machines.</p> <p>Select this option to use Microsoft Volume Shadow Copy Service (VSS) functionality to quiesce all applications with VSS writers before performing the backup. This produces application-consistent backups.</p> <p>Do not select this option if one or more virtual machines selected in the backup specification have the <b>Backup memory file</b> specified.</p> <p>Default: not selected. For details, see <a href="#">“Quiescence” (page 36)</a>.</p>
	<b>Level of error</b>	<p>Available only if <b>Use quiescence snapshots</b> is selected.</p> <p>Specify the level of error to be reported if a quiescence snapshot fails.</p>
<b>Transportation mode</b>		<p>Select the transportation mode to be used when backing up a virtual machine.</p> <ul style="list-style-type: none"> <li>• <b>NBD</b>: Use this mode when your ESX(i) Server systems do not have access to a SAN, but use local storage devices or NAS to store virtual machine disks. This is an unencrypted transportation mode over a local area network that uses the Network Block Device (NBD) driver protocol. This transportation mode is usually slower than Fibre Channel.</li> <li>• <b>NBD (SSL)</b>: Same as NBD except that the communication over the network is encrypted using the Secure Socket Layer (SSL) cryptographic protocol.</li> <li>• <b>Hotadd</b>: Use this mode if your backup host (a client with the Data Protector Virtual Environment Integration component installed) is a virtual machine. Such a configuration enables you to back up other virtual machines residing on datastores visible to the ESX(i) Server that hosts the backup host.</li> <li>• <b>SAN</b>: Use this mode when your ESX(i) Server systems store their virtual machine disks on Fibre Channel SAN or iSCSI SAN. This is an unencrypted transportation mode over Fibre Channel or iSCSI.</li> </ul> <p>For details, see the VMware documentation.</p> <p>This transportation mode requires that the storage volumes on which the virtual machines are located are presented to the client with the <code>Virtual Environment Integration</code> component installed (<b>backup host</b>).</p> <p><b>Caution:</b> Do not reformat these storage volumes. Otherwise, you will delete all your virtual machines.</p> <p>For details on these VMware transportation modes, see the VMware documentation.</p> <p>If you are not concerned which mode is used, select <b>Fastest available</b>.</p> <p>Default: <b>Fastest available</b></p>

## Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.

## 2. Execute:

```
vepa_util.exe
command
--configvm
--virtual-environment { vmware | vCD }
--host AppHostName
--instance DatacenterPath
--vm VMpath VM_OPTIONS

VM_OPTIONS
--snapshots { 0 | 1 | 2 }
--transportation-mode { san | nbd | nbdssl | hotadd | fastest }
--enableCt { 0 | 1 }
--useCt { 0 | 1 }
--quiescence { 0 | 1 }
--quiescenceErrLvl { 0 | 1 }
--optimize-disks { 0 | 1 }
--uuid UUID_of_VM
```

The values { 0 | 1 | 2 } represent the **Disabled**, **Single**, and **Mixed** snapshot handling modes respectively. For details, see the `vepa_util.exe` man page or the *HP Data Protector Command Line Interface Reference*.

To change the virtual machine specific settings back to the common virtual machine settings, execute:

```
vepa_util.exe
command
--configvm
--virtual-environment { vmware | vCD }
--host AppHostName
--instance DatacenterPath
--vm VMpath
--uuid UUID_of_VM
--default
```

The message `*RETVAL*0` indicates successful configuration.

### Example

To configure the virtual machine with the virtual machine path `/MyDatacenter/MyVM` and with the UUID `42375365-ebe1-e9da-7068-7beb727cab19` that resides in the datacenter `/MyDatacenter` and which is registered in the vCenter Server system `vc.company.com`, with the following settings:

- The snapshot handling mode is mixed.
- Before virtual machine disks are backed up, they are optimized.
- Quiescence snapshots are used.
- If a quiescence snapshot fails, the level of error is a warning.
- CBT is enabled and used.
- The transportation mode is fastest available.

execute:

```
vepa_util.exe
command
--configvm
--virtual-environment vmware
--host vc.company.com
--instance /MyDatacenter
--vm /MyDatacenter/MyVM
--snapshots 2
--enableCt 1
```

```
--useCt 1
--quiescence 1
--quiescenceErrLvl 0
--transportation-mode fastest
--optimize-disks 1
--uuid 42375365-ebe1-e9da-7068-7beb727cab19
```

## Customizing the Data Protector behavior with omnirc options

The omnirc options are useful for troubleshooting or overriding other settings affecting the behavior of the Data Protector client. The options that apply to the Virtual Environment integration have the prefix `OB2_VEAGENT`.

For details on how to use Data Protector omnirc options, see the *HP Data Protector Help* index: “omnirc options”.

## Backup

This section contains procedures that are needed to back up virtual machines. For details on backup concepts, see “Backup concepts” (page 25).

### Backup limitations

- To be able to perform your backups, use only letters, numbers, single quotes, spaces, underscores, or hyphens in names of any VMware vCenter, VMware ESX(i), or VMware vCloud Director backup object (for example, virtual machines, datastores, datacenters, vApps, and so on). Special characters (such as +, =, @, and so on) are not supported.

For details, see the following websites:

[http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&displayKC=2017661](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&displayKC=2017661)

<http://www.vmware.com/support/developer/vddk/VDDK-1.2.1-RelnoteOPs.html>

[http://msdn.microsoft.com/en-us/library/windows/desktop/aa365247\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/aa365247(v=vs.85).aspx)

## Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Virtual Environment**, and select **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. Specify the application to be backed up:
  - In the **Client** drop-down list, select a VMware client.

---

**NOTE:** The drop-down list contains all clients that have been imported into the Data Protector cell as VMware vCenter, VMware ESX(i), or VMware vCloud Director clients. They have a corresponding label appended at the end of their names, such as **(VMware vCenter)**, **(VMware ESX(i))**, or **(VMware vCloud Director)**.

If you configured the Data Protector VMware (Legacy) integration as well, the drop-down list also contains the clients that have the Data Protector VMware Integration (Legacy) component installed. These clients have the **(Legacy VMware)** label appended at the end. To back up virtual machines using the Data Protector VMware (Legacy) integration, select one of these clients. For details, see “Data Protector VMware (Legacy) integration” (page 89).

---



If the selected VMware client is not configured correctly, a warning is displayed. Click **OK** to open the Configure Virtual Environment dialog box and provide the connection parameters as described in [“Importing and configuring VMware clients”](#) (page 42).

- In the **Backup host** drop-down list, select a system to be used to control the backup. The list contains all clients that have the Data Protector Virtual Environment Integration component installed.
- In **Datacenter/Organization**, select a datacenter or a VMware vCloud Director organization to back up from.

**NOTE:** If you have selected a standalone ESX(i) Server system in the **Client** option, there is only one datacenter available – **/ha-datacenter**.

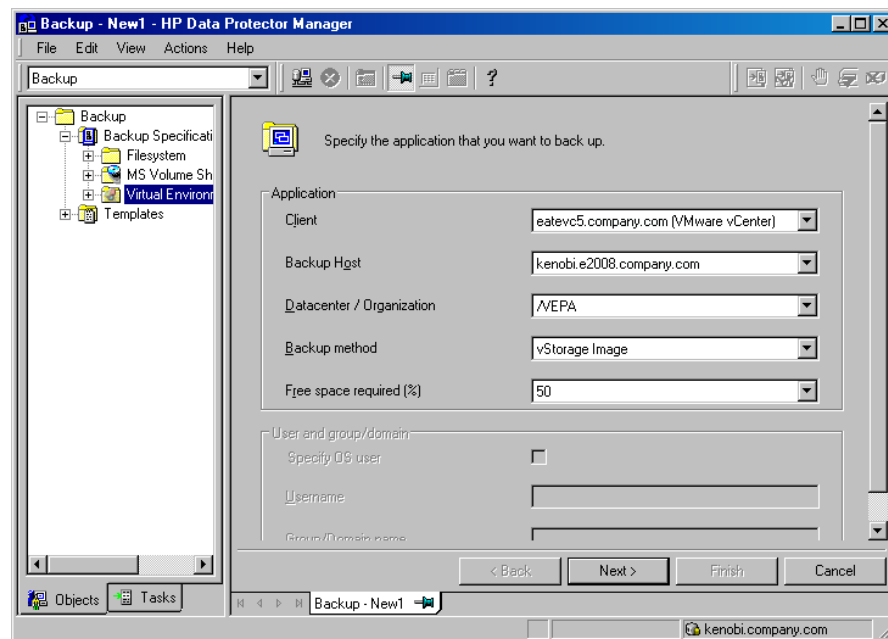
If you have selected a vCenter Server system in the **Client** option, you can select **All Datacenters** to back up virtual machines from different datacenters.

If you have selected a vCloud Director in the **Client** option, you can select **All Organizations** to back up virtual machines from different organizations.

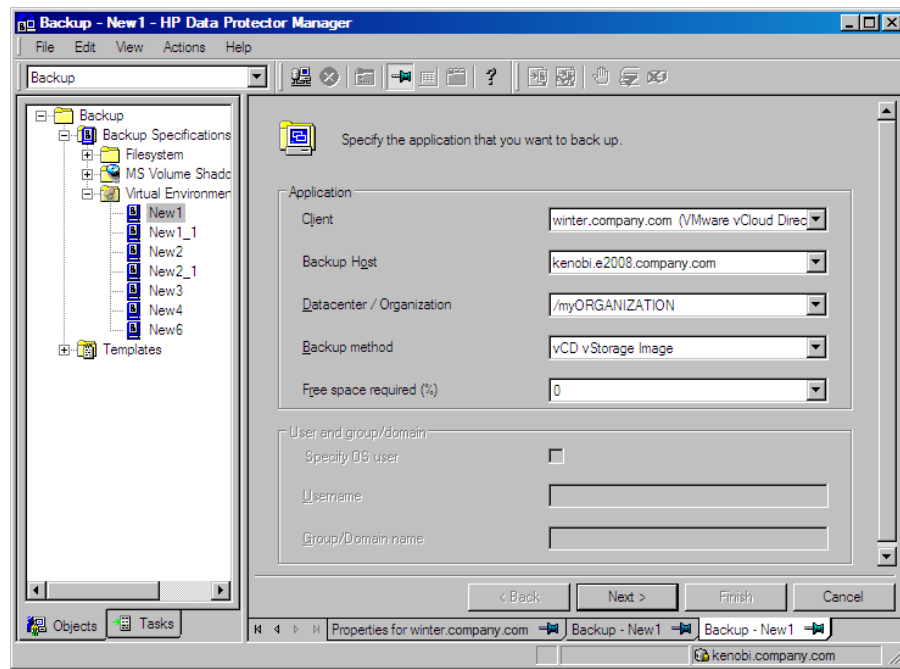
- In **Backup method**, the backup method is displayed:
  - **vStorage Image** for VMware vCenter and VMware ESX(i) clients
  - **vCD vStorage Image** for VMware vCloud Director clients
- In **Free space required [%]**, specify the percentage of disk space that should be free on a datastore before a virtual machine is backed up. The free space is calculated based on the size of the datastore where the virtual machine disks reside.

The check is performed separately for each virtual machine. For details, see [“Free space required option”](#) (page 37).

**Figure 21 Selecting a VMware vCenter Server client, backup host, and datacenter**



**Figure 22** Selecting a VMware vCloud Director client, backup host, and organization



Click **Next**.

**NOTE:** The settings specified in this page of the wizard cannot be changed once the backup specification is saved. To change the settings, you will have to create a new backup specification.

5. Select the objects that you want to back up. In the **Show** drop-down list, simplify your selection by choosing the **Hosts and Clusters** or **VMs and Templates** view. By default, **Hosts and Clusters** is displayed.

**NOTE:** If you switch the view after you have already selected one or more objects for backup, a warning dialog is displayed. Its confirmation clears the already selected objects, clicking **No** does no change to the view.


You can make your selections at different levels:

- For VMware vCenter and VMware ESX(i) clients:
  - ESX/ESXi Servers systems
  - Pools
  - vApps
  - VM folders
  - Individual VMs
  - VM disks
  - VM templates
- For VMware vCloud Director clients:
  - Organizations
  - vDatacenters


- vApps
- Individual VMs

If you select any level above individual VMs (for example, a vApp or a VMware vCloud Director organization), all VMs and VM disks contained in the selected item will be included in the backup specification. If VMs are added within the item after the backup specification is saved, they will also be backed up.


**Figure 23 Selecting VMware objects (vCenter Server client)**



**Figure 24 Switching view while selecting VMware objects (vCenter Server client) – warning**



**Figure 25 Selecting VMware objects (vCloud Director client)**



- ① **IMPORTANT:** In the object tree of a particular VMware client, a virtual machine may be displayed as selected in two different ways:
- The *blue* check mark indicates that the virtual machine is selected for backup in its entirety, including its configuration and all its virtual disks.  
If such virtual machine is backed up, it can be restored even if the original virtual machine does not exist anymore.
  - The *gray or black* check mark indicates that some or all virtual disks belonging to the virtual machine are selected. The virtual machine itself and its configuration is omitted from the backup.  
If such virtual machine is backed up, its disks can only be restored if the original virtual machine is still configured at the time of restore.

If your virtual machines are not configured yet, right-click the client system at the top or any of the virtual machines listed below, and select **Configure Virtual Machines**. For details, see [“Configuring virtual machines”](#) (page 49).

Click **Next**.

6. Select which devices to use for the backup.


To specify device options, right-click the device and select **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and which media pool to use.

Click **Next**.

7. Set backup options.

For information on application-specific backup options, see [“VMware backup options”](#) (page 61).

**Figure 26 Application-specific options**



Click **Next**.

8. Optionally, schedule the backup. See [“Scheduling backup sessions” \(page 64\)](#).

Click **Next**.

9. Save the backup specification, specifying a name and a backup specification group.



**TIP:** Preview your backup specification before using it for a real backup. See [“Previewing backup sessions” \(page 65\)](#).

**Table 11 VMware backup options**

Option	Description
<b>Pre-exec, Post-exec</b>	Specifies which command line to run on the backup host before (pre-exec) or after (post-exec) the backup. Do not use double quotes. Type only the name of the command and ensure that the command resides in the following directory on the backup host: <b>Windows systems:</b> <code>Data_Protector_home\bin</code> <b>Linux systems:</b> <code>/opt/omni/sbin</code>
<b>Backup memory file</b>	If this option is selected, the memory of a running virtual machine is saved into a file and backed up. Do not select this option if one or more virtual machines selected in the backup specification have the <b>Use quiescence snapshots</b> selected.


## Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.


In the Source page, you can modify the backup objects by using the **Show** drop-down list. In the drop-down list, the **Hosts and Clusters** view or the **VMs and Templates** view are available. The view that you have used during the creation of your backup specification has the string (Original) appended. If you switch the view, a warning dialog is displayed and its confirmation clears the already selected objects.

**NOTE:** When modifying a backup specification created with one of the previous versions of Data Protector, you can switch between **Selected**, **All**, **Hosts and Clusters**, and **VMs and Templates** views. The view **All** is available for the legacy backup specifications only and provides a legacy browsing mechanism. Selecting **Hosts and Clusters** or **VMs and Templates** and clicking **Yes** in the warning dialog clears the previous backup object selection and upgrades your browsing mechanism. After saving the backup specification only the **Hosts and Clusters** and **VMs and Templates** views are available.

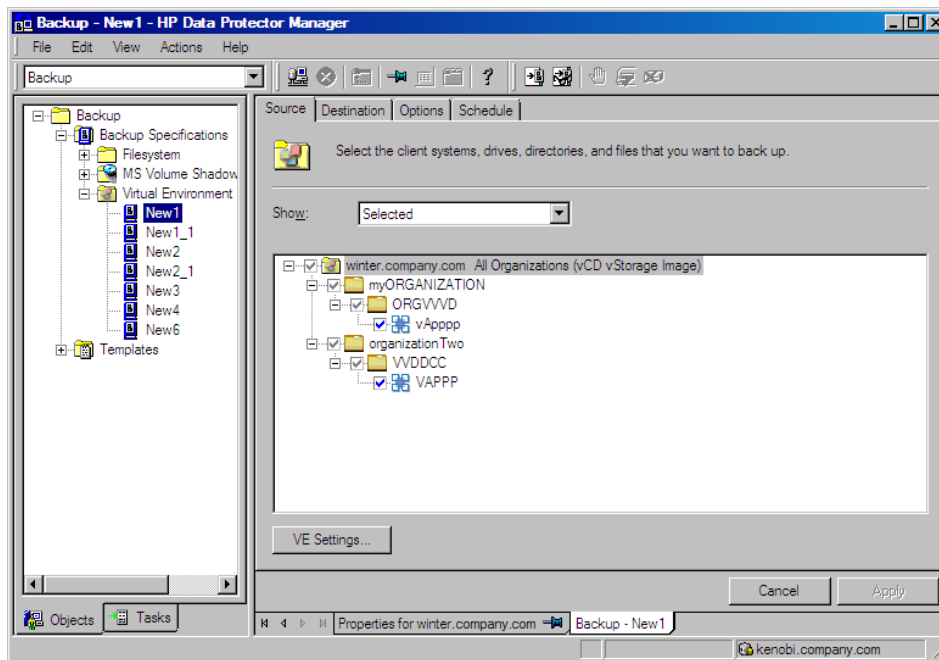
**Figure 27** Modifying a backup specification (VMware vCenter Server client)



**Figure 28** Switching view while modifying a backup specification (VMware vCenter Server client) – warning

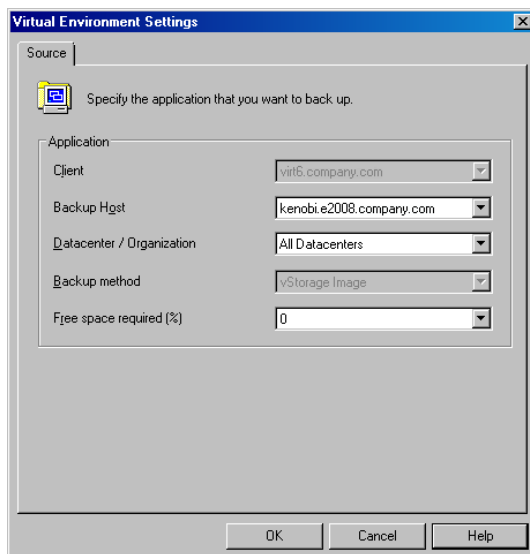


**Figure 29 Modifying a backup specification (VMware vCloud Director client)**

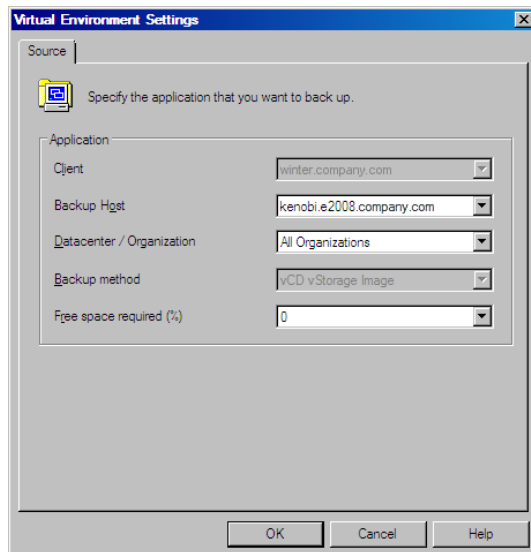


To display the virtual environment settings, in the Results Area, click **VE Settings**. Not all settings can be modified.

**Figure 30 Virtual environment settings (VMware vCenter Server client)**



**Figure 31 Virtual environment settings (VMware vCloud Director client)**



## Scheduling backup sessions

You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

### Scheduling example

To schedule differential backups at 8:00, 13:00, and 18:00 during week days:

1. In the Schedule property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under Recurring, select **Weekly**. Under Time options, select **8:00**. Under Recurring options, select **Mon, Tue, Wed, Thu, and Fri**. See “[Scheduling a backup session](#)” (page 65). Under Session options, select **Differential** from the **Backup type** drop-down list.

Click **OK**.

3. Repeat [Step 1](#) and [Step 2](#) to schedule differential backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.



**Figure 32 Scheduling a backup session**

**Schedule Backup**

Specify the desired backup time, frequency, duration, and type.

**Recurring**

☐ None  
☐ Daily  
☒ Weekly  
☐ Monthly

**Time options**

Time: 8:00 AM  
☐ Use starting  
1/25/2011

**Recurring options**

Every 1 week(s) on

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

**Session options**

Backup type: Differential  
Network load: ☒ High ☐ Medium ☐ Low  
Backup protection: Default

OK Cancel Help

## Previewing backup sessions

Preview the backup session to test it. You can use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Virtual Environment**. Right-click the backup specification you want to preview and select **Preview Backup**.
3. Specify the **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful preview.

### Using the Data Protector CLI

1. Log in to any client with the Data Protector User Interface component installed.
2. Open the command prompt and change to the directory in which the `omnib` command is located.
3. Execute:

```
omnib -veagent_list BackupSpecificationName -test_bar
```

## What happens during the preview?

The following are tested:

- Communication between the backup host and Data Protector
- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

## Starting backup sessions

Interactive backups are run on demand. They are useful for performing urgent backups or restarting failed backups.

To start a backup interactively, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **Virtual Environment**. Right-click the backup specification you want to use and select **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful backup session.

## Using the Data Protector CLI

1. Log in to any client with the Data Protector User Interface component installed.
2. Open the command prompt and change to the directory in which the `omnib` command is located.
3. Execute:

```
omnib -veagent_list BackupSpecificationName [-barmode  
VirtualEnvironmentMode] [ListOptions]
```

where *VirtualEnvironmentMode* is one of the following backup types:

```
full|diff|incr
```

The default is `full`.

For *ListOptions*, see the `omnib` man page or the *HP Data Protector Command Line Interface Reference*.

### Examples

To start a full backup using the backup specification `MyVirtualMachines`, execute:

```
omnib -veagent_list MyVirtualMachines -barmode full
```

To start a differential backup using the same backup specification, execute:

```
omnib -veagent_list MyVirtualMachines -barmode diff
```

## Preparing for disaster recovery

To do a disaster recovery, you need backups of the following VMware objects:

**Table 12 What must be backed up**

VMware object	How to back it up
ESX/ESXi Server console	<p><b>ESX Server systems:</b></p> <ol style="list-style-type: none"><li>1. Ensure that the Data Protector <code>Disk Agent</code> component is installed on all the ESX Server systems.</li><li>2. In the Backup context of the Data Protector GUI, right-click <b>Filesystem</b> and select <b>Add backup</b> to create a backup specification of the filesystem type. In the Source page of the backup specification, select ESX Server consoles of all ESX Server systems.  For details on what to back up, see the topic “ESX Server Configuration Backup and Restore procedure” at <a href="http://kb.vmware.com/selfservice/microsites/microsite.do">http://kb.vmware.com/selfservice/microsites/microsite.do</a>.</li><li>3. Start a backup using the newly created backup specification.</li></ol> <p><b>ESXi Server systems:</b></p> <p>With ESXi Server systems, it is not possible to install a Data Protector <code>Disk Agent</code>, so you will need to use VMware utilities to back up the configuration.</p> <p>A tool, <code>esxcfg-cfgbackup</code> is available from VMware. For information, see the VMware website.</p>
vCenter configuration database (applicable only for VirtualCenter environments)	<p>The vCenter configuration database can be an Oracle database or Microsoft SQL Server database. To back up the database, use the corresponding Data Protector integration. For example, if it is an Oracle database, follow these steps:</p> <ol style="list-style-type: none"><li>1. Ensure that the Data Protector <code>Oracle Integration</code> component is installed on the vCenter Server system.</li><li>2. In the Backup context of the Data Protector GUI, right-click <b>Oracle Server</b> and select <b>Add backup</b> to create a backup specification of the Oracle type. In <b>Application database</b>, type the name of the vCenter configuration database.  Continue with the backup specification creation as described in the <i>HP Data Protector Integration Guide for Oracle and SAP</i>.</li><li>3. Start a backup using the newly created backup specification.</li></ol>
VMware virtual machines	Back up the virtual machines as described in this chapter.

## Restore

This section contains procedures that are needed to restore virtual machines. For details on restore concepts, see “[Restore concepts](#)” (page 39).

### Restore limitations

- When restoring a VM to a datacenter with a datastore shared among several inventory objects (hosts or clusters) or several cluster nodes, and the restore option for its subsequent registration is selected, the restored VM may not register at the original inventory location:
  - If several inventory objects share the datastore, the VM is registered with the first available host or cluster.
  - If several cluster nodes share the datastore, the VM is registered with the first available cluster node.

If the VM must be registered at its original inventory location, migrate it appropriately after the restore session completes.

- Before a restore of a virtual machine backed up with the vCD vStorage Image backup method is started, the vApp where the selected VM resides enters maintenance mode for consistency reasons. As a consequence all VMs in this vApp are shut down.

- Restoring virtual machines to a specific datastore is not supported using the vCD Storage Image backup method.
- In a vSphere environment, do not use the left parenthesis (the ( character) in names of vSphere distributed port groups. Data Protector is unable to connect a restored virtual machine to a non-original vSphere network that uses distributed switches and is assigned such a distributed port group. In this case, after the restore, you need to manually connect the restored virtual machine to the desired network using vSphere Client.
- Recovering virtual machines after a restore to a directory is not supported for incremental or differential backups. The VMware Converter that is used for moving the restored VM images to an ESX Server or ESXi Server system recognizes only the full backup type.
- Recovering virtual machines from a backup chain (for example, full, incremental, incremental, incremental...) is supported only for a restore to a datacenter.

## Finding information for restore

You can find information about backup objects in the Data Protector IDB, such as which backup type and media were used, and which messages were displayed during the backup. To retrieve this information, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

In the Internal Database context, expand **Objects** or **Sessions**.

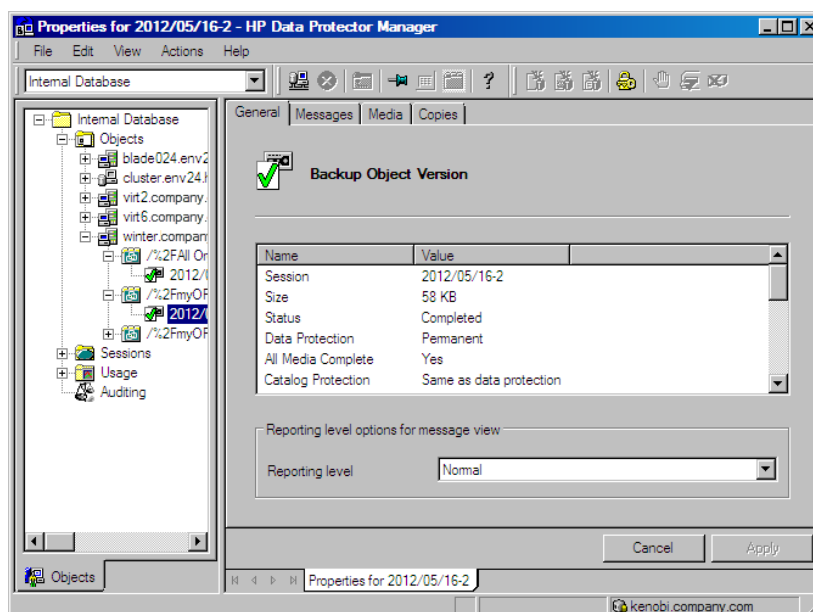
If you expand **Objects**, backup objects are sorted according to the virtual machine for which they were created. For example:

- Backup objects for the virtual machine /vm/mach1 from the datacenter ELDatacentro are listed under /%2FEldatacentro/0/%2Fvm%2Fmach1.
- Backup objects for the virtual machine /vm1Org22 from vApp vAppORG22, from vDatacenter vDCOrg22, from the VMware vCloud Director organization ORG22 are listed under /%2FORG22/8/%2FORG22%2FvDCOrg22%2FvAppORG22%2Fvm1Org22.

If you expand **Sessions**, backup objects are sorted according to the session in which they were created. For example, backup objects created in the session 2012/07/10-82 are listed under 2012/07/10-82.

To view details on a backup object, right-click the backup object and select **Properties**.

**Figure 33 Backup object information**





---

**TIP:** To view the messages displayed during the session, click the **Messages** tab.

---

## Using the Data Protector CLI

1. Log in to any client with the Data Protector User Interface component installed.
2. Open the command prompt and change to the directory in which the `omnidb` command is located.
3. Get a list of VMware backup objects created in a backup session with the session ID *SessionID*:  

```
omnidb -session SessionID
```
4. Get details on a backup object with the backup object name *BackupObjectName*:  

```
omnidb -veagent BackupObjectName -session SessionID -catalog
```

Here is one example of a backup object name:

```
gabriel.company.com::/%2FE1Datacentro/0/%2Fvm%2Fharbour
```

For details, see the `omnidb` man page or the *HP Data Protector Command Line Interface Reference*.

## Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Virtual Environment**, expand the relevant client and click the datacenter or the VMware vCloud Director organization from which you backed up.
3. In the Source page, Data Protector displays VMware objects.  
By using the **From** and **To** options, you can narrow the scope of displayed virtual machines to those backed up within the specified time interval.  
Select the objects that you want to restore.

---

**NOTE:** Data Protector restores the full restore chain for each selected VMware object, beginning with the last full backup session (even if that full backup is outside the specified time interval) and ending with the last backup session performed during the specified time interval.

---

Figure 34 Selecting VMware objects for restore (vCenter Server client)

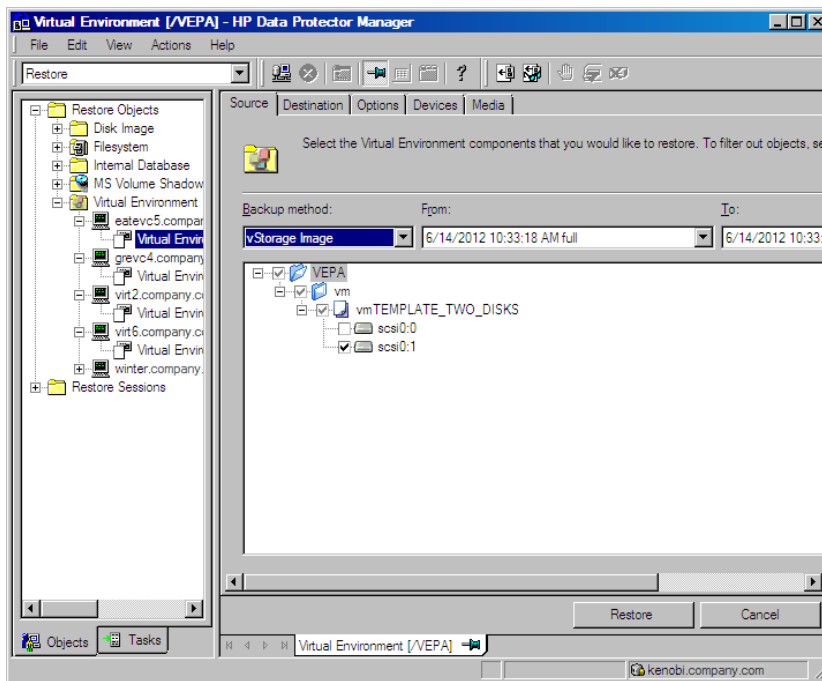
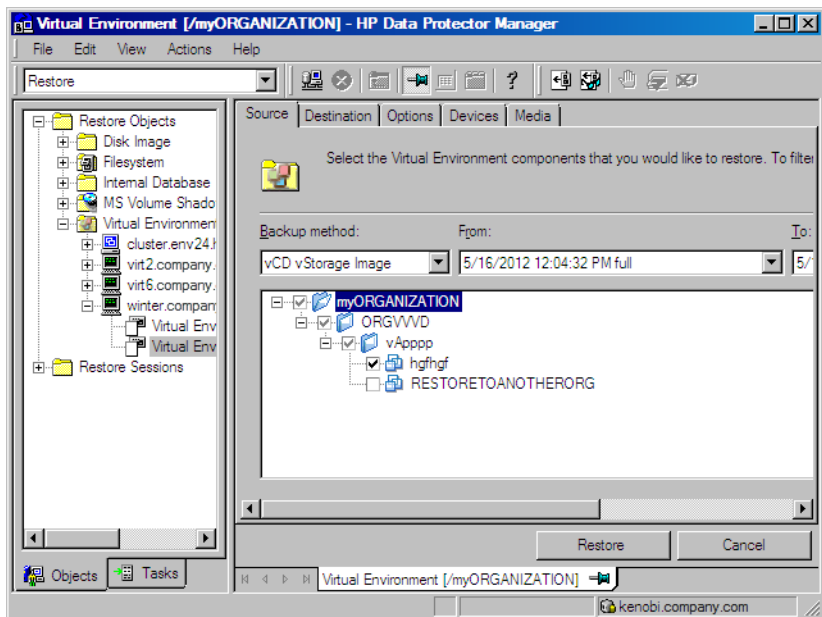
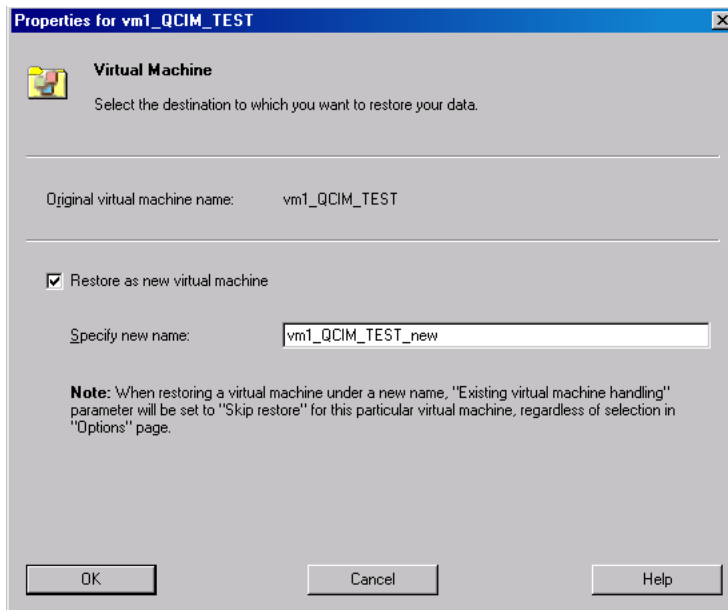


Figure 35 Selecting VMware objects for restore (vCloud Director client)



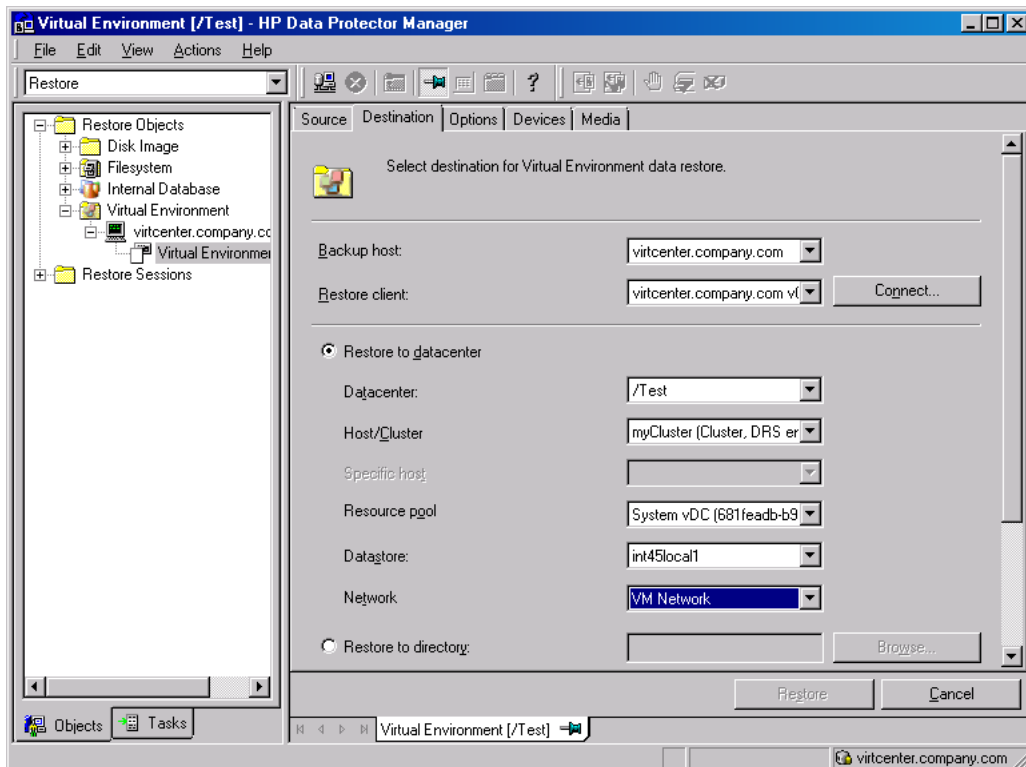
Optionally, by right-clicking the selected virtual machine and then clicking **Restore As / Into ...**, you can restore it as a new virtual machine. A new dialog window opens to specify the name of the virtual machine.

**Figure 36 Restoring as a new virtual machine**

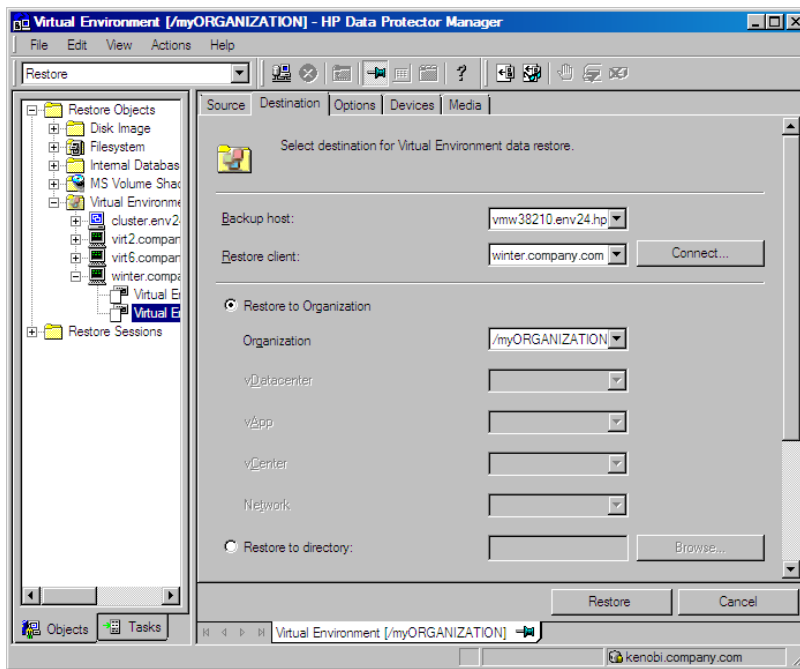


4. In the Destination page, specify the restore destination. For details, see [“Restore destination \(VMware vCenter Server and VMware ESX\(i\) Server clients\)”](#) (page 73) or [“Restore destination \(VMware vCloud Director client\)”](#) (page 74).

**Figure 37 Restore destination (VMware vCenter Server client)**

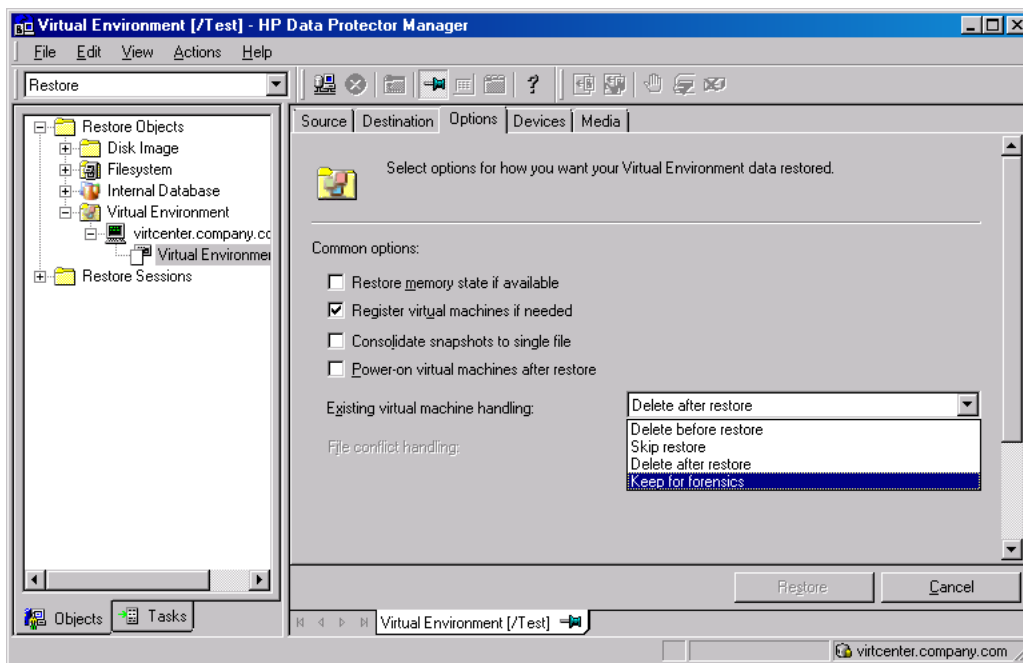


**Figure 38 Restore destination (VMware vCloud Director client)**



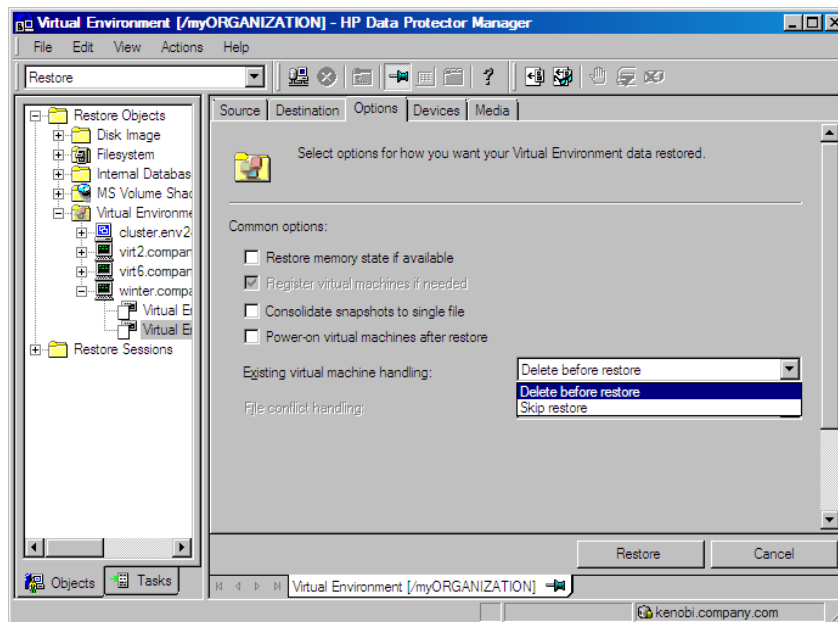
5. In the Options page, specify the VMware restore options. For details, see [“Restore options”](#) (page 75).

**Figure 39 Restore options (VMware vCenter Server and VMware ESX(i) Server clients)**





**Figure 40 Restore options (VMware vCloud Director client)**



6. In the **Devices** page, select devices to use for restore.
7. Click **Restore**.
8. In the **Start Restore Session** dialog box, click **Next**.
9. Specify **Report level** and **Network load**.

Click **Finish** to start the restore.

The message `Session completed successfully` is displayed at the end of a successful session.

**NOTE:** If the restore fails, see “Cleaning up a datastore after a failed restore” (page 84).

**Table 13 Restore destination (VMware vCenter Server and VMware ESX(i) Server clients)**

GUI/CLI option	Description
<b>Backup host/</b> -barhost	Specifies the client with the Virtual Environment Integration installed to control the restore session. By default, the same client as used for the backup is selected.
<b>Restore client/</b> -apphost	Specifies the client that the selected virtual machine objects should be registered and restored to. By default, the client from which the virtual machines were backed up is selected. To change the client configuration, click the <b>Connect</b> button.
<b>Restore to datacenter/</b> -instance -newinstance	Select this option to restore a virtual machine to a datacenter. By default, virtual machines are restored to the original datacenter.
<b>Host/Cluster/</b> -host/cluster	Select the ESX(i) Server system or the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original ESX(i) Server system or cluster.
<b>Specific host/</b> -specificHost	Select the specific ESX(i) Server system in the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original ESX(i) Server system.
<b>Resource pool/</b> -resourcePool	Select the resource pool on the ESX(i) Server system or the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original resource pool.

**Table 13 Restore destination (VMware vCenter Server and VMware ESX(i) Server clients) (continued)**

GUI/CLI option	Description
<b>Datastore/</b> -store	Specifies the datastore to which the virtual machines should be restored. You can choose among all datastores that are accessible from the selected restore target host. If you leave this option empty, the virtual machines are restored to the original datastore.
<b>Restore to directory/</b> -directory	<p>Select this option to restore virtual-machine files to a directory (outside of the datacenter) on the backup host. You can use the <b>Browse</b> button to find the target directory.</p> <p>After such a restore, the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX Server or ESXi Server system, using the VMware Converter as described in <a href="#">“Recovering virtual machines after restore to a directory”</a> (page 78).</p>

**Table 14 Restore destination (VMware vCloud Director client)**

GUI/CLI option	Description
<b>Backup host/</b> -barhost	Specifies the VMware vCloud Director client with the Virtual Environment Integration installed to control the restore session. By default, the same client as used for the backup is selected.
<b>Restore client/</b> -apphost	<p>Specifies the client that the selected virtual machine objects should be registered and restored to. By default, the client from which the virtual machines were backed up is selected.</p> <p>To change the client configuration, click the <b>Connect</b> button.</p>
<b>Restore to organization/</b> -neworganization	Select this option to restore virtual machines to an organization. By default, virtual machines are restored to the original organization, but you can specify a different destination organization.
<b>vDatacenter/</b> -virtual_datacenter_path -virtual_datacenter_uuid	Specifies the vDatacenter to which the virtual machines should be restored. By default, the virtual machines are restored to the original vDatacenter.
<b>vApp/</b> -vapp_path -vapp_uuid	<p>Specifies the vApp to which the virtual machines should be restored. By default, the virtual machines are restored to the original vApp.</p> <p>Note that the virtual machines are restored as a new vApp if the original vApp is no longer available or all virtual machines of the selected vApp are restored.</p>
<b>vCenter/</b> -vcenter_path -vcenter_uuid	<p>Specifies the vCenter to which the virtual machines should be restored. By default, the virtual machines are restored to the original vCenter.</p> <p>If this option is not selected, the virtual machines are restored to the vCenter used by the selected vDatacenter or vApp.</p>
<b>Network/</b> -network_name -network_uuid	<p>Select a network to enable virtual machines communication.</p> <p>If an individual virtual machine is restored into an existing vApp, the vApp network is selected.</p> <p>If all virtual machines of the selected vApp are restored, the Organization network is selected.</p>
<b>Restore to directory/</b> -directory	<p>Select this option to restore virtual-machine files to a directory (outside of the organization) on the backup host. You can use the <b>Browse</b> button to find the target directory.</p> <p>After such a restore, the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX Server or ESXi Server system, using the VMware Converter as described in <a href="#">“Recovering virtual machines after restore to a directory”</a> (page 78).</p>

**Table 15 Restore options**

GUI/CLI option		Description
<b>Restore memory state if available/</b> -memory		Available if <b>Restore to datacenter</b> is selected. Select this option to restore the memory file if it is included in the backup.
<b>Register virtual machines if needed/</b> -register		Available if <b>Restore to datacenter</b> is selected. Select this option to register restored virtual machines. If this option is not selected, you need to manually recover the restored virtual machines as described in <a href="#">“Recovering virtual machines manually” (page 77)</a> . Default: selected.
<b>Consolidate snapshots to single file/</b> -consolidate		Select this option to commit all snapshots (including non-Data Protector ones) to the virtual machine base once a virtual machine is restored. Available if <b>Restore to datacenter</b> is selected.
<b>Power-on virtual machines after restore/</b> -poweron		Select this option to power the virtual machines on once they are restored. Available if <b>Restore to datacenter</b> is selected.
<b>Existing virtual machine handling</b>	Specifies Data Protector's behavior when restoring existing virtual machines.	
	<b>Delete before restore/</b> -deletebefore	Select this option to delete an existing virtual machine before it is restored, and then restore it from new. The existing virtual machine is deleted even if it resides in a different datacenter from your target datacenter.  This is the space-efficient option, but is less secure, since the old virtual machine is not available if the restore fails, so select it with caution.
	<b>Skip restore/</b> -skip	Select this option to skip the restore of an existing virtual machine. This allows you to restore missing virtual machines without affecting existing ones.
	<b>Delete after restore /</b> -deleteafter	Select this option to delete an existing virtual machine after it is restored. The existing virtual machine is deleted even if it resides in a different datacenter than your target datacenter. If the restore fails, the existing virtual machine is not deleted.  This option is not available for VMware vCloud Director client. Default: selected.
	<b>Keep for forensics/</b> -keep_for_forensics	Select this option to mark an existing virtual machine with a timestamp. The virtual machine which is kept for forensics is powered off after the restore and remains at the original location. It does not affect consecutive backups of the original virtual machine.  This option is not available for VMware vCloud Director and Microsoft Hyper-V clients.
<b>File conflict handling</b>	Specifies Data Protector's behavior when restoring existing files.	
	<b>Overwrite/</b> -overwrite	Select this option to overwrite existing files with those from the backup. Default: selected.
	<b>Keep latest/</b> -latest	Select this option to leave an existing file intact if it is more recent than the one from the backup. Otherwise, the existing file is overwritten with the one from the backup.
	<b>Skip/</b> -skip	Select this option to preserve an existing file (the file is not restored from the backup).

## Restoring using the Data Protector CLI

1. Log in to any client with the Data Protector User Interface component installed.
2. Open the command prompt and change to the directory in which the `omnir` command is located.
3. Execute:

### VMware vCenter Server or VMware ESX(i) Server client

```
omnir -veagent
-virtual-environment vmware
-barhost BackupHost
-apphost OriginalVMwareClient
-instance OriginalDatacenter
-method vStorageImage
[-session BackupID]
VirtualMachine [VirtualMachine...]
[VMwareClient | Directory]

VirtualMachine
-vm VMPATH [-new_name NewVirtualMachineName] [-disk DiskName [-disk Disk...]]

VMwareClient
[-newinstance TargetDatacenter]
[-store TargetDatastore]
[-network_name TargetNetwork]
[-destination TargetVMwareClient]
[-consolidate] [-memory] [-register] [-poweron]
[-deletebefore | -deleteafter | -skip | -keep_for_forensics]

Directory
-directory RestoreDirectory
[-overwrite | -skip | -latest]
```

### VMware vCloud Director client

```
omnir -veagent
-virtual-environment vcd
-barhost BackupHost
-apphost OriginalVMwareClient
-instance OriginalDatacenter
-method vCDvStorageImage
[-session BackupID]
VirtualMachine [VirtualMachine...]
[NewOrganization | Directory]

VirtualMachine
-vm VMPATH [-disk DiskName [-disk Disk...]]

NewOrganization
[-neworganization TargetOrganization]
[[-virtual_datacenter_path | -virtual_datacenter_uuid ]
TargetVDC]
[[-vapp_path | -vapp_uuid ] TargetVApp]
[[-vcenter_path | -vcenter_uuid ] TargetVCenter]
[[-network_name | -network_uuid ] TargetNetwork]
[-consolidate] [-memory] [-register] [-poweron]
[-deletebefore | -skip]

Directory
-directory RestoreDirectory
[-overwrite | -skip | -latest]
```

For a description of all the options, see the `omnir` man page or the *HP Data Protector Command Line Interface Reference*.

- ❗ **IMPORTANT:** A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session. Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you have to specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.
- The `omnir` syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

---

**NOTE:** If the restore fails, see “Cleaning up a datastore after a failed restore” (page 84).

---

### Example (Restoring virtual machines to a datacenter)

Suppose you want to restore the virtual machine `/vm/machineA` and the individual disks (`scsi0:0` and `scsi0:1`) of the virtual machine `/vm/machineB`. At the time of backup, the virtual machines were running on the ESX Server systems that belonged to the datacenter `/MyDatacenter` managed by the vCenter Server system `vcenter.company.com`. The virtual machines were backed up with the `vStorage Image` backup method.

To restore them to the original location, using the backup session `2011/01/11-1`, and to ensure that the newly restored virtual machines are put online when the session completes, execute:

```
omnir -veagent -virtual-environment vmware -barhost
backuphost.company.com -apphost vcenter.company.com -instance
/MyDatacenter -method vStorageImage -session 2011/1/11-1 -vm /vm/machineA
-vm /vm/machineB -disk scsi0:0 -disk scsi0:1 -memory -poweron
```

### Example (Restoring virtual machines to a directory)

Suppose the virtual machines `/MyVirtualMachines/machineA` and `/MyVirtualMachines/machineB` were backed up in the session `2011/02/12-5` from the datacenter `/MyDatacenter` that is managed by the vCenter Server system `vcenter.company.com`, using the `vStorage Image` backup method. To restore the virtual machines outside of the datacenter, to the directory `C:\tmp` on the backup host `backuphost.company.com`, execute:

```
omnir -veagent -virtual-environment vmware -barhost
backuphost.company.com -apphost vcenter.company.com -instance
/MyDatacenter -method vStorageImage -session 2011/2/12-5 -vm
/MyVirtualMachines/machineA -vm /MyVirtualMachines/machineB -directory
c:\tmp
```

## Recovering virtual machines manually

There are two different scenarios in which you need to recover virtual machines manually after they have been restored with Data Protector:

- If you have restored the virtual machines to a directory on a backup host (**Restore to directory**). For details, see “Recovering virtual machines after restore to a directory” (page 78).
- If you have restored the virtual machines to a datacenter (**Restore to datacenter**) without selecting the restore option `Register virtual machines` if needed. For details, see “Recovering virtual machines after restore to a datacenter” (page 83).

## Recovering virtual machines after restore to a directory

The steps for recovering virtual machines after restore to a directory depend on the format in which the virtual machine configuration file was backed up.

### Recovering with the VM configuration file in the VMX format

Suppose the virtual machine `helios` was restored to the directory `C:\tmp\helios` on the backup host using the following backup session:

- Backup method: **vStorage Image**
- Backup type: Incremental
- CBT: Enabled and used

To move the virtual machine files manually to the ESX(i) Server system `dioxide.company.com` managed by the vCenter Server system `bmwvc2.company.com`, using the VMware Converter:

1. Display the contents of the directory `C:\tmp\helios`:

```
helios.vmdk
helios.vmx
helios.vmdk
scsi0-0.cbt
scsi0-0.meta
helios-flat.vmdk
helios.vmx-1
helios.vmdk-1
scsi0-0.cbt-1
scsi0-0.meta-1
helios.vmx-2
helios.vmdk-2
scsi0-0.cbt-2
scsi0-0.meta-2
```

Note that all files backed up in the last full, differential, and the selected incremental session are restored.

2. Share the folder `C:\tmp\helios` so that it can be accessed from the system with the VMware Converter installed.
3. Log in to the system with the VMware Converter installed and open the VMware Converter user interface.
4. Click **Convert Machine** to open the Conversion wizard.
5. In the Source System page, select **VMware Workstation or other VMware virtual machine** for the source type, browse to the `C:\tmp\helios` directory, and select the `helios.vmx` file. Click **Next**.

**Figure 41 Conversion (Source System)**

The screenshot shows the 'Conversion' window with the 'Source System' tab selected. The window title is 'Conversion'. Below the title bar, the text 'Source System' is followed by 'Select the source system you want to convert'. On the left, a sidebar lists the following options: 'Source System' (selected), 'Destination System', 'Destination Virtual Machine', 'Destination Location', 'Options', and 'Summary'. The main area contains the following fields and controls:

- Source:** C:\helios\helios.vmx
- Destination:** none
- Select source type:** VMware Workstation or other VMware virtual machine (dropdown menu)
- Convert a virtual machine from VMware Workstation, VMware Player, VMware product.**
- Browse for source virtual machine or image** (section header)
- Virtual machine file:** C:\helios\helios.vmx (text field)
- Browse...** (button)
- [View source details...](#) (link)

**NOTE:** In our example, the VMware Converter is installed on the backup host.

6. In the Destination System page, select **VMware Infrastructure virtual machine** for the destination type and provide the login credentials for the vCenter Server system.  
Click **Next**.

**Figure 42 Conversion (Destination System)**

The screenshot shows the 'Conversion' window with the 'Destination System' tab selected. The window title is 'Conversion'. Below the title bar, the text 'Destination System' is followed by 'Select a host for the new virtual machine'. On the left, a sidebar lists the following options: 'Source System' (link), 'Destination System' (selected), 'Destination Virtual Machine', 'Destination Location', 'Options', and 'Summary'. The main area contains the following fields and controls:

- Source:** C:\helios\helios.vmx
- Destination:** none
- Select destination type:** VMware Infrastructure virtual machine (dropdown menu)
- Creates a new virtual machine for use on a VMware Infrastructure pr**
- VMware Infrastructure server details** (section header)
- Server:** bmvvc2 (dropdown menu)
- User name:** Administrator (text field)
- Password:** (password field with masked characters)

7. In the Destination Virtual Machine page, specify the name under which the virtual machine should be recovered.  
Click **Next**.

**Figure 43 Conversion (Destination Virtual Machine)**

**Conversion**

**Destination Virtual Machine**  
Select the destination VM name and folder

Source: C:\helios\helios.vmx      Destination: helios on bmw

Name:

Inventory for: bmwvc2      Search for name with:

bmwvc2		VM name
	New Datacenter 4.1 [31 VMs]	(w2k3sp2 JP) CLONE
	Discovered virtual machine	GRE_41_2003x64EE_25
		GRE_41_2003x64EE_30
		GRE_41_2008x32EE_21

[Refresh](#)

[Help](#)   [Export diagnostic logs...](#)   [< Back](#)   [Next >](#)

8. In the Destination Location page, select the destination ESX(i) Server system and datastore.

**Figure 44 Conversion (Destination Location)**

**Conversion**

**Destination Location**  
Select the location for the new virtual machine

Source: C:\helios\helios.vmx      Destination: helios on dioxide

Inventory for: bmwvc2      Total source disks size: 9 GB

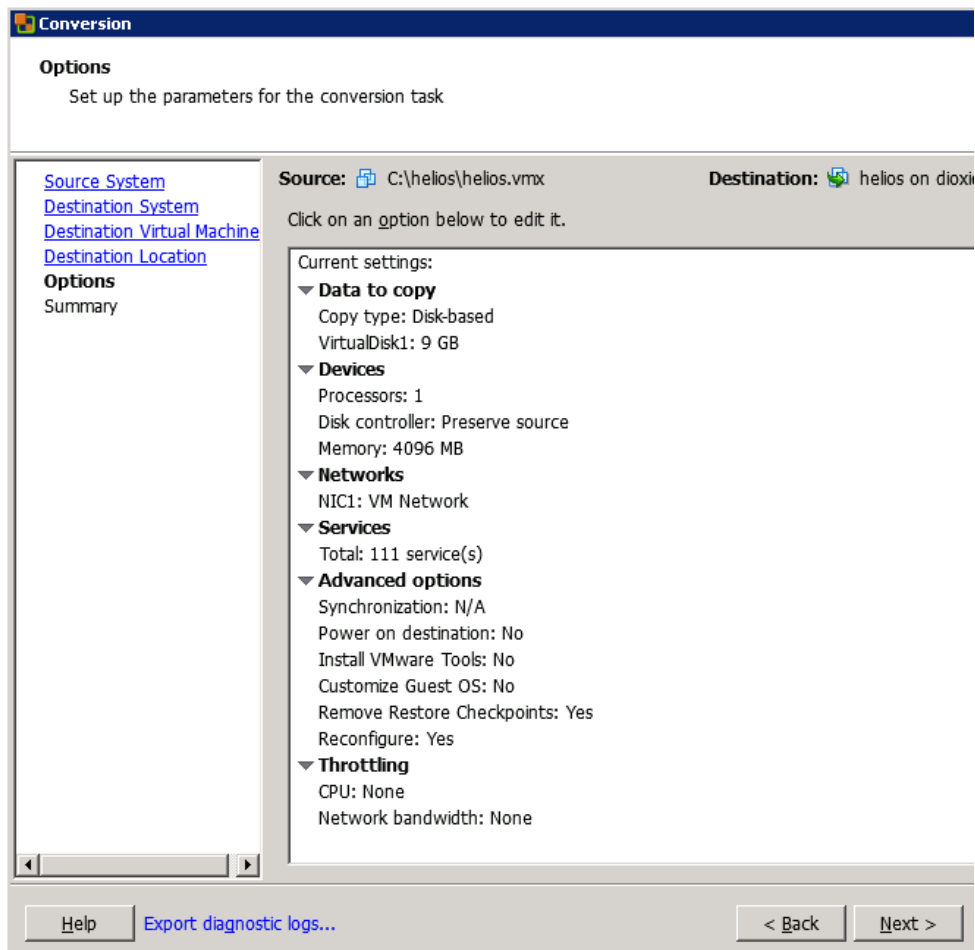
New Datacenter 4.1		Datastore
	bmw3.company.com	ergo:iscsi storage carbon dioxide
	bmw4.company.com	Capacity: 276.25 GB
	carbon.company.com	Free: 262.87 GB
	dioxide.company.com	Type: VMFS

Virtual machine version  
Version 7

9. In the Options page, edit the options and click **Next**.



**Figure 45 Conversion (Options)**



10. In the Summary page, review your selection and click **Finish**.

**Figure 46 Conversion (Summary)**

**Conversion**

**Summary**  
Review the conversion parameters

**Source:** C:\helios\helios.vmx **Destination:** helios on bmwvc2

[Source System](#)  
[Destination System](#)  
[Destination Virtual Machine](#)  
[Destination Location](#)  
[Options](#)  
**Summary**

**Source system information**

Source type:	VMware Workstation or other VMware virtual machine
Path:	C:\helios\helios.vmx
No throttling information	

**Destination system information**

Virtual machine name:	helios
Host/Server:	bmwvc2
Connected as:	Administrator
VM folder:	New Datacenter 4.1
Host system:	dioxide.company.com
Resource pool:	Default
Power on after conversion:	No
Number of CPUs:	1
Physical memory:	4096 MB
Network:	Preserve NIC count
NIC1	Connected
	VM Network
Storage:	Disk-based cloning
Number of disks:	1
Create disk 0 as:	Monolithic flat disk
Configuration files datastore:	ergo:iscsi storage carbon dioxide

**Destination customization**

Install VMware Tools:	No
Customize guest OS:	No
Remove restore checkpoints:	Yes
Reconfigure virtual machine:	Yes

[Help](#) [Export diagnostic logs...](#) [< Back](#) [Finish](#)

11. Open the Datastore Browser and upload the files created in the incremental and differential backup sessions to the virtual machine directory:

```
helios.vmx-1  
helios.vmdk-1  
scsi0-0.cbt-1  
scsi0-0.meta-1  
helios.vmx-2  
helios.vmdk-2  
scsi0-0.cbt-2  
scsi0-0.meta-2
```

12. Power the virtual machine on.

## Recovering with the VM configuration file in the XML format

Follow the procedure:

1. Open vSphere Client and log in to an ESX(i) Server or vCenter Server system.  
If the virtual machine is still configured, remove all its hard disks:
  - a. In the inventory object tree, right-click the virtual machine and select **Edit Settings...**
  - b. In the Virtual Machine Properties window, in the **Hardware** tab, select each hard disk and click **Remove**.
  - c. Click **OK** to confirm the removal.If the virtual machine is no longer present, configure a new virtual machine without hard disks, and use the name of the original virtual machine.  
In either case, remember the associated datastore name.
2. Upload the virtual machine files that were created during the backup session:
  - a. In the inventory objects tree, select the ESX(i) Server system that hosts the virtual machine.
  - b. Click the **Configuration** tab and select **Storage** under Hardware.
  - c. Right-click the datastore name and select **Browse Datastore...**
  - d. In the Datastore Browser window, in the folder tree, select the virtual machine folder, and click a corresponding icon on the window toolbar. Select **Upload File...** or **Upload Folder...** as appropriate.
  - e. Select all applicable files and complete the upload.
3. Add hard disks to the virtual machine by reusing their backup copies:
  - a. In the inventory object tree, right-click the virtual machine and select **Edit Settings...**
  - b. In the Virtual Machine Properties window, click **Add**.
  - c. In the Add Hardware window, select **Hard Disk** and click **Next**.
  - d. Select **Use an existing virtual disk** and click **Next**.
  - e. Click **Browse**.
  - f. In the Browse Datastores window, browse to the appropriate datastore and open the virtual machine folder. Select the virtual disk file and click **OK**.
  - g. Follow the Add Hardware wizard to complete the process.
  - h. Repeat the substeps from b to g for each additional hard disk for which a backup copy exists.
4. Power the virtual machine on.

## Recovering virtual machines after restore to a datacenter

If you have restored a virtual machine to a datacenter without selecting the option Register virtual machines if needed:

1. Open the Datastore Browser and browse to the restored virtual machine directory.
2. Right-click the virtual machine \* .vmtx file and select **Add to Inventory**.
3. Follow the wizard and click **Finish**.

## Recovering virtual machines after restore to an organization

For a detailed procedure on how to recover virtual machine to a vCloud Director organization, see VMware vCloud Director Administrator's Guide.

## Restoring using another device

You can restore using a device other than that used for backup. For details, see the *HP Data Protector Help* index: "restore, selecting devices for".

## Cleaning up a datastore after a failed restore

Sometimes, when a virtual machine restore fails, Data Protector creates extra files on the virtual machine datastore. If these files are not deleted, corrupt virtual machine backups may be created in subsequent sessions and, consequently, restore from such a backup also fails.

Suppose the virtual machine `MyVirtualMachine` failed to be restored. To clean up the datastore after the restore:

1. Open the VMware vSphere client.
2. Right-click the virtual machine and select **Delete from disk**.
3. Open the **Datastore Browser**.

The directory `MyVirtualMachine` should no longer be there.

Check if there are any extra directories:

`MyVirtualMachine_1`

`MyVirtualMachine_2`

and so on.

Right-click each such directory and select **Delete from disk**.

## Disaster recovery

Disaster recovery is very complex, involving different products from different vendors. Check the instructions from the guest operating systems and VMware on how to prepare for it.

The following are the main steps needed to recover a virtual machine after a disaster:

1. Reinstall the VMware environment. The configuration should be the same as during the backup.
2. Install Data Protector in the newly configured environment.
3. Restore the service console of the ESX Server system on which the virtual machine was running to the newly configured ESX Server system from a Data Protector filesystem backup.

For details on what to restore, see the topic “ESX Server Configuration Backup and Restore procedure” at <http://kb.vmware.com/selfservice/microsites/microsite.do>.

For details on how to restore from a filesystem backup, see the *HP Data Protector Help*.

4. Restore the original vCenter database (if needed). For details, see the Data Protector integration that was used to back up the database.
5. Restore the virtual machine from a Data Protector Virtual Environment backup as described in this chapter.

## Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run a backup or restore session, a monitor window shows the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the `User Interface` component installed, using the `Monitor` context.

To monitor a session, see the *HP Data Protector Help* index: “viewing currently running sessions”.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Virtual Environment integration.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HP Data Protector Help* index: “patches” on how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, and other information.

## Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors reported in the `debug.log` on the backup host.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HP Data Protector Help*.

Additionally, if your backup failed:

- Check the configuration of the vCenter Server or standalone ESX(i) Server system as described in “[Checking the configuration of VMware clients](#)” (page 48).

## Problems

### Problem

#### **An incremental or differential CBT backup session fails**

When performing an incremental or differential backup session with the Use changed block tracking option enabled, the session fails with an error similar to the following:

```
[Critical] From: OB2BAR_VEPA_BAR@droid.company.com "/New Datacenter 4.1"
Time: 2/10/2011 11:14:52 AM
Virtual Machine 'ddd': Could not gather changed blocks on disk scsi0:0 ...
```

The reason may be that you performed a restore session and forgot to run a full backup session to start a new backup chain.

To make sure changed block tracking is working properly, certain requirements need to be met. For details, see

[http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US=displayKC=1020128](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US=displayKC=1020128)

### Action

1. Run a full backup session.
2. Run an incremental or differential backup session.

### Problem

#### **After a restore or move to a different folder, backups are not performed correctly**

After you have restored or moved a virtual machine to a different folder, the virtual machine is not backed up correctly. For example, instead of an incremental backup, a full backup is performed.

The reason for this is that the datacenter configuration file has been updated. As a result, it contains two virtual machine sections with the same UUID; this is from where the inconsistencies arise.

### Action

Re-configure the virtual machine:

1. Open the backup specification.
2. In the Source page, right-click the VMware client and select **Configure Virtual Machines**.
3. Click **OK**.

## Problem

### A restore session is using LAN instead of SAN

If the backup host used to control a restore session is a virtual machine, Data Protector automatically switches to LAN transportation mode.

## Action

To use SAN transportation mode for restore, configure a backup host on a physical system (that is, install the Virtual Environment Integration component on a physical system) and select this system as a backup host for the restore session.

For details on transportation modes, see the *VMware Virtual Disk Programming Guide*.

## Problem

### A restore session using SAN transportation mode fails

A restore session that is using SAN transportation mode fails with a message similar to the following:

```
[Critical] From: OB2BAR_VEAgent@dpi00019.company.com
"/BlrVirtual01_ESX401" Time: 13-03-2011 12:22:57
Virtual Machine 'Win2k3_x64_dpi00002': Error restoring item
\alf9f4e3-482d-4b7f-afcb-cb16babe1980\%2FBlrVirtual01_ESX401\vms
\%2FBlrVirtual01_ESX401%2Fhost%2Fclus01%2FWin2k3_x64_dpi00002\
images\3\scsi2:15.
```

This may happen if a storage volume that is shared between the backup host and an ESX(i) Server system is read-only.

## Action

1. Log in to the backup host and open the command prompt.
2. Execute diskpart.

```
C:\Users\Administrator>diskpart
```

```
Microsoft DiskPart version 6.1.7600
Copyright (C) 1999-2008 Microsoft Corporation.
On computer: TPC134
```

3. Set the SAN policy to onlineAll.

```
DISKPART> san policy=onlineAll
```

```
DiskPart successfully changed the SAN policy for the current
operating system.
```

4. Select the disk (storage volume) that should be used for restore.

```
DISKPART> list disk
```

Disk ###	Status	Size	Free	Dyn	Gpt
-----	-----	-----	-----	---	---
Disk 0	Online	136 GB	1024 KB		
Disk 1	Offline	14 GB	14 GB		
Disk 2	Offline	14 GB	14 GB		
Disk 3	Offline	14 GB	14 GB		
Disk 4	Offline	14 GB	14 GB		
Disk 5	Offline	50 GB	50 GB		
Disk 6	Offline	14 GB	14 GB		
Disk 7	Offline	14 GB	14 GB		

```
DISKPART> select disk 1
```

5. Bring the disk online.

```
DISKPART> online disk
```

```
DiskPart successfully onlineed the selected disk.
```

6. Ensure the disk is not read-only.

List the disk properties.

```
DISKPART> detail disk
```

```
HP OPEN-V SCSI Disk Device
Disk ID: 00000000
Type    : FIBRE
Status  : Online
Path    : 0
Target  : 0
LUN ID  : 0
Location Path : UNAVAILABLE
Current Read-only State : Yes
Read-only : Yes
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
```

There are no volumes.

Clear the read-only attribute.

```
DISKPART> attribute disk clear readonly
```

Disk attributes cleared successfully.

List the disk properties again.

```
DISKPART> detail disk
```

```
HP OPEN-V SCSI Disk Device
Disk ID: 00000000
Type    : FIBRE
Status  : Online
Path    : 0
Target  : 0
LUN ID  : 0
Location Path : UNAVAILABLE
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
```

There are no volumes.

Exit the session.

```
DISKPART> exit
```

7. Restart the restore session.

### Problem

#### **Slow performance of vepa\_util.exe browse command on newer Red Hat Enterprise Linux (RHEL) versions**

When executing the `vepa_util.exe browse` command on a newer RHEL version its performance is significantly slower than on other operating systems.

### Action

The root cause of the problem is that on newer versions of RHEL systems name service cache daemon is not enabled by default.

Start the name service cache daemon by invoking the following command: `/etc/init.d/nscd start`. To enable automatic daemon start-up during system start-up, execute: `chkconfig nscd on`.

#### Problem

#### **Restore of a virtual machine using an ESX(i) Server system ends with a corrupted VM guest operating system**

In a vSphere environment, when you restore a virtual machine to a `/ha-datacenter` using an ESX(i) 5.0 Server system or later as a restore client, the restore session ends successfully, but the guest operating system on the virtual machine is corrupted.

#### Action

When selecting the objects for restore, in the **Destination** page, in the **Restore client** drop-down list, select a vCenter Server instead of an ESX(i) Server system.

#### Problem

#### **Backing up many VMs from VMware vCloud Director at once fails**

When backing up many virtual machines (for example, 30) from numerous vApps in the VMware vCloud Director at the same time, the backup session may fail due to a limitation on the number of simultaneously opened HTTP connections to the vCenter Server system.

#### Action

Reduce the limit of opened parallel connections by setting the omnirc option `OB2_VEAGENT_VCENTER_CONNECTION_LIMIT` to a lower number, for example, to 5. By default, the option is set to 10.



---

## 2 Data Protector VMware (Legacy) integration

### Introduction

This chapter explains how to configure and use the Data Protector VMware (Legacy) integration (**VMware (Legacy) integration**). Data Protector integrates with VMware vSphere or Virtual Infrastructure, including ESX Server and ESXi Server<sup>2</sup> systems (**ESX(i) Server systems**), and VirtualCenter Server systems, to back up and restore the following **VMware objects**:

- Virtual machines
- Filesystems of virtual machines<sup>3</sup>

The following backup methods are available:

- Snapshot
- Suspend
- VCBimage<sup>4</sup>
- VCBfile<sup>4</sup>

During backup, virtual machines can be powered off (**offline backup**) or actively used (**online backup**). The **Suspend** backup method supports only offline backup.

The **Snapshot**, **Suspend**, and **VCBimage** methods are used to back up virtual machines. The **VCBfile** method is used to back up filesystems of virtual machines.

For details, see “Backup methods” (page 95).

Data Protector offers interactive and scheduled backups of the following types:

- Full
- Incremental
- Differential

For details on the backup types see, “Backup types” (page 105).

Virtual machines can be restored:

- To the original ESX(i) Server systems
  - To any Data Protector client with the `VMware Integration (Legacy)` component installed
- After such a restore, you need to manually move the restored virtual machine images to an ESX(i) Server system, using the VMware Converter.

Filesystems of virtual machines can be restored to any Windows system (physical or virtual) that has the `VMware Integration (Legacy)` component installed.

This chapter provides information specific to the VMware (Legacy) integration. For limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HP Data Protector Help*.

### Integration concepts

Data Protector integrates with VMware vSphere or Virtual Infrastructure through the Data Protector VMware (Legacy) integration agent, which channels communication between the Data Protector Session Manager and the clients in the VMware environment. The Data Protector VMware (Legacy)

2. Since Data Protector components cannot be installed on ESXi Server systems, not all backup and restore functionality is available for virtual machines running on ESXi Server systems. For details, see “ESXi Server systems” (page 92).
3. Currently, you can back up only filesystems of virtual machines running Windows. For a detailed list of guest operating systems supported for the VCBfile backup method, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.
4. This method requires a Windows system with the VMware Consolidated Backup software installed (backup proxy system).

integration agent communicates with the VMware vSphere or Virtual Infrastructure through VI SDK, a web-service API.

## Supported environments

Data Protector supports environments where ESX(i) Server systems are managed through a VirtualCenter Server system (**VirtualCenter environments**) as well as environments with standalone ESX Server systems (**standalone ESX Server environments**). Mixed environments, in which some of the ESX(i) Server systems are managed through a VirtualCenter Server system and some are standalone, are also supported. You can even have multiple VirtualCenter Server systems in your environment, each managing its own set of ESX(i) Server systems.

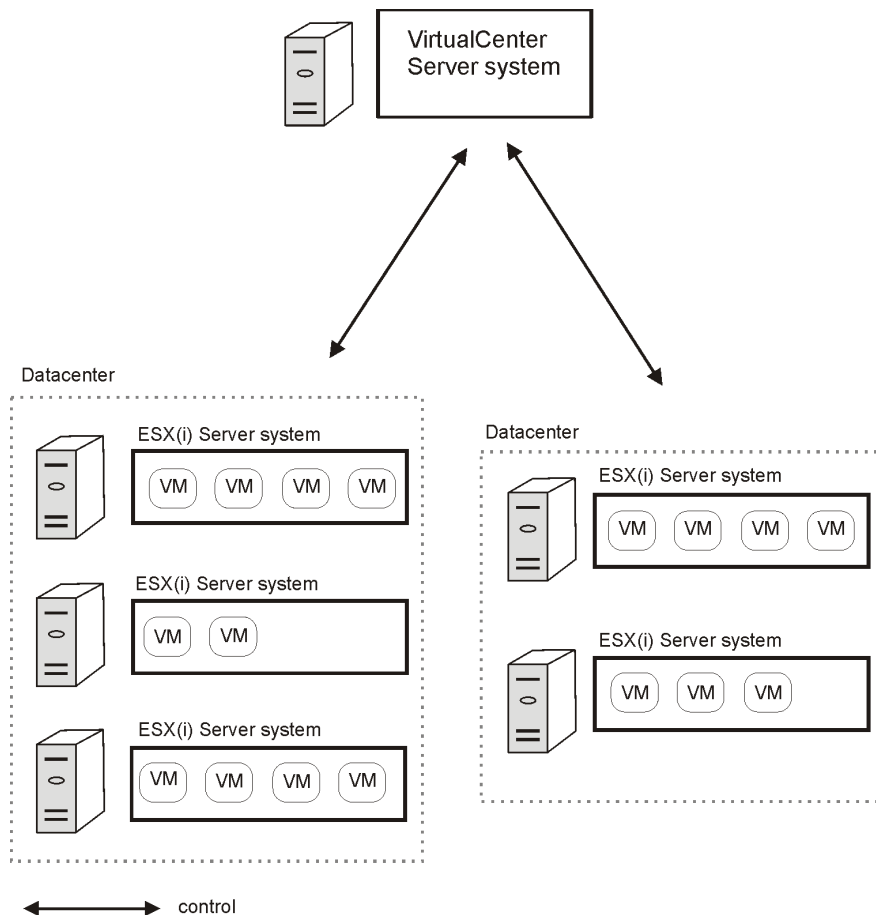
**NOTE:** Since Data Protector installation components cannot be installed on ESXi Server systems, not all backup and restore functionality is available for virtual machines running on ESXi Server systems. For details, see [“ESXi Server systems” \(page 92\)](#).

### VirtualCenter environment

In a VirtualCenter environment, Data Protector communicates with the VMware vSphere or Virtual Infrastructure through the VirtualCenter Server system. All backup and restore requests are sent there.

In one session, you can back up virtual machines from only one datacenter.

**Figure 47 VirtualCenter environment**




ESX(i) Server system	VMware platform capable of hosting multiple virtual machines.
VM	Virtual machine. Virtualized x86 or x64 PC environment, in which a guest operating system and associated application software can run.
Datacenter	An organizational unit that consists of one or more ESX(i) Server systems and the related storage for virtual machines (datastores). Datastores can reside on local disks/RAID, iSCSI or SAN storage.

## Migration of virtual machines

In a VirtualCenter environment, Data Protector supports migration of virtual machines between ESX(i) Server systems that belong to the same datacenter. Virtual machines migrate from one ESX(i) Server system to another for various reasons:

- If ESX(i) Server systems are configured in a **VMware high availability cluster**, virtual machines automatically migrate when the original ESX(i) Server system fails.
- If ESX(i) Server systems are configured in a **VMware load balancing cluster**, virtual machines automatically migrate to ESX(i) Server systems with less workload.
- You can start a migration of a virtual machine manually, using the VMware vSphere or Virtual Infrastructure client.

**Figure 48 Migration of virtual machines**



Whatever the reason for the migration, you do not need to create a new backup specification afterwards. Data Protector will automatically find the migrated virtual machines and back them up.

Data Protector does not support migration of virtual machines between ESX(i) Server systems that belong to different datacenters.

## VirtualCenter Server system in a cluster

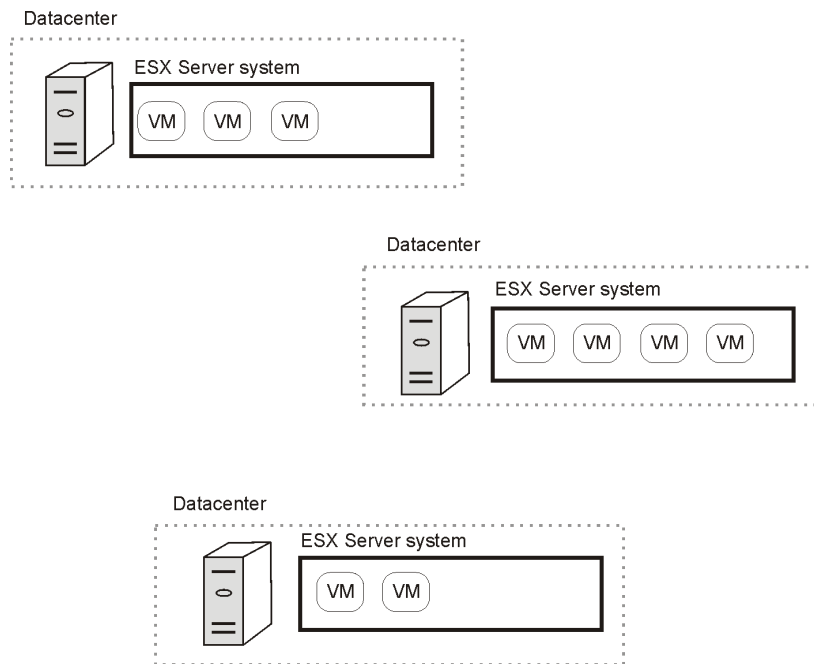
Data Protector also supports environments with a VirtualCenter Server system in a Microsoft Cluster Service cluster. After a failover in such a cluster, you do not need to change the backup specification. However, if the failover occurs during a backup or restore session, the session fails and has to be restarted.

## Standalone ESX Server environment

In a standalone ESX Server environment, Data Protector communicates with the VMware vSphere or Virtual Infrastructure through an ESX Server system. All backup and restore requests are sent there.

In one session, you can back up virtual machines from only one datacenter (ESX Server system).

**Figure 49 Standalone ESX Server environment**



## ESXi Server systems

ESXi Server systems are supported only in VirtualCenter environments; they are not supported in standalone environments. The reason is that Data Protector installation components cannot be installed on ESXi Server systems. Therefore, to be able to back up virtual machines from an ESXi Server system, the ESXi Server system must be managed through a VirtualCenter Server system. For details, see [“VirtualCenter environment” \(page 90\)](#).

### Backup

When you create a backup specification, select the VirtualCenter Server system in the **Client** option and the **VCBfile** or **VCBimage** backup method. See [Step 4](#).

---

**NOTE:** The **Snapshot** and **Suspend** backup methods are not supported; they cannot be used to back up virtual machines running on ESXi Server systems.

---

### Restore

During restore, select the VirtualCenter Server system in the **Restore client** option. See [“Restore options \(Snapshot, Suspend, VCBimage\)” \(page 125\)](#) and [“Restore options \(VCBfile\)” \(page 126\)](#).

---

**NOTE:** Since the Data Protector VMware Integration (Legacy) component cannot be installed on ESXi Server systems, virtual machines backed up from ESXi Server systems can be restored back to the original ESXi Server systems only if there is an ESX Server system (with the VMware Integration (Legacy) component installed) in the same datacenter, having access to the virtual machine datastores. Alternatively, the virtual machines can first be restored to a directory and then moved to the ESXi Server systems using the VMware Converter.

---

## Limitations in VMware vSphere environments

In VMware vSphere environments, virtual machines cannot be backed up online using the **Snapshot** backup method. To back up the virtual machines online, use the **VCBimage** or **VCBfile** backup method, or use the Data Protector Virtual Environment integration. For details on the latter, see “Data Protector Virtual Environment integration” (page 20).

## Data Protector installation components

### Data Protector Cell Manager

The Data Protector Cell Manager can be installed on a virtual machine, VirtualCenter Server system, backup proxy system, or a separate system outside the VMware vSphere or Virtual Infrastructure environment.

- ❗ **IMPORTANT:** During VCB backup sessions, the performance of the backup proxy system may seriously degrade. If the same system is the Data Protector Cell Manager, other sessions running in the Data Protector cell may be considerably affected as well.

### Data Protector VMware Integration (Legacy) component

The Data Protector VMware Integration (Legacy) component must be installed on the following clients:

- All ESX Server systems from which you plan to back up virtual machines
- VirtualCenter Server systems (if they exist)
- Backup proxy systems (if you plan to use the **VCBfile** and **VCBimage** backup methods)
- Windows systems (physical or virtual) to which you plan to restore filesystems of virtual machines

The component consists of the following parts:

- `vmware_bar.exe`, activated during backup and restore
- `util_vmware.exe`, activated during configuration and mounting on backup proxy systems

### Data Protector Media Agents

Data Protector Media Agents can be installed on ESX Server systems, VirtualCenter Server systems, backup proxy systems, or separate systems outside the VMware vSphere or Virtual Infrastructure environment.

## Backup concepts

Using the Data Protector VMware (Legacy) integration, you can back up the following VMware objects:

- Virtual machines
- Filesystems of virtual machines

## What is backed up?

Data Protector identifies datacenters and virtual machines by their VMware vSphere or Virtual Infrastructure inventory path. A standalone ESX Server system has only one datacenter `/ha-datacenter` and only one folder `/vm`, which stores all the virtual machines.

*Example:*

Datacenter: `/ha-datacenter`

Virtual machine: `/vm/myvm1`

In a VirtualCenter environment, you can organize virtual machines and datacenters within folders that you create yourself. However, once you move a datacenter or a virtual machine to another folder, you need to create a new backup specification because Data Protector no longer finds the datacenter or virtual machine under the specified path.

*Example:*

Virtual machine: `/vm/myfolder1/myfolder2/.../myvm2`

Datacenter: `/myfolder/mydatacenter`

Depending on the backup method that you select, you can back up either a virtual machine or a filesystem of a virtual machine.

### Virtual machines

Virtual machines can be backed up using the **Snapshot**, **Suspend**, or **VCBimage** backup method. A virtual machine backup contains the following files:

`vmx`

Virtual machine configuration file.

`vm disk`

Virtual machine disk files. There are two types: metadata files and binary extent files. When you create a snapshot of a virtual machine, all the subsequent changes made to the virtual machine are recorded to new metadata and binary extent delta files. A separate metadata and binary extent delta file is created for each virtual machine disk.

`vm snapshot/vm suspend`

Snapshot or suspend memory file (not applicable for the **VCBimage** backup method).

`vm snapshot description`

Snapshot description file (not applicable for the **VCBimage** backup method).

`nvram`

Non-volatile RAM file, describing the BIOS of the virtual machine.

### Virtual machine disks

Data Protector supports backup of individual virtual machine disks. In this case, all the virtual machine files are backed up, except the `vm disk` files of virtual machine disks that are excluded from backup. You can run full, incremental, and differential backups. However, after you restore individual virtual machine disks, you may need to follow additional steps to fully recover the virtual machine. For details, see [“Recovering virtual machines manually”](#) (page 130).

### Filesystems

Filesystems of virtual machines can be backed up using the **VCBfile** backup method. Currently, you can back up filesystems of virtual machines running Windows. This backup type is similar to the common filesystem backup, enabling you to select individual files and folders.

## Backup methods

Data Protector offers four different backup methods:

- **Snapshot**—see “Snapshot method” (page 95)
- **Suspend**—see “Suspend method” (page 101)
- **VCBimage**—see “VCBimage method” (page 102)
- **VCBfile**—see “VCBfile method” (page 103)

**Table 16 Backup method overview**

Backup method	How is backup consistency achieved?	Needs a backup proxy system	Supported backup types			Backs up all snapshot branches (including the user-created)
			Full	Diff	Incr	
Snapshot	Data Protector creates a virtual machine snapshot.		✓	✓	✓	✓
Suspend	Data Protector suspends the virtual machine.		✓	✓	✓	✓
VCBimage	The backup proxy system that is involved in the backup creates a virtual machine snapshot.	✓	✓			
VCBfile		✓	✓	✓	✓	


For details on the supported backup types, see “Backup types” (page 105).

### Snapshot method

A virtual machine snapshot is an operation that puts the virtual machine into a consistent state. As a result, all subsequent changes made to the virtual machine disks are recorded to separate files. Note that the snapshot operation is not supported by all virtual machine disks. For details, see the VMware documentation.

During a **Snapshot** backup, Data Protector first creates a snapshot and then copies the consistent state to Data Protector media. Snapshots created by Data Protector (**DP snapshots**) are distinguished from other snapshots by the label `_DP_SNAP_` and a description that contains the product name and a timestamp. Therefore, avoid using this label for snapshots that you create for other purposes.

**Figure 50 Snapshot tree**



Existing virtual machine snapshots reduce the overall performance of a virtual machine. For this reason, Data Protector removes DP snapshots once they are no longer needed. The number of DP

snapshots that remain on the snapshot tree depends on the selected snapshot handling mode and backup type. The following snapshot handling modes are available:

- **Disabled:** This mode supports only full backups. The snapshot that is made at the beginning of a backup session is used only to create a consistent state. After the data transfer completes, the snapshot is removed. For details, see [“Snapshot mode: disabled” \(page 96\)](#).
- **Single:** This mode supports full, differential, and incremental backups. The following backup chains are supported:  
Full, differential, differential, differential,...  
Full, incremental, incremental, incremental,...  
It means that you cannot mix incremental and differential backups within the same backup chain. The snapshot that is made at the beginning of a backup session is used to create a consistent state. After the backup completes, one DP snapshot remains on the snapshot tree. It is needed to track changes made since the last full or incremental backup. For details, see [“Snapshot mode: single” \(page 97\)](#).
- **Mixed:** This mode supports full, differential, and incremental backups. All backup chains are supported. For example:  
Full, incremental, incremental, differential, incremental, differential,...  
The snapshot that is made at the beginning of a backup session is used to create a consistent state. After the backup completes, up to two DP snapshots remain on the snapshot tree. One is needed to track changes made since the last full backup and the other to track changes made since the last backup (incremental or differential). For details, see [“Snapshot mode: mixed” \(page 99\)](#).



#### **IMPORTANT:**

A backup chain gets broken if you do any of the following operations:

- Delete a snapshot
- Revert to a snapshot
- Create a non-Data Protector snapshot
- Change snapshot handling mode
- Add a new virtual machine disk or rename an existing one
- Restore the virtual machine

After you make such changes, you must run a full backup again to start a new backup chain. Otherwise, subsequent incremental and differential sessions fail.


---

#### **Snapshot mode: disabled**

A full backup in **Disabled** mode progresses as shown in the following figure.



**Figure 51 Full backup (disabled mode)**




**Table 17 Legend**

T <sub>n</sub>	Boxes denoted by T <sub>i</sub> show how the virtual machine snapshot tree changes in time.
Virtual machine base	The rectangle denoted by Virtual machine base represents the virtual machine base or the last non-Data Protector created virtual machine delta file on the active branch.
VM <sub>Δ</sub>	A rectangle denoted by VM <sub>Δ</sub> represents a virtual machine delta file created by a Data Protector snapshot.
DP snapshot	A rectangle denoted by DP snapshot represents a process (snapshot operation triggered by Data Protector). This process closes the current active state to become a read-only file. At the same time, it creates a new delta file, which becomes the active state. The active state is denoted by round corners.

### Snapshot mode: single

A full backup in **Single** mode progresses in the same way as a full backup in **Disabled** mode, with the exception that the DP snapshot is not removed at the end (you end up with one DP snapshot). A subsequent differential backup is shown in the following figure.


**Figure 52 Differential backup (single mode)**



DP snapshot1 remains on the snapshot tree to track changes made since the last full backup.

A backup chain consisting of a full backup that is followed by incremental sessions progresses in the same way, with the exception that, at the end of an incremental session, DP snapshot1 is removed instead of DP Snapshot2 (see [Figure 53 \(page 99\)](#)).

**Figure 53 Incremental backup (single mode)**



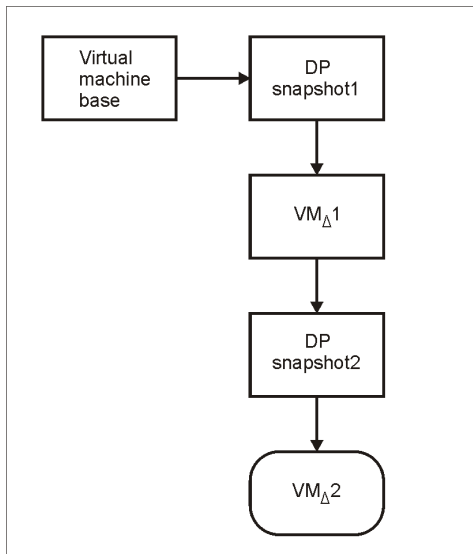
DP snapshot2 remains on the snapshot tree to track changes made since the last incremental backup.

### Snapshot mode: mixed

A full backup in **Mixed** mode progresses in the same way as a full backup in **Single** mode (you end up with one DP snapshot). A subsequent differential or incremental backup progresses in the same way as a differential or incremental backup in **Single** mode, with the exception that no DP snapshot is removed at the end (you end up with two DP snapshots).

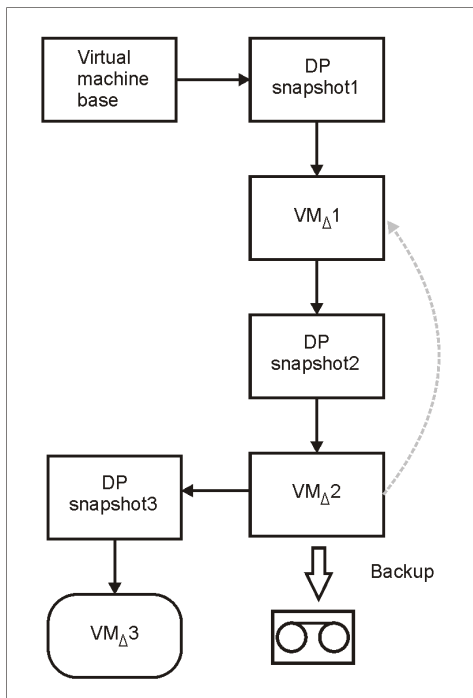
The progress of a subsequent incremental backup is shown below.

T<sub>0</sub>




START

T<sub>1</sub>



1. A new snapshot is created (DP snapshot3).
2. Changes made between DP snapshot2 and DP snapshot3 are backed up.
3. DP snapshot2 is removed (VM<sub>Δ</sub>2 is committed to the parent file VM<sub>Δ</sub>1).

**Figure 54 Incremental backup (mixed mode)**



DP snapshot1 and DP snapshot3 remain on the snapshot tree to track changes made since the last full and the last backup respectively.

A subsequent differential backup progresses in the same way as an incremental session described in “[Incremental backup \(mixed mode\)](#)” (page 101) with the following exceptions:

- DP snapshot2 is removed before DP snapshot3 is created.
- Instead of changes made between DP snapshot2 and DP snapshot3, changes made between DP snapshot1 and DP snapshot3 are backed up.


### [Suspend method](#)

The suspend operation saves the current memory state of a virtual machine to a file and puts the virtual machine offline. This functionality is similar to the Hibernate power mode in Windows operating systems. Note that the suspend operation is supported by all virtual machines. Therefore, this is the only backup method that can be used if your virtual machines do not support the snapshot operation.

During a **Suspend** backup, Data Protector suspends the virtual machine (if it is online) and then copies the virtual machine files to Data Protector media. If specified, the memory state is also backed up. After the backup, the virtual machine is powered on (if it was online), the memory is read from the memory file, and the virtual machine resumes the original state.

Using this method, you can run full, incremental, and differential backups. Data Protector uses the file modification time as an incremental or differential backup criterion. It means that only the files with the modification time changed since the last (full) backup are backed up.

**Figure 55 Snapshot and Suspend backup methods**



### VCBimage method

For the **VCBimage** backup method, you need to have at least one backup proxy system configured in your environment. A backup proxy system is a Windows system that has the VMware Consolidated Backup (**VCB**) software installed. For details on the VCB software, see the VMware documentation.


During a **VCBimage** backup, Data Protector invokes VCB to mount virtual machines on a backup proxy host. Before a virtual machine is mounted, VCB creates a virtual machine snapshot to put the virtual machine into a consistent state. Once mounted, the virtual machine is copied (exported) to the backup proxy local disk.

**NOTE:** Depending on virtual machine disk sizes, the copy operation can be very time-consuming. It may take longer than the default Data Protector Session Manager timeout, which is 10 minutes. If the timeout is reached, the session is automatically aborted. To solve the problem, extend the timeout by resetting the Data Protector `SmWaitForFirstBackupClient` global option. For details of how to set the option, see the *HP Data Protector Help* index: "global options".

After the virtual machine copy is created, it is transferred to Data Protector media. At the end, the virtual machine is dismounted. Consequently, the virtual machine copy is removed from the backup proxy. The virtual machine snapshot is removed as well.

The **VCBimage** backup method backs up only the current state of a virtual machine. Information about the snapshot tree and the changes made on non-active snapshot branches are not included in the backup.

**Figure 56 VCBimage backup method**




### VCBfile method

For the **VCBfile** backup method, you need to have at least one backup proxy system configured in your environment. A backup proxy system is a Windows system that has the VMware Consolidated Backup (VCB) software installed. For details on the VCB software, see the VMware documentation.

During a **VCBfile** backup, Data Protector invokes VCB to mount NTFS filesystems of Windows virtual machines on a backup proxy host. Before a filesystem is mounted, VCB creates a virtual machine snapshot to put the files into a consistent state. Once the filesystem is mounted, the files are transferred directly to Data Protector media while the backup proxy is only referencing them. At the end, the filesystem is dismounted and the virtual machine snapshot is removed.

The **VCBfile** backup method enables you to back up NTFS filesystems of virtual machines running Windows. Filesystems of other guest operating systems cannot be backed up.

Figure 57 VCBfile method





## Backup types

The integration provides the following backup types:

**Table 18 Backup types**

Full	<b>Snapshot and Suspend:</b> Backs up the complete virtual machine, including the virtual machine snapshot tree and the virtual machine memory file (if specified). <b>VCBimage:</b> Backs up the complete virtual machine. The virtual machine snapshot tree and the virtual machine memory file are not included. <b>VCBfile:</b> Backs up all selected files and folders.
Incr	<b>Snapshot and Suspend:</b> Backs up changes made to the virtual machines since the last backup of any type. <b>VCBimage:</b> Not available. <b>VCBfile:</b> Backs up only those files and folders that have changed since the last backup of any type.
Differential	<b>Snapshot and Suspend:</b> Backs up changes made to virtual machines since the last full backup. <b>VCBimage:</b> Not available. <b>VCBfile:</b> Backs up only those files and folders that have changed since the last full backup.

How Data Protector identifies what to back up during incremental and differential backups depends on the selected backup method:

- **Snapshot:** Data Protector uses DP snapshots to identify the changes.
- **Suspend and VCBfile:** Data Protector checks file modification time to identify the changes.

## Considerations

- **Disk space:** Virtual machine operations that are performed during backup require additional disk space on the datastores. Data Protector checks for each virtual machine or filesystem separately whether the required virtual machine operation can be safely performed (whether enough disk space is available). If not, the backup of that particular virtual machine or filesystem is skipped.

**Table 19 Disk space requirements**

Backup method	Required disk space on datastores	Explanation
Snapshot	The sum of the sizes of all the virtual machine disks, plus the size of the virtual machine RAM if specified.	After a snapshot is created, changes made to the virtual machine disks are recorded to separate files (one delta file is created for each virtual machine disk). A delta file can grow up to the total size of the virtual machine disk.
Suspend	The size of the virtual machine RAM.	The Suspend method does not create any delta files. However, a copy of the memory state is always made, even if the memory file is not selected for backup.
VCBimage	The sum of the sizes of all the virtual machine disks.	Same as for the Snapshot method.
VCBfile		

For the **VCBimage** and **VCBfile** backup methods, disk space is needed also on the backup proxy system for mounting virtual machines and filesystems. VCB checks whether enough disks space is available and informs Data Protector of it. If not, Data Protector skips the backup of that particular virtual machine or filesystem.

To disable the check for available disk space that is performed during backup of a particular datacenter, set the Data Protector omnirc option `OB2_VMWARE_SKIPSPACECHECK` on the backup proxy system to 1.

For details of how to set the option, see the *HP Data Protector Help* index: “omnirc options” and “Setting the environment variables using GUI” (page 106).

### Using the VCBImage backup method in optimized mode

If you use the **VCBImage** backup method in optimized mode, the virtual machine disk files copied to the backup proxy server are optimized. Unused portions of the virtual machine data are skipped, resulting in smaller and shorter backup.

To use the **VCBImage** backup method in optimized mode, set the Data Protector omnirc option `OB2_VMWARE_VCBOPTIMISED` on the backup proxy system to 1.

For details of how to set the option, see the *HP Data Protector Help* index: “omnirc options” and “Setting the environment variables using GUI” (page 106).

### Setting the environment variables using GUI

To set the variable using the Data Protector GUI:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **VMware (Legacy)**. Click a backup specification for selected VMware datacenter.

---

**NOTE:** If also configured the Data Protector Virtual Environment integration, right-click **Virtual Environment**.

---

3. In the Source property page, right-click the VMware datacenter and select **Set Environment Variables....**
  4. In the Advanced dialog box, add the variable (for example, `OB2_VMWARE_SKIPSPACECHECK`, or `OB2_VMWARE_VCBOPTIMISED`) and set its value (for example, to 1). Apply the change by clicking **OK**.
- **Concurrent sessions:** Backup sessions that use the same devices or back up the same datacenter cannot run concurrently. If multiple sessions are started, one session waits for the other to complete.

## Restore concepts

You can restore virtual machines and filesystems of virtual machines using the Data Protector GUI or CLI.

### Restore of virtual machines

Virtual machines backed up with the **Snapshot**, **Suspend**, or **VCBImage** method can be restored in two ways, depending on whether you select the **Restore to datacenter** or **Restore to default directory** option.

#### Restore to datacenter

If you select the **Restore to datacenter** option, virtual machines are restored to the original ESX(i) Server systems. For the restore, you can also specify the following:

- Whether existing virtual machine files should be overwritten with files from the backup or left intact.
- Whether the restored virtual machine files should be consolidated when the restore completes.
- Whether the restored virtual machines should be registered in the datacenter, and powered on.

Usually, you specify the original datacenter in the **Restore to datacenter** option. Specifying a different datacenter is allowed only in the case of a disaster recovery when the new datacenter name differs from the original one. Do not restore a virtual machine to a different datacenter in

order to clone the virtual machine because this leads to licensing and network problems (mostly because the virtual machine UUID gets duplicated).

### Restore to default directory

If you select the **Restore to default directory** option, virtual machines are restored to the client specified in the **Restore client** option. This destination client can be any client with the VMware Integration (Legacy) component installed, but typically it is a VCB proxy system. During such a restore, the virtual machine files are restored to the directory specified in the destination client `omnirc` option `OB2_VMWARE_PATH` (for example, `/tmp/`). Therefore, before you start the restore session, ensure that the `omnirc` option `OB2_VMWARE_PATH` is set. For details of how to set `omnirc` options, see the *HP Data Protector Help* index: "omnirc options".

After such a restore (**restore outside a datacenter**), the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX(i) Server system, using the VMware Converter as described in ["Recovery after restoring outside a datacenter"](#) (page 130).

- ❗ **IMPORTANT:** If the connection with the target ESX(i) Server system cannot be established, the virtual machines are automatically restored to the client specified in the **Restore client** option as if **Restore to default directory** had been selected (see ["Restore to default directory"](#) (page 107)). However, if the `omnirc` option `OB2_VMWARE_PATH` on the restore client has not been set, the restore fails.

### Restore of filesystems

Filesystems of virtual machines backed up with the **VCBfile** method can be restored to the original virtual machine or to any Windows system (physical or virtual) that has the VMware Integration (Legacy) component installed. You can also specify what to do if the files to be restored already exist on the destination client: you may keep or overwrite them.

### Restore considerations

- **Restore chain:** When you restore a virtual machine from an incremental or differential **Snapshot** or **Suspend** session, Data Protector automatically restores the complete backup chain, starting with the last full backup, which is then followed by the last differential and all subsequent incremental backups (if they exist) up to the selected session.
- **Backup proxy:** A restore session from a **VCBfile** or **VCBimage** backup does not involve a backup proxy system. VMware objects are restored directly to specified clients.
- **Concurrent sessions:** Restore sessions that use the same devices or restore the same datacenter cannot run concurrently.
- **Different backup methods:** VMware objects that were created using different backup methods cannot be restored in the same session.

### Restoring an optimized VCBimage backup session

Before restoring data backed up using the **VCBimage** backup method in optimized mode, ensure that the Data Protector `omnirc` options `OB2_VMWARE_VCBOPTIMISED` and `OB2_VMWARE_PATH` are set on the restore client. The optimized VCB files can only be restored with the `vcbRestore` utility. For details, see the VMware documentation.

Set the `OB2_VMWARE_VCBOPTIMISED` option to 1 and the `OB2_VMWARE_PATH` option to a specified directory on a destination client where the virtual machine files will be restored (for example, `/tmp/`).

For details of how to set the `omnirc` options, see the *HP Data Protector Help* index: "omnirc options".

- 
- ❗ **IMPORTANT:** Since Data Protector cannot differentiate between optimized and non-optimized **VCBimage** backup sessions from the backup session information, you have to decide whether the restore should be done in optimized mode. It is possible to restore either optimized or non-optimized **VCBimage** backup sessions by setting the `OB2_VMWARE_VCBOPTIMISED` and `OB2_VMWARE_PATH` options.

However, note that restoring a non-optimized backup session by setting these options will take longer, because after restoring the files to the temporary location, the `vcbRestore` utility is used to convert the disk files.

---

## Configuring the integration

Configure the integration as follows:

- Configure users as described in “Configuring VMware users” (page 109).
- Provide Data Protector with login information to VMware management clients as described in “Configuring VMware management clients” (page 109).
- If your ESX Server systems are configured in a cluster, check the cluster settings as described in “Configuring clusters” (page 109).
- For each virtual machine, specify details of how to perform various backup methods as described in “Configuring virtual machines” (page 112).

## Prerequisites

- Ensure that you have a correctly installed and configured VMware environment.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.
  - For information on installing, configuring, and using the VMware vSphere or Virtual Infrastructure, see the VMware documentation.  
For the **VCBfile** and **VCBimage** backup methods, ensure that you have at least one backup proxy system configured in your environment. For details, see the VMware documentation.
- Ensure that you have correctly installed Data Protector. For information on how to install Data Protector in various architectures, see the *HP Data Protector Installation and Licensing Guide*.
- For limitations, see “Limitations and recommendations” in the *HP Data Protector Product Announcements, Software Notes, and References*.

## Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the VMware vSphere or Virtual Infrastructure and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every VMware client (ESX Server system, VirtualCenter Server system, backup proxy system) in your environment.

## Configuring VMware users

1. Identify or configure the following users:

**Table 20 VMware users**

VirtualCenter users (VirtualCenter environment)	For each VirtualCenter Server system, identify the Windows operating system user who administers the VirtualCenter Server.
ESX Server users (standalone ESX Server environment)	For each standalone ESX Server system, identify an operating system user who has read, write, and execute permissions on the related datastores.

2. Add all these users to the Data Protector `admin` or `operator` user group.  
For details on adding users, see the *HP Data Protector Help* index: "adding users".

**NOTE:** If your VirtualCenter Server system is in a cluster, add the users from both nodes to the Data Protector `admin` or `operator` user group.

## Configuring clusters

### ESX Server systems in a cluster

If your ESX Server systems are configured in a high availability cluster, open the cluster settings dialog box in the VMware vSphere or Virtual Infrastructure client and select the option **Allow virtual machines to be powered on even if they violate availability constraints**.

**NOTE:** This option is applicable only for the **Suspend** backup method.

If the option is not selected, you may encounter problems when trying to back up virtual machines that have migrated to other ESX Server systems after a failover. More specifically, Data Protector may not be able to power the virtual machines on at the end of the session.

## Configuring backup proxy systems

If your virtual machines reside on **iSCSI datastores**, restart the Data Protector Inet service on the related backup proxy system under a network domain user account that has read-write permissions for the following directories:

- C:\Program Files
- C:\Windows\Temp\vmware-system (this directory is created when you run the vcbMounter for the first time)

**NOTE:** If the Data Protector Inet service on the backup proxy system runs under the Local System account, **VCBimage** and **VCBfile** backup sessions fail because this account does not have permissions to use the network, which is needed to access iSCSI datastores.

For information on changing the Data Protector Inet account, see the *HP Data Protector Help* index: "Inet, changing account".

## Configuring VMware management clients

In VirtualCenter environments, Data Protector communicates with the VMware vSphere or Virtual Infrastructure through a VirtualCenter Server system and in standalone ESX Server environments, through an ESX Server system. From now on, this client will be called the **VMware management client**.

For each VMware management client, you need to provide the following login information:

- Username
- Password
- Web root (optional)
- Port (optional)

For details on the parameters, see the following section.

You can provide the information in two different ways:

- Specify the login information manually (**Standard security**). In this case, the information is saved in the VMware management client specific configuration file on the Cell Manager. The file is named *VMwareManagementClient%\_OB2\_GLOBAL*.

---

**NOTE:** Before the information is saved, Data Protector first tests the connection. If the connection fails, the information is not saved in the configuration file, leaving the previous login information (if it exists) in the configuration file intact.

---


- Instruct Data Protector to use the login information contained in a local file on the VMware management client (**Integrated security**). At the beginning of a backup session, the Session Manager starts the VMware (Legacy) integration agent on the VMware management client and the agent provides the information from the file.

To configure VMware management clients, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

You specify the login information in the Configure VMware dialog box. The dialog box is displayed automatically when you create the first backup specification. To change or check the parameters later, open any backup specification for this VMware management client, go to the Source page, right-click the client at the top, and click **Configure**.

**Figure 58 Configuring a VirtualCenter Server system**



Select either **Integrated security** or **Standard security**.

If you select **Standard security**, specify the following:

- **Username and Password:** Specify an operating system user account that has the following privileges:

*VMware Virtual Infrastructure:*

- `System.View`
- `System.Anonymous`
- `Folder.Create`
- `VirtualMachine.State.CreateSnapshot`
- `VirtualMachine.State.RemoveSnapshot`
- `VirtualMachine.Interact.Suspend`
- `VirtualMachine.Interact.PowerOff`
- `VirtualMachine.Interact.PowerOn`
- `VirtualMachine.Inventory.Create`
- `VirtualMachine.Inventory.Delete`

For VMware vSphere privileges, see the corresponding section in [“Data Protector Virtual Environment integration” \(page 20\)](#) ([“Importing and configuring VMware clients” \(page 42\)](#)).

- **Web service:** Optionally, change the web service entry point URI. Default: `/sdk`
- **Port:** Optionally, change the TCP port number of the VMware vSphere or Virtual Infrastructure web service server.

Default: 443 (SSL-encrypted HTTP), 80 (unencrypted HTTP).

By default, the HTTP/S (SSL-encrypted HTTP) is used. To switch to unencrypted HTTP, configure the VMware management client to allow HTTP connections and set the Data Protector `omnirc` option `OB2_VMWARE_HTTP` to 1. For information of how to set the `omnirc` option, see the *HP Data Protector Help* index: “`omnirc` options”.

If you leave the **Port** option empty, the value is read from the following file, depending on your VMware management client:

**VirtualCenter Server system:** Windows registry: `HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware VirtualCenter\`

**ESX Server system:** `/etc/vmware/hostd/config.xml`

If you select **Integrated security**, ensure that the `backuptools.conf` file on the VMware management client contains the required data. The file location depends on your VMware management client:

**VirtualCenter Server system:** `C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\backuptools.conf`.

**Standalone ESX Server system:** `/etc/vmware/backuptools.conf`.

On the VirtualCenter Server system, the file does not exist by default. It has to be created manually at the specified location. The file content should be similar to the following:

```
USERNAME="Administrator"
PASSWORD="vnm46578"
```

## Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in “Configuring VMware users” (page 109).
2. Execute:

For **Integrated security**:

```
util_vmware.exe  
-config  
-security 1
```

For **Standard security**:

```
util_vmware.exe  
-config  
-security 0  
-user Username  
-password Password  
[-webroot WebServiceRoot]  
[-port WebServicePort]
```

The message `*RETVAL*0` indicates successful configuration.

For a description of the options, see the `util_vmware.exe` man page or the *HP Data Protector Command Line Interface Reference*.

## Configuring virtual machines

For each virtual machine, specify details of how to perform various backup methods:

- For the **Snapshot** method, specify how to handle virtual machine snapshots that are created during backup. Note that not all snapshot handling modes support incremental and differential backups.
- For the **VCBfile** and **VCBimage** backup methods, specify which backup proxy system and mount points should be used to back up virtual machines or filesystems.
- The **Suspend** backup method has no specifics.

You can configure each virtual machine separately or all together. Configuration settings for virtual machines of the same datacenter are saved in a separate configuration file on the Cell Manager. The file is named `VMwareManagementClient%DatacenterPath`. It is used for all backup sessions involving this particular datacenter.

To configure virtual machines, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

You can specify details of how a virtual machine backup should be performed when you create or modify a backup specification. In the Source page of a backup specification, right-click the client system at the top or any of the virtual machines listed below and click **Configure Virtual Machines**.



**Figure 59 Configuring virtual machines**

**Configure Virtual Machine(s)**

Settings

**Virtual machine configuration**

Here you can specify common VM settings and override them for particular virtual machines.

Configure virtual machine

/vm/capricorn1\*

☐ Use common settings for selected VM

Snapshot handling

☐ Disabled  
Only full backup is possible.

☒ Single  
Either incremental or differential possible in same backup chain.

☐ Mixed  
Incremental and differential possible in same backup chain.

Backup proxy configuration

Backup proxy: gabriel.company.com

☐ Specify mountpoint:

OK Cancel Help

In the Configure Virtual Machines dialog box, specify the following options:

**Table 21 Virtual machine options**

<b>Configure virtual machine</b>		Select whether you want to specify common virtual machine settings ( <b>Common VM Settings</b> ) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.
	<b>Use common settings for selected VM</b>	Available only if a virtual machine is selected. Select this option to apply the common virtual machine settings for the selected virtual machine. Default: selected
	<b>Use default settings</b>	Available only if <b>Common VM Settings</b> is selected. Select this option to set default values for the common virtual machine settings. Default: selected
<b>Snapshot handling</b>	<b>Disabled</b> (default)	This mode supports only full backups. The virtual machine snapshot that is created during backup is removed at the end of the session. For details, see <a href="#">“Snapshot mode: disabled”</a> (page 96).
	<b>Single</b>	This mode supports full, incremental, and differential backups. However, you cannot mix incremental and differential backups within the same backup chain.  Data Protector keeps one DP snapshot for backup purposes. For details, see <a href="#">“Snapshot mode: single”</a> (page 97).
	<b>Mixed</b>	This mode supports full, incremental, and differential backups, in all combinations.  Data Protector keeps up to two DP snapshots for backup purposes. For details, see <a href="#">“Snapshot mode: mixed”</a> (page 99).
<b>Backup proxy configuration</b>	<b>Backup proxy</b>	Select a backup proxy system to be used for the <b>VCBfile</b> and <b>VCBimage</b> backup methods. Note that Data Protector lists all the systems that have the VMware Integration (Legacy) component installed, including those that may not be backup proxy systems.
	<b>Specify mountpoint</b>	During a VCB backup session, virtual machine disks are mounted on a backup proxy system. Select this option to specify a different mount point directory on the backup proxy system. This is particularly useful for the <b>VCBimage</b> backup method, during which the virtual machine disks are copied to a local disk on the backup proxy system. This option enables you to specify a mount point where you have enough disk space. It also enables you to mount virtual machines on different disks, which improves backup performance.  Default: <code>Data_Protector_home\tmp</code>

## Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in [“Configuring VMware users”](#) (page 109).

## 2. Execute:

```
util_vmware.exe
  -config
  -instance DatacenterPath
  -vm VMpath VM_OPTIONS [-vm VMpath VM_OPTIONS...]
```

```
VM_OPTIONS
  -snapshots { 0 | 1 | 2 }
  -proxy BackupProxy
  -mount ProxyMountPoint
  -default
```

The values { 0 | 1 | 2 } represent **Disabled**, **Single**, and **Mixed** snapshot handling modes respectively. For details, see the `util_vmware.exe` man page or the *HP Data Protector Command Line Interface Reference*.

To change virtual machine specific settings back to common virtual machine settings, execute:

```
util_vmware.exe -config -instance DatacenterPath -vm VMpath -default
[-vm VMpath -default ...]
```

The message `*RETVAL*0` indicates successful configuration.



---

**TIP:** You can join the options for configuring virtual machines and the options for configuring the VMware management client in the same command line.

---

### Example

Suppose you want to set **Integrated security** for the VirtualCenter Server system `virtualcenter2.company.com`. In addition, you want to specify **Single** snapshot handling mode and use the backup proxy system `proxy2.company.com` for the virtual machine `/vm/myfolder/myvm` that belongs to the `datacenter/Mydatacenters/Datacenter1`.

To achieve all this, log in to the VirtualCenter Server system `virtualcenter2.company.com`, execute:

```
util_vmware.exe -config -security 1 -instance /Mydatacenters/Datacenter1
-vm /vm/myfolder/myvm -snapshots 1 -proxy proxy2.company.com
```

## Checking the configuration of VMware clients

To verify the connection, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

You can verify the connection to the VMware management client after you have created at least one backup specification.

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **VMware (Legacy)**. Click the backup specification for the VMware management client to be checked.

---

**NOTE:** If you also configured the Data Protector Virtual Environment integration, right-click **Virtual Environment**.

---

3. Right-click the VMware management client and click **Check configuration**.

### Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in *"Configuring VMware users"* (page 109).

2. Execute:

```
util_vmware.exe -chkconf
```

For a description of the options, see the `util_vmware.exe` man page or the *HP Data Protector Command Line Interface Reference*.

## Backup

This section contains procedures that are needed to back up virtual machines. For details on backup concepts, see [“Backup concepts” \(page 93\)](#).

### Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **VMware (Legacy)**, and click **Add Backup**.

---

**NOTE:** If you also configured the Data Protector Virtual Environment integration, right-click **Virtual Environment**.

---

3. In the Create New Backup dialog box, click **OK**.
4. In **Client**, select a VMware management client. Note that the **Client** drop-down list offers all clients that have the *VMware Integration (Legacy)* component installed. If your VMware management client is a VirtualCenter Server system that is configured in a Microsoft Cluster Service cluster, select its virtual hostname.

---

**NOTE:** If you also configured the Data Protector Virtual Environment integration, select the client that has the **(Legacy VMware)** label appended at the end.

---

If the selected VMware management client is not configured for use with Data Protector, a warning is displayed that the configuration check failed. Click **OK** to open the Configure VMware dialog box and provide the connection parameters as described in [“Configuring VMware management clients” \(page 109\)](#).

In **Application database**, select a datacenter that you want to back up from (standalone ESX Server system has only one datacenter (/ha-datacenter)).

In **Backup method**, select a backup method.

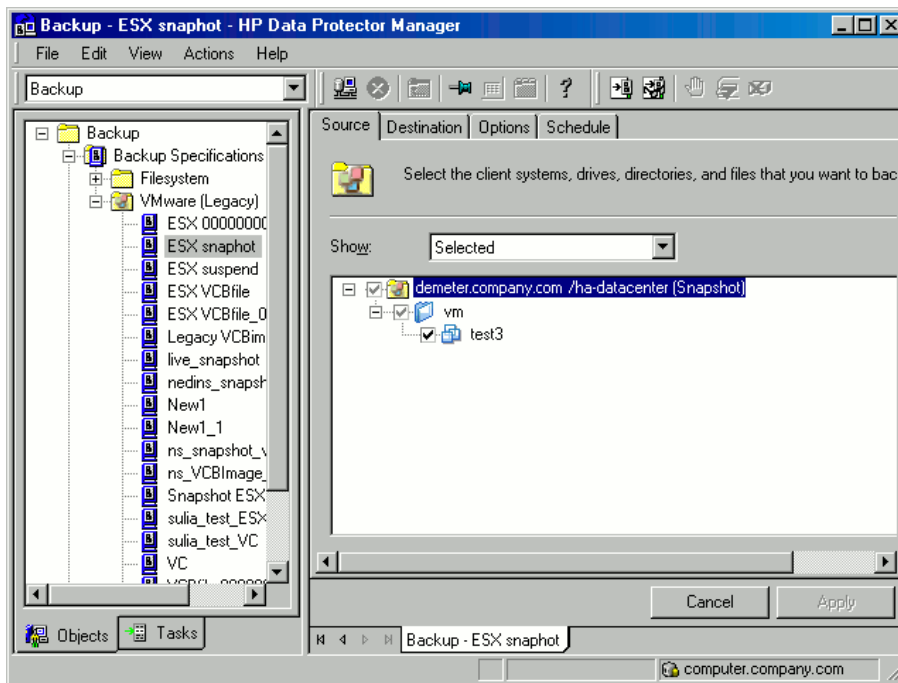
Specify the **User and group/domain** options, which are available if an ESX Server or Windows Server 2008 VirtualCenter Server system is selected:

- **Specify OS:** On ESX Server clients, it is mandatory to specify the **Username** and **Group/Domain name** options. On Windows Server 2008 VirtualCenter Server systems, it is not mandatory to specify these options and if they are not specified, the backup runs under the Local System Account.
- **Username, Group/Domain name:** Specify the operating system user account under which you want the backup session to run (for example, the user name `root`, group `root`, or user name `Administrator`, domain `DP`). Ensure that this user has been added to the Data Protector `admin` or `operator` user group. This user becomes the backup owner.  
If you selected an ESX Server client, ensure that this user also has the read, write, and execute permissions on the ESX Server datastores.

Click **Next**.

5. **Snapshot, Suspend, and VCBimage** backup methods: Select virtual machines or individual virtual machine disks that you want to back up.

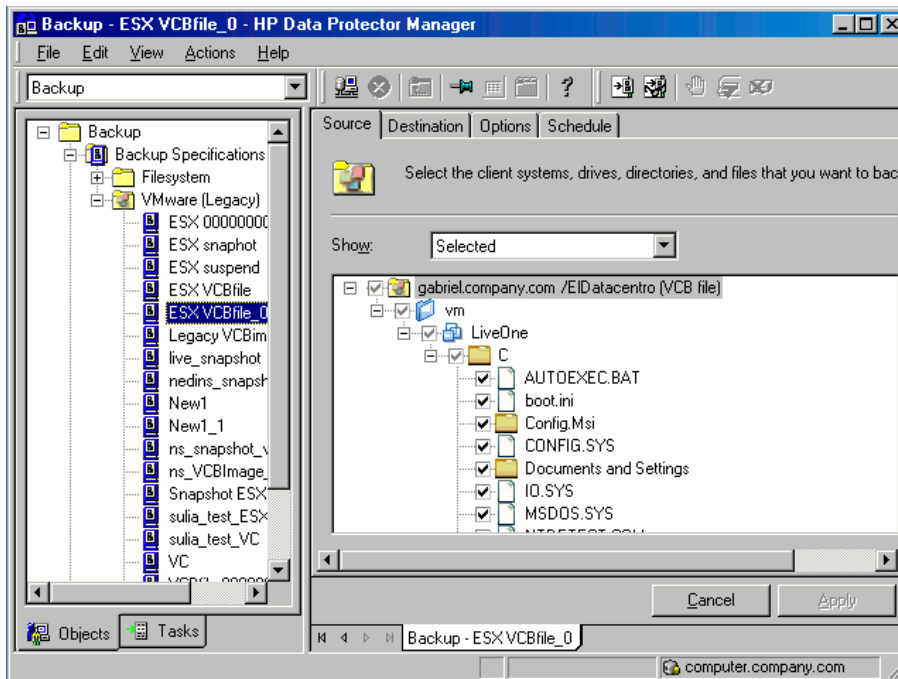
**Figure 60 Selecting VMware objects (Snapshot, Suspend, VCBimage)**



**VCBimage** backup method: To use the method in optimized mode, right-click the VMware datacenter, select **Set Environment Variables...**, and set the OB2\_VMWARE\_VCBOPTIMISED variable to 1.

For details, see [“Setting the environment variables using GUI”](#) (page 106).

**Figure 61 Selecting VMware objects (VCBfile)**



**VCBfile** backup method: Right-click a virtual machine and click **Mount filesystem** to mount virtual machine filesystems. This may take some time. Then, select the files and folders that you want to back up.

---

**NOTE:** This note applies to all backup methods: If you select the client system at the top, all virtual machines in the datacenter will be selected for backup, including the virtual machines that you create after the backup specification is saved.

---

If your virtual machines are not configured yet, right-click the client system at the top or any of the virtual machines listed below, and click **Configure Virtual Machines**. For details, see [“Configuring virtual machines” \(page 112\)](#).

Click **Next**.

6. Select devices to use for the backup.


To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and the media pool you will use.

Click **Next**.

7. Set backup options.

For information on application-specific backup options, see [“VMware backup options” \(page 119\)](#).

**Figure 62 Application-specific options**



Click **Next**.

8. Optionally, schedule the backup. See [“Scheduling backup sessions” \(page 119\)](#).

Click **Next**.

9. Save the backup specification, specifying a name and a backup specification group.



---

**TIP:** Preview your backup specification before using it for real. See [“Previewing backup sessions” \(page 120\)](#).

---

**Table 22 VMware backup options**

Option	Description
<b>Pre-exec, Post-exec</b>	<p>The command specified here is run by <code>vmware_bar.exe</code> on the VMware management client before the backup (pre-exec) or after it (post-exec).</p> <p>Do not use double quotes. Type only the name of the command and ensure that the command resides in the following directory:</p> <p><b>VirtualCenter Server system:</b> <code>Data_Protector_home\bin</code></p> <p><b>Standalone ESX Server system:</b> <code>/opt/omni/bin</code></p>
<b>Backup memory file</b>	<p>Available only for the <b>Snapshot</b> and <b>Suspend</b> backup methods.</p> <p>If this option is selected, the memory of a running virtual machine is saved into a file and backed up. Note that the backup takes considerably longer if this option is selected.</p>

## Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

## Scheduling backup sessions

You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

### Scheduling example

To schedule differential backups at 10:15, 13:00, and 18:00 during week days:

1. In the Schedule property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under Recurring, select **Weekly**. Under Time options, select **10:15**. Under Recurring options, select **Mon, Tue, Wed, Thu, and Fri**. See “[Scheduling a backup session](#)” (page 120). Under Session options, select **Differential** from the **Backup type** drop-down list.  
Click **OK**.
3. Repeat [Step 1](#) and [Step 2](#) to schedule differential backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

**Figure 63 Scheduling a backup session**

**Schedule Backup**

Specify the desired backup time, frequency, duration, and type.

**Recurring**

☐ None  
☐ Daily  
☒ Weekly  
☐ Monthly

**Time options**

Time: 10:15 AM  
☐ Use starting  
1/21/2011

**Recurring options**

Every 1 week(s) on

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

**Session options**

Backup type: Differential  
Network load: ☒ High ☐ Medium ☐ Low  
Backup protection: Permanent

OK Cancel Help

## Previewing backup sessions

Preview the backup session to test it. You can use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **VMware (Legacy)**. Right-click the backup specification you want to preview and click **Preview Backup**.

---

**NOTE:** If you also configured the Data Protector Virtual Environment integration, right-click **Virtual Environment**.

---

3. Specify the **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful preview.

### Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in [“Configuring VMware users”](#) (page 109).
2. Execute:  

```
omnib -vmware_list BackupSpecificationName -test_bar
```

## What happens during the preview?

The following are tested:

- Communication between the VMware management client and Data Protector
- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices



## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

To start a backup, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **VMware (Legacy)**. Right-click the backup specification you want to use and click **Start Backup**.

---

**NOTE:** If you also configured the Data Protector Virtual Environment integration, right-click **Virtual Environment**.

---

3. Specify **Backup type** and **Network load**. Click **OK**.

The message `Session completed successfully` is displayed at the end of a successful backup session.

### Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in [“Configuring VMware users”](#) (page 109).
2. Execute:

```
omnib -vmware_list BackupSpecificationName [-barmode  
VMwareMode] [ListOptions]
```

where *VMwareMode* is one of the following backup types:

`full|incr|diff`

The default is `full`.

For *ListOptions*, see the `omnib` man page or the *HP Data Protector Command Line Interface Reference*.

#### Examples

To start a full backup using the backup specification `snapshot2`, execute:

```
omnib -vmware_list snapshot2 -barmode full
```

To start a differential backup using the same backup specification, execute:

```
omnib -vmware_list snapshot2 -barmode diff
```

## Preparing for disaster recovery

To do a disaster recovery, you need backups of the following VMware objects:

**Table 23 What must be backed up**

VMware object	How to back it up
ESX Server console	<ol style="list-style-type: none"><li>1. Ensure that the Data Protector Disk Agent component is installed on all the ESX Server systems.</li><li>2. In the Backup context of the Data Protector GUI, right-click <b>Filesystem</b> and click <b>Add backup</b> to create a backup specification of the filesystem type. In the Source page of the backup specification, select ESX Server consoles of all ESX Server systems. For details on what to back up, see the topic “ESX Server Configuration Backup and Restore procedure” at <a href="http://kb.vmware.com/selfservice/microsites/microsite.do">http://kb.vmware.com/selfservice/microsites/microsite.do</a>.</li><li>3. Start a backup using the newly created backup specification.</li></ol>
VirtualCenter configuration database (applicable only for VirtualCenter environments)	<p>The VirtualCenter configuration database can be an Oracle database or Microsoft SQL Server database. To back up the database, use the corresponding Data Protector integration. For example, if it is an Oracle database, follow these steps:</p> <ol style="list-style-type: none"><li>1. Ensure that the Data Protector Oracle Integration component is installed on the VirtualCenter Server system.</li><li>2. In the Backup context of the Data Protector GUI, right-click <b>Oracle Server</b> and click <b>Add backup</b> to create a backup specification of the Oracle type. In <b>Application database</b>, type the name of the VirtualCenter configuration database. Continue with the backup specification creation as described in the <i>HP Data Protector Integration Guide for Oracle and SAP</i>.</li><li>3. Start a backup using the newly created backup specification.</li></ol>
VMware virtual machines	Back up the virtual machines as described in this chapter.

## Restore

This section contains procedures that are needed to restore virtual machines. For details on restore concepts, see “[Restore concepts](#)” (page 106).

## Finding information for restore

You can find information about backup objects in the Data Protector Internal Database (IDB), such as which backup type and media were used, and which messages were displayed during the backup. To retrieve this information, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

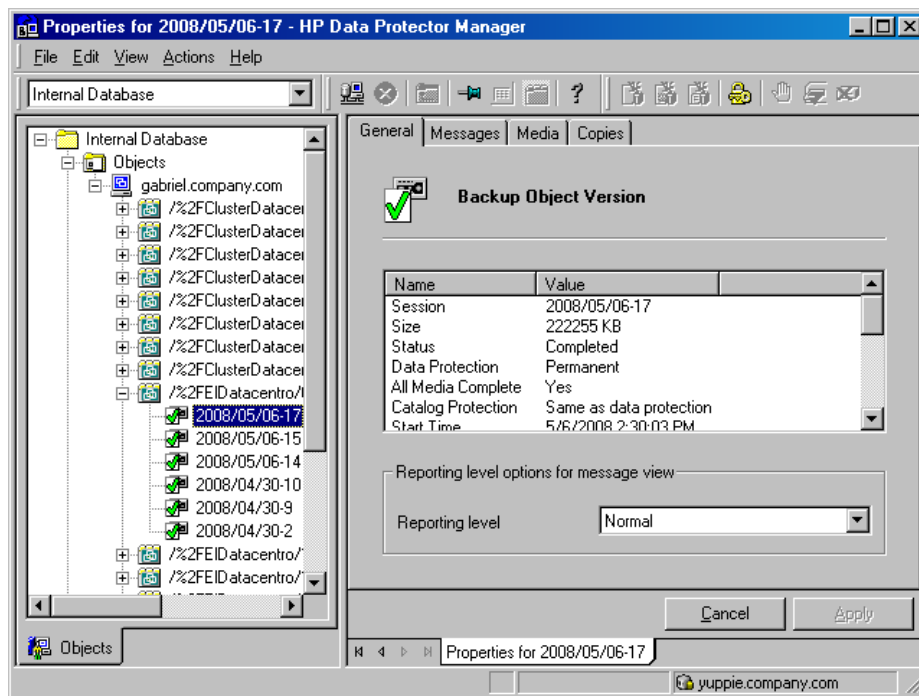
In the Internal Database context, expand **Objects** or **Sessions**.

If you expand **Objects**, backup objects are sorted according to the virtual machine for which they were created. For example, backup objects for the virtual machine /vm/mach1 from the datacenter ELDatacentro are listed under /%2FEldatacentro/0/%2Fvm%2Fmach1.

If you expand **Sessions**, backup objects are sorted according to the session in which they were created. For example, backup objects created in the session 2010/08/15-7 are listed under 2010/08/15-7.

To view details on a backup object, right-click the backup object and click **Properties**.

Figure 64 Backup object information



**TIP:** To view the messages displayed during the session, click the **Messages** tab.

## Using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in [“Configuring VMware users”](#) (page 109).
2. Get a list of VMware backup objects created in a particular backup session:

```
omnidb -session SessionID
```

3. Get details on a particular backup object:

```
omnidb -vmware BackupObjectName -session SessionID -catalog
```

Here is one example of a backup object name:

```
gabriel.company.com::/%2FEIDatacentro/0/%2Fvm%2Fharbour
```

For details, see the `omnidb` man page or the *HP Data Protector Command Line Interface Reference*.

## Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Virtual Environment** and **VMware (Legacy)**, then expand the VMware management client and click the datacenter you want to restore.
3. In the Source page, Data Protector displays VMware objects that were backed up with the method specified in **Backup method**. To display VMware objects that were backed up with a different method, change **Backup method**.

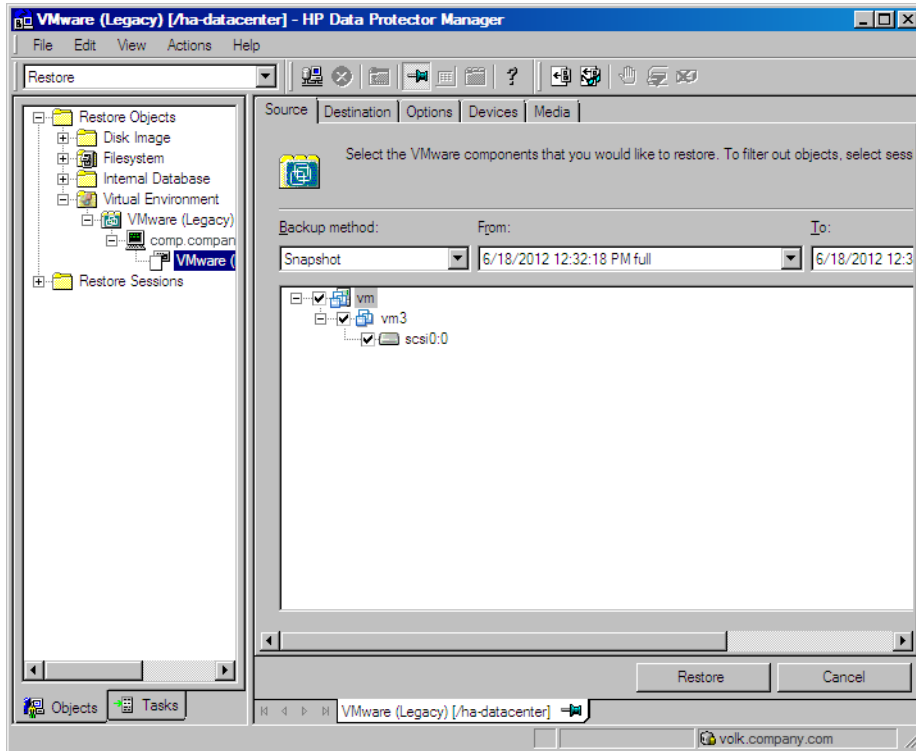
**NOTE:** VMware objects that were created using different backup methods cannot be restored in the same session.

By using the **From** and **To** options, you narrow the scope of displayed virtual machines to those backed up within the specified time interval.

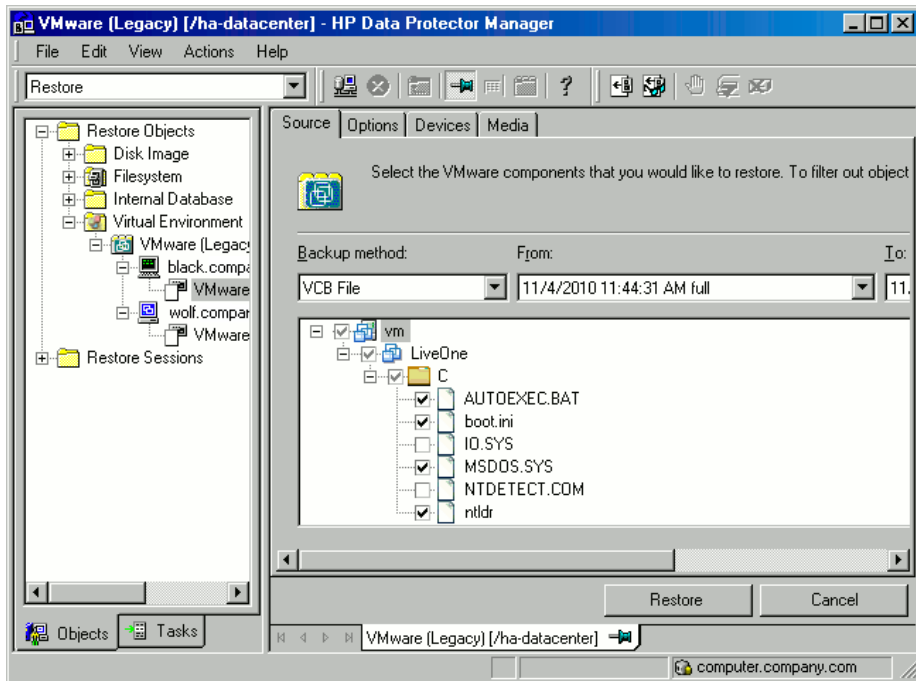
Select the VMware objects that you want to restore.

**NOTE:** Data Protector restores each selected VMware object from the last backup session performed during the specified time interval.

**Figure 65 Selecting VMware objects for restore (Snapshot, Suspend, VCBimage)**

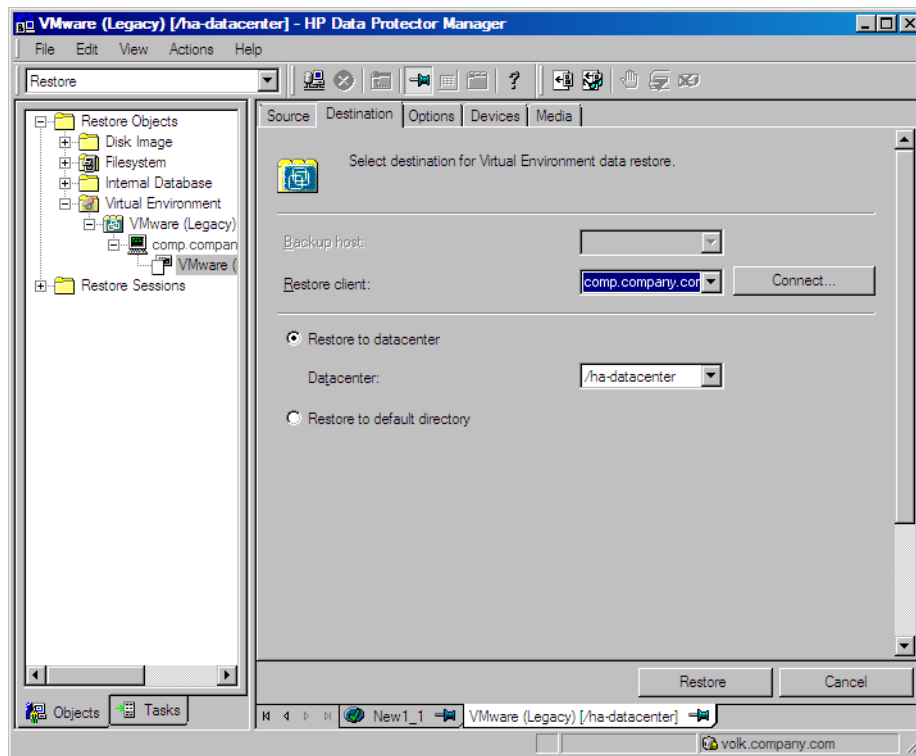


**Figure 66 Selecting VMware objects for restore (VCBfile)**



4. In the **Destination** page, specify the restore destination. For details, see [“VMware restore destination”](#) (page 127).

**Figure 67 Restore destination (Snapshot, Suspend, VCBimage)**



5. In the Options page, specify the VMware restore options. For details, see “VMware restore options” (page 127).

**Figure 68 Restore options (Snapshot, Suspend, VCBimage)**

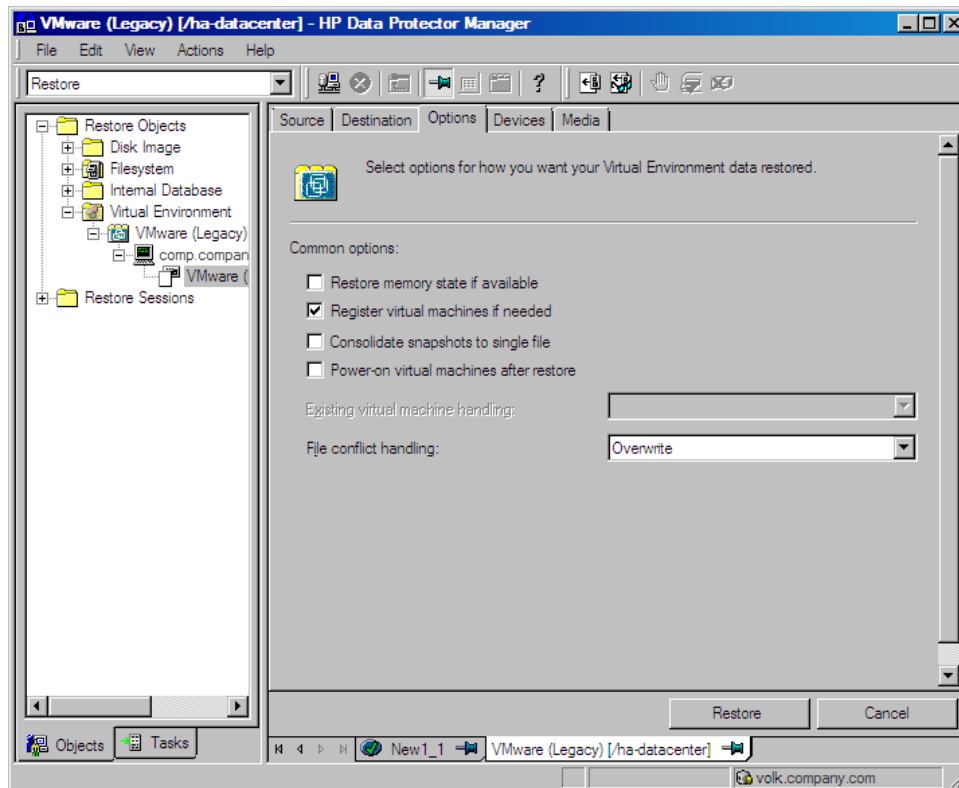
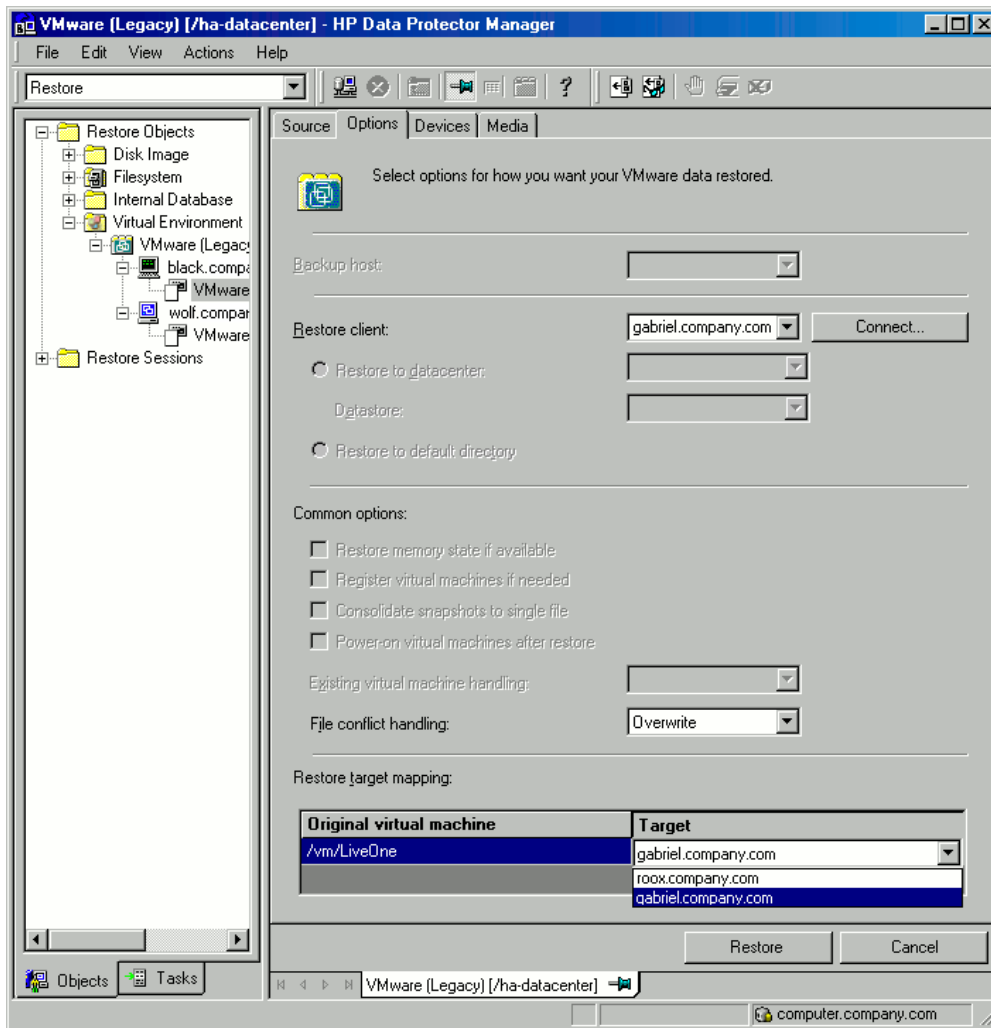


Figure 69 Restore options (VCBfile)



- ❗ **IMPORTANT:** If you restore from a **Snapshot** backup in which individual virtual machine disks were backed up, clear the option **Register virtual machines if needed**. Otherwise, the restore fails.

Whenever you restore with the option **Register virtual machines if needed** cleared, you need to manually recover the virtual machine after the session completes. For details, see [“Recovery after restoring to a datacenter”](#) (page 135).

6. In the **Devices** page, select the devices to be used for the restore.  
For more information on how to select devices for a restore, see the *HP Data Protector Help* index: “restore, selecting devices for”.
  7. Click **Restore**.
  8. In the **Start Restore Session** dialog box, click **Next**.
  9. Specify **Report level** and **Network load**.  
Click **Finish** to start the restore.
- The message `Session completed successfully` is displayed at the end of a successful session.

**Table 24 VMware restore destination**

GUI/CLI option	Description
<b>Backup host</b>	Not available.
<b>Restore client/</b> -destination	<p>Specifies the client that the restore session is started on. By default, the session is started on the same client on which the backup was started.</p> <p>Change the default only in the case of a disaster recovery when you restore from a <b>Suspend</b>, <b>Snapshot</b>, or <b>VCBimage</b> backup and the new VMware management client does not have the same name as the original one. Ensure that the new client has access to all the required datastores.</p> <p>You can also change the default if you want to restore the virtual machines outside a datacenter, in which case you must specify the client to restore to.</p> <p>If you select a client that has not been configured for use with the Data Protector VMware (Legacy) integration, a warning is displayed. Click <b>OK</b> and configure the client as described in <a href="#">“Configuring VMware management clients”</a> (page 109). However, to restore outside a datacenter, you do not need to configure the destination client, just click <b>Cancel</b>.</p> <p>To change the configuration of the selected client, click the <b>Connection</b> button.</p>
<b>Restore to datacenter/</b> -newinstance, -instance	<p>Specifies the datacenter that the virtual machines are restored to. By default, the virtual machines are restored to the original datacenter. Change the default only in the case of a disaster recovery, when the new datacenter does not have the same name as the original one.</p> <p>This option is not available if you restore from a <b>VCBfile</b> backup.</p>
<b>Restore to default directory</b>	<p>Select this option to restore virtual machines outside datacenter, to the client selected in <b>Restore client</b>, and to the directory specified in the OB2_VMWARE_PATH omnirc option on the selected client.</p> <p>This option is not available if you restore from a <b>VCBfile</b> backup.</p>

**Table 25 VMware restore options**

GUI/CLI option	Description
<b>Restore memory state if available/</b> -memory	<p>Select this option to restore the virtual machine memory file as well, if it was backed up.</p> <p>This option is not available if <b>Restore to default directory</b> is selected or if you restore from a <b>VCBfile</b> backup.</p> <p>Default: not selected.</p>
<b>Register virtual machines if needed/</b> -register	<p>Select this option to register restored virtual machines.</p> <p>Clear this option if you restore from a <b>Snapshot</b> backup in which individual virtual machine disks were backed up. Otherwise, the restore fails.</p> <p>If this option is not selected, you need to manually recover the restored virtual machines as described in <a href="#">“Recovering virtual machines manually”</a> (page 130).</p> <p>This option is not available if <b>Restore to default directory</b> is selected or if you restore from a <b>VCBfile</b> backup.</p> <p>Default: not selected.</p>
<b>Consolidate snapshots to single file/</b> -consolidate	<p>This option is applicable only if <b>Register virtual machines if needed</b> is selected. Select this option to commit all existing snapshots (including non-Data Protector ones) to the virtual machine base when the restore completes.</p> <p>This option is not available if <b>Restore to default directory</b> is selected or if you restore from a <b>VCBfile</b> backup.</p> <p>Default: not selected.</p>

**Table 25 VMware restore options** *(continued)*

GUI/CLI option	Description						
<b>Power-on virtual machines after restore/</b> <code>-poweron</code>	<p>This option is applicable only if <b>Register virtual machines if needed</b> is selected. Select this option to put the newly-restored virtual machines online when the session completes.</p> <p>This option is not available if <b>Restore to default directory</b> is selected or if you restore from a <b>VCBfile</b> backup.</p> <p>Default: not selected.</p>						
<b>File conflict handling/</b> <code>-overwrite [older]</code>	<p>Specifies Data Protector's behavior if virtual machines or files to be restored already exist in the destination. The file handling options have different meanings, depending on whether you restore from a <b>Snapshot</b>, <b>Suspend</b>, <b>VCBimage</b>, or <b>VCBfile</b> backup.</p> <p>If you restore from a <b>Snapshot</b>, <b>Suspend</b>, or <b>VCBimage</b> backup and you select <b>Restore to default directory</b>, the description for restore from a <b>VCBfile</b> backup applies.</p> <table> <tr> <td><b>Overwrite</b> (default)</td><td> <b>Snapshot, Suspend, VCBimage:</b> Data Protector unregisters such virtual machines, deletes their files, and then restores them from the backup.  <b>VCBfile:</b> Data Protector overwrites such files with those from the backup. </td></tr> <tr> <td><b>Keep latest</b></td><td> <b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup.  <b>VCBfile:</b> Data Protector leaves such files intact only if they are more recent than those from the backup. Otherwise, such files are overwritten with those from the backup. </td></tr> <tr> <td><b>Skip</b></td><td> <b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup.  <b>VCBfile:</b> Data Protector leaves such files intact. They are not restored from the backup. </td></tr> </table>	<b>Overwrite</b> (default)	<b>Snapshot, Suspend, VCBimage:</b> Data Protector unregisters such virtual machines, deletes their files, and then restores them from the backup. <b>VCBfile:</b> Data Protector overwrites such files with those from the backup.	<b>Keep latest</b>	<b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup. <b>VCBfile:</b> Data Protector leaves such files intact only if they are more recent than those from the backup. Otherwise, such files are overwritten with those from the backup.	<b>Skip</b>	<b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup. <b>VCBfile:</b> Data Protector leaves such files intact. They are not restored from the backup.
<b>Overwrite</b> (default)	<b>Snapshot, Suspend, VCBimage:</b> Data Protector unregisters such virtual machines, deletes their files, and then restores them from the backup. <b>VCBfile:</b> Data Protector overwrites such files with those from the backup.						
<b>Keep latest</b>	<b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup. <b>VCBfile:</b> Data Protector leaves such files intact only if they are more recent than those from the backup. Otherwise, such files are overwritten with those from the backup.						
<b>Skip</b>	<b>Snapshot, Suspend, VCBimage:</b> Data Protector leaves such virtual machines intact. They are not restored from the backup. <b>VCBfile:</b> Data Protector leaves such files intact. They are not restored from the backup.						
<b>Target hosts for restore/</b> <code>-target</code>	<p>This option is available only if you restore from a <b>VCBfile</b> backup.</p> <p>Specifies where to restore filesystems of virtual machines. The target system can be any Windows system (physical or virtual).</p>						

## Restoring using the Data Protector CLI

1. Log in to the VMware management client under a user account that is configured as described in [“Configuring VMware users”](#) (page 109).
2. Execute:

### **Restore of virtual machines:**

```

omnir -vmware
      -barhost OriginalVMwareManagementClient
      -instance OriginalDatacenter
      -method {snapshot | suspend | vcbimage}
      [-session BackupID | -fromsession BackupID -untilsession BackupID]
      VirtualMachines [-disk Disk [-disk Disk...]]
      [VirtualMachines [-disk Disk [-disk Disk...]]...]
      [-destination RestoreClient]
      [-newinstance NewDatacenter]
      [-consolidate]
      [-memory]
      [-register]
      [-poweron]
      [-overwrite [older]]

```



## Restore of filesystems:

```
omnir -vmware
-barhost OriginalVMwareManagementClient
-instance OriginalDatacenter
-method vcbfile
[-session BackupID | -fromsession BackupID -untilsession BackupID]
VirtualMachines -target TargetClient [-file File [-file File...]]
[VirtualMachines -target TargetClient [-file File [-file File...]]...]
[-destination RestoreClient]
[-overwrite [older]]
```

where *VirtualMachines* must be specified in one of the following ways:

```
{
-all [-exclude VMFolder [-exclude VMFolder...]] |
-vmfolder VMFolder [-exclude VMFolder [-exclude VMFolder...]] |
-vm VM
}
```

For a description of all the options, see the *omnir* man page or the *HP Data Protector Command Line Interface Reference*.

---

**NOTE:** *BackupID* is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the object copy session.

The *omnir* syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

---

### Example (Restoring virtual machines to a datacenter)

Suppose you want to restore the virtual machine `/vm/machineA` and the individual disks (`scsi0:0` and `scsi0:1`) of the virtual machine `/vm/machineB`. At the time of backup, the virtual machines were running on the ESX Server systems that belonged to the datacenter `/MyDatacenter` managed by the VirtualCenter Server system `Virtualcenter.company.com`. The virtual machines were backed up with the *Suspend* backup method.

You want to restore them to the original location, using the backup session `2010/07/14-1`. If, in this session, virtual machine memory files were also backed up, you want to restore them as well. You also want to ensure that the newly restored virtual machines are put online when the session completes.

Execute:

```
omnir -vmware -barhost Virtualcenter.company.com -instance /MyDatacenter
-method suspend -session 2010/7/14-1 -vm /vm/machineA -vm /vm/machineB
-disk scsi0:0 -disk scsi0:1 -memory -poweron
```

### Example (Restoring virtual machines outside a datacenter)

Suppose the virtual machines `/MyVirtualMachines/machineA` and `/MyVirtualMachines/machineB` were backed up in the session `2010/8/14-5` from the datacenter `/MyDatacenter` that is managed by the VirtualCenter Server system `Virtualcenter.company.com`, using the *VCBimage* backup method. To restore the virtual machines outside the datacenter, to the directory `C:\tmp` on the VCB proxy system `proxy.company.com`:

1. Log in to the VCB proxy system.
2. Set the omnirc option OB2\_VMWARE\_PATH to C:\tmp.
3. Execute:

```
omnir -vmware -barhost VirtualCenter.company.com -instance
/MyDatacenter -method vcbimage -session 2010/8/14-5 -vm
/MyVirtualMachines/machineA -vm /MyVirtualMachines/machineB
-destination proxy.company.com -newinstance "None"
```

### Example (Restoring filesystems of virtual machines)

Suppose you want to restore all filesystems of all the virtual machines contained in the VMware vSphere or Virtual Infrastructure inventory folder /MyVirtualMachines, except the filesystems of the virtual machine /MyVirtualMachines/MachineA. The restore destination is the Windows client computer1.company.com.

In addition, you want to restore the C:\Documents and Settings folder and the file C:\Test\Schedule.txt of the virtual machine /MyVirtualMachines2/MachineB back to the same virtual machine MachineB.company.com. The virtual machines were backed up from the datacenter /MyDatacenter that was managed by the VirtualCenter Server system VirtualCenter.company.com. You want to restore from the last backup session.

Execute:

```
omnir -vmware -barhost VirtualCenter.company.com -instance /MyDatacenter
-method vcbfile -vmfolder /MyVirtualMachines -exclude
/MyVirtualMachines/MachineA -target computer1.company.com -vm
/MyVirtualMachines2/MachineB -target MachineB.company.com -file
"C/Documents and Settings" -file C/Test/Schedule.txt
```

## Recovering virtual machines manually

There are two different scenarios when you need to manually recover virtual machines after they have been restored with Data Protector:

- If you have restored the virtual machine files to a directory on a Data Protector client due to one of the following reasons:
  - You selected **Restore to default directory** restore option.
  - You selected **Restore to datacenter** restore option, but the connection with the target ESX(i) Server system could not be established.

See ["Recovery after restoring outside a datacenter"](#) (page 130).

- If you have restored the virtual machine files to the original ESX(i) Server systems by selecting the **Restore to datacenter** restore option, but without selecting the restore option **Register virtual machines if needed** (see ["Recovery after restoring to a datacenter"](#) (page 135)).

### Recovery after restoring outside a datacenter

Suppose the virtual machine FreeBSD 8.0 beta2 i386 was backed up in the **VCBimage** backup session 2010/07/07-2 and you have restored the virtual machine to the directory C:\tmp on the VCB proxy system droid.company.com.

To manually move the restored virtual machine image to the ESXi Server system srsen.company.com managed by the VirtualCenter Server system volk.company.com, using the VMware Converter:

1. Display the content of the restored directory C:\tmp\FreeBSD 8.0 beta2 i386:

```
catalog_2010_07_17_0002
FreeBSD 8.0 beta2 i386.vmdk_2010_07_17_0002
FreeBSD 8.0 beta2 i386.vmx_2010_07_17_0002
scsi0-0-0-FreeBSD 8.0 beta2 i386-flat.vmdk_2010_07_17_0002
```

2. Delete the backup session IDs appended at the end of the filenames:

```
catalog
FreeBSD 8.0 beta2 i386.vmdk
FreeBSD 8.0 beta2 i386.vmx
scsi0-0-0-FreeBSD 8.0 beta2 i386-flat.vmdk
```

3. Open the catalog file with a text editor:

```
catalog
version= esx-3.0
state= poweredOn
display_name= "FreeBSD 8.0 beta2 i386"
uuid= "503d98d9-cad4-5c0e-9132-b35de35b900d"
disk.scsi0:0.filename= "scsi0-0-0-FreeBSD 8.0 beta2 i386.vmdk"
disk.scsi0:0.diskname= "[Storage_VmWare_all] FreeBSD 8.0 beta2
i386/FreeBSD 8.0 beta2 i386.vmdk"
config.vmx= "[Storage_VmWare_all] FreeBSD 8.0 beta2 i386/FreeBSD
8.0 beta2 i386.vmx"
host= srsen.company.com
timestamp= "Fri Jul 17 05:52:17 2010"
config.suspenddir= "[Storage_VmWare_all] FreeBSD 8.0 beta2 i386/"
config.snapshotdir= "[Storage_VmWare_all] FreeBSD 8.0 beta2 i386/"
config.file0= "FreeBSD 8.0 beta2 i386.vmsd"
config.file1= "FreeBSD 8.0 beta2 i386.vmx"
config.file2= "FreeBSD 8.0 beta2 i386.nvram"
config.logdir= "[Storage_VmWare_all] FreeBSD 8.0 beta2 i386/"
config.log0= "vmware.log"
folderpath= "/Datacenters/europa/vm"
resourcepool= "/Datacenters/europa/host/srsen.company.com/Resources"
```

The name specified in the disk.scsi0:0.filename catalog variable and the name of the base metadata file residing in the folder must match. Therefore, rename the base metadata file FreeBSD 8.0 beta2 i386.vmdk to scsi0-0-0-FreeBSD 8.0 beta2 i386.vmdk:

```
catalog
scsi0-0-0-FreeBSD 8.0 beta2 i386.vmdk
FreeBSD 8.0 beta2 i386.vmx
scsi0-0-0-FreeBSD 8.0 beta2 i386-flat.vmdk
```

4. Share the folder C:\tmp\FreeBSD 8.0 beta2 i386 so that it is accessible from the VirtualCenter Server system.
5. Log in to the VirtualCenter Server system and open the VMware Converter user interface.
6. Click the **Convert Machine** button to open the **Conversion** wizard.
7. In the Specify Source page, select **Backup image or third-party virtual machine** for the source type and \\droid\tmp\FreeBSD 8.0 beta2 i386\FreeBSD 8.0 beta2 i386.vmx for the virtual machine.

**Figure 70 Conversion (Specify Source)**

**Conversion**

1 Specify Source 2 Specify Destination 3 View/Edit Options 4 Ready to Complete

**Machine Type**

Source: none Destination: none

Select source type: Backup image or third-party virtual machine

Convert a backup image or a third-party virtual machine.

Browse for source virtual machine or image

Virtual machine file: \\droid\ntp\FreeBSD 8.0 beta2 i386\FreeBSD 8.0 beta2 i386.vmx Browse...

[View source details...](#)

Click **Next**.

8. In the Specify Destination (Destination Type) page, provide the login information for the VirtualCenter Server system.

**Figure 71 Conversion (Specify Destination 1)**

**Conversion**

1 Specify Source 2 Specify Destination 3 View/Edit Options 4 Ready to Complete

**Destination Type** Host/Resource

Source: \\droid\ntp\FreeBSD...0 beta2 i386\FreeBSD 8.0 beta2 i386.vmx Destination: none

Select destination type: VMware Infrastructure virtual machine

Creates a new virtual machine for use on a VMware Infrastructure product.

VMware Infrastructure server details

Server: wolf


User name: Administrator

Password: .....

Click **Next**.

9. In the Specify Destination (Host/Resource) page, specify the destination ESXi Server system `srsen.company.com`.

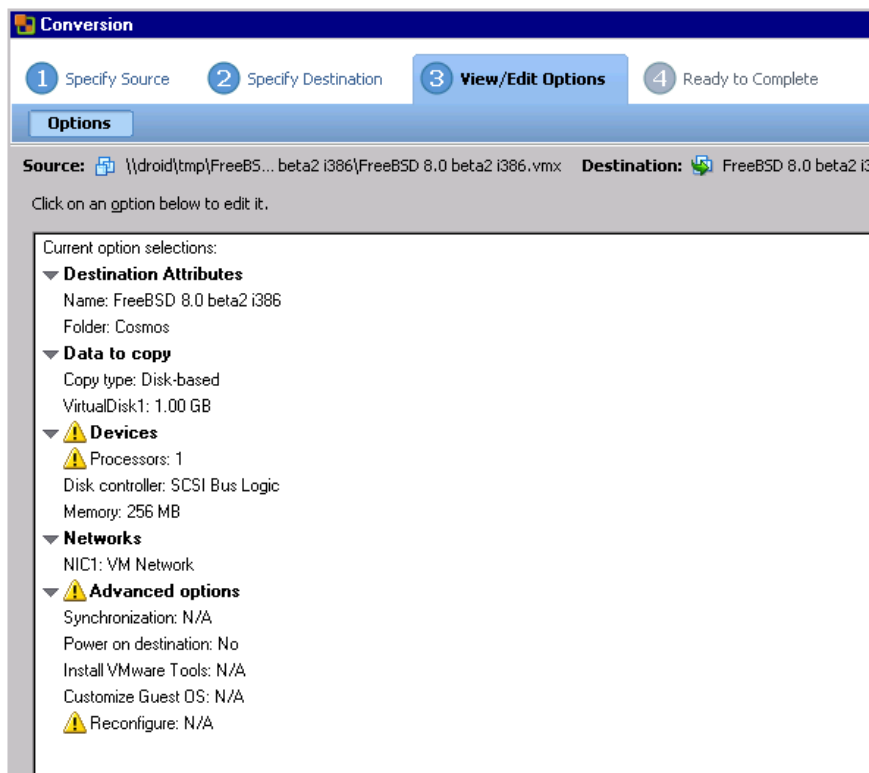
**Figure 72 Conversion (Specify Destination 2)**



Click **Next**.

10. In the View/Edit Options page, modify the options (optional).

**Figure 73 Conversion (View/Edit Options)**



Click **Next**.

11. In the Ready to Complete page, review your selection and click **Finish**.

**Figure 74 Conversion (Ready to Complete)**

**Conversion**

1 Specify Source 2 Specify Destination 3 View/Edit Options 4 **Ready to Complete**

**Summary**

Source: \\droid\{tmp\FreeBS... beta2 i386\FreeBSD 8.0 beta2 i386.vmx Destination: FreeBSD 8.0 beta2 i386

**Source System Information**

Source type: Backup image or third-party virtual machine  
Path: \\droid\{tmp\FreeBSD 8.0 beta2 i386\FreeBSD 8.0 beta2 i386.vmx

**Destination System Information**

Virtual machine name: FreeBSD 8.0 beta2 i386  
Host/Server: volk  
Connected as: Administrator  
VM folder: None  
Host system: seren.company.com  
Resource pool: Default  
Power on after conversion: No  
Number of CPUs: 1  
Physical memory: 256 MB  
Network: Preserve NIC count  
NIC1: Connected  
VM Network  
Storage: Disk-based Cloning  
Number of disks: 1  
Create disk 0 as: Monolithic Flat Disk  
Configuration files Datastore: Storage2\_Seren

**Destination Customization**

Not customized

12. After a while, the VMware Converter informs you about the successful recovery.

**Figure 75 Recovery completed**

**VMware vCenter Converter Standalone**

File View Task Administration Help

Convert Machine Configure Machine

Show: All tasks Recent tasks

Task ID	Source	Destination	Status	Start time	End time
5	...FreeBSD 8.0 be...	seren.company.com	Completed	31/07/2009 16:5...	31/07/2009 17:11:41

Task ID 5: Completed Source: \\droid...FreeBSD 8.0 beta2 i386.vmx Destination: seren.compa...

**Summary** Task Progress

**Conversion Status**

Type: Convert virtual machine to virtual machine  
Created: 31/07/2009 16:52:41 by Administrator  
Status: Completed  
Started: 31/07/2009 16:52:42  
Completed: 31/07/2009 17:11:41  
Running time: 18 minute(s)

**Log Highlights** Export Logs...

31/07/2009 17:11:40 Updating the boot sectors on the virtual disk  
31/07/2009 17:11:39 Completed cloning disk 0 on the virtual disk  
31/07/2009 16:52:52 Starting to clone disk 0 on the virtual disk  
31/07/2009 16:52:51 Created the virtual machine 'FreeBSD 8.0 beta2 i386'  
31/07/2009 16:52:41 Task: Convert virtual machine.

## Recovery after restoring to a datacenter

The following procedure enables you to import restored virtual machine disks into a new or an existing virtual machine. Use this procedure if you have restored individual virtual machine disks or virtual machines to the original ESX(i) Server system (**Restore to datacenter**), but without selecting the option **Register virtual machines if needed**. Note that after such a restore, the original virtual machines are no longer registered. So if you decide to import the disks into new virtual machines, the new virtual machines may have the same names as the original ones.

However, before you import the disks, you have to prepare the disks first, depending on the backup you restored from:

- **Snapshot:** Link and consolidate virtual machine disk files.
- **Suspend:** Consolidate virtual machine disk files.
- **VCBimage:** No preparatory steps are required.

### Linking virtual machine disk files

The procedure for linking virtual machine disk files is described using the following example:

**Table 26 Virtual machine information**

Virtual machine	helios
Virtual machine disks	<ul style="list-style-type: none"><li>• <code>scsi0:0</code> (1.2 GB) base metadata file: <code>helios.vmdk</code> base extent file: <code>helios-flat.vmdk</code></li><li>• <code>scsi1:2</code> (1.0 GB) base metadata file: <code>helios_1.vmdk</code> base extent file: <code>helios_1-flat.vmdk</code></li></ul>
Datastore	storage2

Suppose you backed up the virtual machine disk `scsi1:2` in the following **Snapshot** backup sessions:

1. **Full** (2010/09/10-1)
2. **Incr** (2010/09/10-2)
3. **Differential** (2010/09/10-3)
4. **Incr** (2010/09/10-4)

When you restore the virtual machine disk `scsi1:2` from the session 2010/09/10-4, Data Protector automatically restores the complete restore chain:

1. 2010/09/10-1 (full backup)
2. 2010/09/10-3 (differential backup)
3. 2010/09/10-4 (incremental backup)

After the restore, you need to manually update the content of `scsi1:2` metadata files by correcting the extent file names and metadata parent file names:

1. Display the content of the restored directory /vmfs/volumes/storage2/helios/:

```
.  
..  
helios_1-000001-delta.vmdk.2010_09_10_0003  
helios_1-000001.vmdk.2010_09_10_0003  
helios_1-000002-delta.vmdk.2010_09_10_0004  
helios_1-000002.vmdk.2010_09_10_0004  
helios_1-flat.vmdk.2010_09_10_0001  
helios_1.vmdk.2010_09_10_0001  
helios.nvram.2010_09_10_0001  
helios.nvram.2010_09_10_0003  
helios.nvram.2010_09_10_0004  
helios-Snapshot1.vmsn.2010_09_10_0001  
helios-Snapshot4.vmsn.2010_09_10_0003  
helios-Snapshot5.vmsn.2010_09_10_0004  
helios.vmsd.2010_09_10_0001  
helios.vmsd.2010_09_10_0003  
helios.vmsd.2010_09_10_0004  
helios.vmx.2010_09_10_0001  
helios.vmx.2010_09_10_0003  
helios.vmx.2010_09_10_0004  
helios.vmx.2010_09_10_0001  
helios.vmx.2010_09_10_0003  
helios.vmx.2010_09_10_0004  
helios.vmx.2010_09_10_0001  
helios.vmx.2010_09_10_0003  
helios.vmx.2010_09_10_0004
```

2. Identify the scsi1:2 metadata files:

```
helios_1.vmdk.2010_09_10_0001  
helios_1-000001.vmdk.2010_09_10_0003  
helios_1-000002.vmdk.2010_09_10_0004
```

Identify the corresponding extent files:

```
helios_1-flat.vmdk.2010_09_10_0001  
helios_1-000001-delta.vmdk.2010_09_10_0003  
helios_1-000002-delta.vmdk.2010_09_10_0004
```

The correct sequence in which the metadata files should be linked is determined from timestamps appended at the end of filenames. The metadata file `helios_1.vmdk.2010_09_10_0001` is the parent of `helios_1-000001.vmdk.2010_09_10_0003`, which is the parent of `helios_1-000002.vmdk.2010_09_10_0004`.

3. Open the base metadata file `helios_1.vmdk.2010_09_10_0001`:

```
# Disk DescriptorFile  
version=1  
CID=d0fb6c81  
parentCID=ffffffff  
createType="vmfs"  
  
# Extent description  
RW 2097152 VMFS "helios_1-flat.vmdk"  
# The Disk Data Base  
#DDB  
ddb.virtualHWVersion = "4"  
ddb.geometry.cylinders = "512"  
ddb.geometry.heads = "128"  
ddb.geometry.sectors = "32"  
ddb.adapterType = "lsilogic"
```

Change the extent file name `helios_1-flat.vmdk` to `helios_1-flat.vmdk.2010_09_10_0001`.

4. Open the snapshot metadata file `helios_1-000001.vmdk.2010_09_10_0003`:

```
# Disk DescriptorFile  
version=1  
CID=d0fb6c81
```



```
parentCID=d0fb6c81
createType="vmfsSparse"
parentFileNameHint="helios_1.vmdk"
# Extent description
RW 2097152 VMFSSPARSE "helios_1-000001-delta.vmdk"
# The Disk Data Base
#DDB
```

Change the parent metadata file name `helios_1.vmdk` to `helios_1.vmdk.2010_09_10_0001`. Change the extent file name `helios_1-000001-delta.vmdk` to `helios_1-000001-delta.vmdk.2010_09_10_0003`.

5. Open the snapshot metadata file `helios_1-000002.vmdk.2010_09_10_0004`:

```
# Disk DescriptorFile
version=1
CID=d0fb6c81
parentCID=d0fb6c81
createType="vmfsSparse"
parentFileNameHint="helios_1-000001.vmdk"
# Extent description
RW 2097152 VMFSSPARSE "helios_1-000002-delta.vmdk"
# The Disk Data Base
#DDB
```

Change the parent metadata file name `helios_1-000001.vmdk` to `helios_1-000001.vmdk.2010_09_10_0003`. Change the extent file name `helios_1-000002-delta.vmdk` to `helios_1-000002-delta.vmdk.2010_09_10_0004`.

### Consolidating virtual machine disks

1. Connect to the ESX Server system to which the virtual machine was restored.
2. Execute:

```
vmkfstools --clonevirtualdisk LastMetadataFile RecoveredDisk
```

For example:

```
vmkfstools --clonevirtualdisk helios_1-000002.vmdk.2010_09_10_0004
helios_1_recovered.vmdk
```

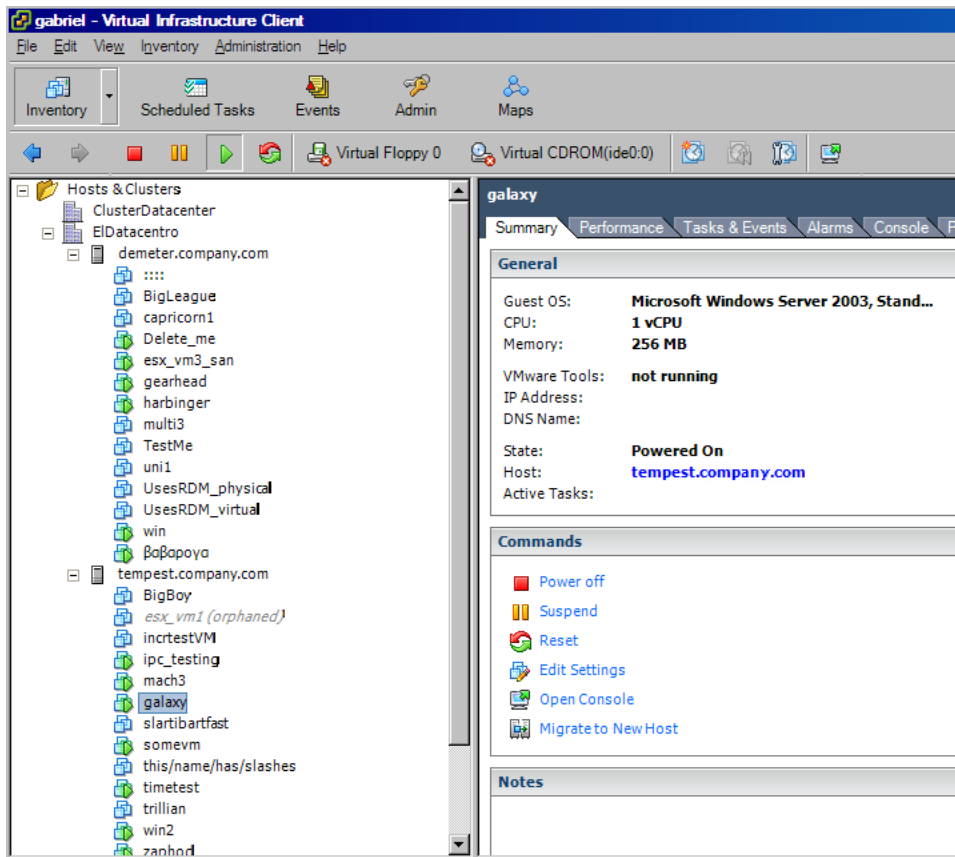
As a result, the file `helios_1_recovered.vmdk` is created in the directory `/vmfs/volumes/storage2/helios/`.

## Importing virtual machine disks

To import the virtual machine disk `helios_1_recovered.vmdk` into the virtual machine `galaxy`:

1. Open a VMware vSphere or Virtual Infrastructure Client.


**Figure 76 Virtual Infrastructure Client**



Right-click the virtual machine `galaxy` and click **Edit Settings**.


2. In the Virtual Machine Properties dialog box, select **Hard Disk 1** and click **Add**.

**Figure 77 Virtual machine properties**



3. In the Select Device Type page of the Add Hardware wizard, select **Hard disk** and click **Next**.
4. In the Select a Disk page of the Add Hardware wizard, select **Use an existing virtual disk** and click **Next**.
5. In the Select Existing Disk page of the Add Hardware wizard, click **Browse**.
6. In the Browse Datastores dialog box, browse to /vmfs/volumes/storage2/helios, select helios\_1\_recovered.vmdk and click **Open**.

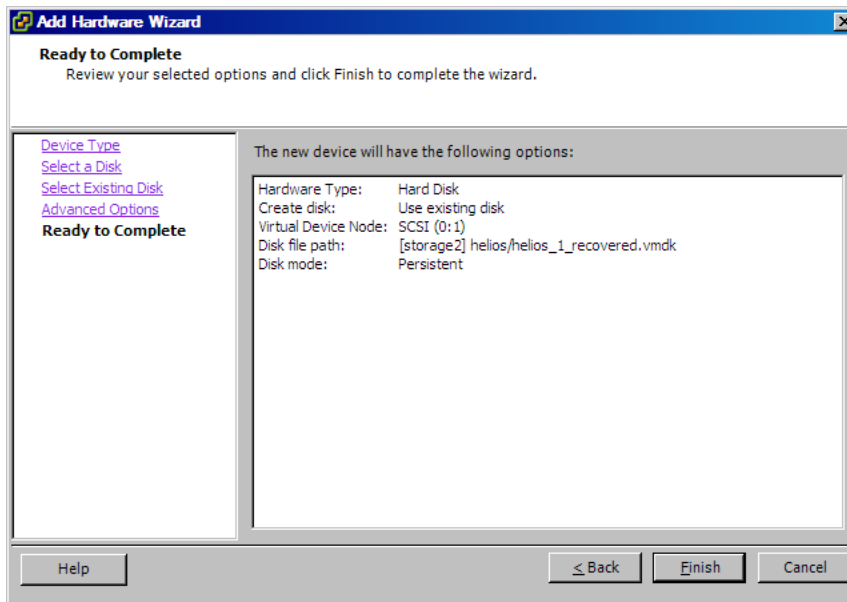
**Figure 78 Browsing datastores**



7. In the Advanced Options page of the Add Hardware wizard, select a Virtual Device Node for the new disk and click **Next**.

8. In the Ready to Complete page of the Add Hardware wizard, review your selection and click **Finish**.

**Figure 79 Add hardware summary**



## Restoring using another device

You can perform a restore using a device other than that used for the backup. For details, see the *HP Data Protector Help* index: “restore, selecting devices for”.

## Disaster recovery

Disaster recovery is very complex, involving different products from different vendors. Check the instructions from the guest operating systems and VMware on how to prepare for it.

The following are the main steps needed to recover a virtual machine after a disaster:

1. Reinstall the VMware environment. The configuration should be the same as during the backup. Although the new VMware management client and datacenters can have different names from before, ensure that the datastores have the same names as the original ones. Otherwise, the restore fails.
2. Install Data Protector in the newly configured environment.
3. Restore the service console of the ESX Server system on which the virtual machine was running to the newly configured ESX Server system from a Data Protector filesystem backup.  
For details on what to restore, see the topic “ESX Server Configuration Backup and Restore procedure” at <http://kb.vmware.com/selfservice/microsites/microsite.do>.  
For details of how to restore from a filesystem backup, see the *HP Data Protector Help*.
4. Restore the original VirtualCenter database (if needed). For details, see the Data Protector integration that was used to back up the database.
5. Restore the virtual machine from a Data Protector VMware backup as described in “Restore” (page 122).

## Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run a backup or restore session, a monitor window shows the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the `User Interface` component installed, using the `Monitor` context.

To monitor a session, see the *HP Data Protector Help* index: “viewing currently running sessions”.

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the VMware (Legacy) integration.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

### Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HP Data Protector Help* index: “patches” for information of how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, and other information.

### Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors reported in the `debug.log`.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HP Data Protector Help*.

Additionally, if your backup failed:

- Check the configuration of the VirtualCenter Server or standalone ESX Server system as described in “[Configuring VMware management clients](#)” (page 109).

## Problems

### Problem

#### **Suspend backup fails with the error “IPC Cannot Connect”**

A **Suspend** backup of virtual machines running on ESX Server systems that are configured in a VirtualCenter environment fails with the `IPC Cannot Connect` error.

### Action

Ensure that the Data Protector VMware (Legacy) Integration agent is not blocked by the ESX Server firewall. For information of how to limit the range of port numbers from which dynamically allocated ports are selected, see the *HP Data Protector Help* index: “firewall support”.

### Problem

#### **Creation of a VCBfile backup specification fails with the error “IPC Cannot Connect”**

When you create a backup specification for a VCBfile backup and you try to mount filesystems of virtual machines in the Source page, Data Protector displays the `IPC Cannot Connect` error.

### Action

Ensure that the Data Protector VMware (Legacy) Integration agent is not blocked by the ESX Server firewall. For information of how to limit the range of port numbers from which dynamically allocated ports are selected, see the *HP Data Protector Help* index: “firewall support”.

## Problem

### Suspend backup fails with the error “insufficient resources to satisfy failover level”

This error may occur if your ESX Server systems are configured in a high availability cluster and one of the ESX Server systems fails. Consequently, the virtual machines from the failed system migrate elsewhere. If you start a **Suspend** backup of such virtual machines, the power-on operation at the end of the backup fails with an error similar to the following:

```
[Critical] From: OB2BAR_VMWARE_BAR@gabriel.company.com
"/ClusterDatacenter" Time: 7.4.2010 16:13:50 Virtual machine
'/vm/vmsan1': operation failed: Error: {
localizedMessage='Insufficient resources to satisfy configured
failover level for HA.';
```

## Action

1. Open the cluster settings dialog box in the VMware vSphere or Virtual Infrastructure client and select the option **Allow virtual machines to be powered on even if they violate availability constraints**.
2. Power the virtual machines on and restart the backup (if needed).

## Problem

### VCB backup fails

When you start a **VCBimage** or **VCBfile** backup session, the VCB mount cannot be performed and the session fails with the following error:

Cannot query guest OS information. The mount directory path is invalid.

In addition, if you restart the session, the Data Protector GUI stops responding.

## Action

VCB mount fails if the guest operating system is not installed on the first virtual hard disk. For details, see known issues in the VMware vSphere Consolidated Backup Release Notes:

[http://www.vmware.com/support/vsphere4/doc/vsp\\_vcb\\_15\\_u2\\_rel\\_notes.html](http://www.vmware.com/support/vsphere4/doc/vsp_vcb_15_u2_rel_notes.html).

## Problem

### Orphaned virtual machines

When you open the VMware vSphere or Virtual Infrastructure client while the restore session is in progress, the virtual machines that are being restored are shaded and the note (orphaned) is added next to them.

This happens if you back up virtual machines from a datacenter that is managed by a VirtualCenter Server system, and then restore the virtual machines, specifying a different client on which to start the restore. Specifically, you change the restore option **Restore client** from the VirtualCenter Server system to the ESX Server system on which the virtual machines were running. Consequently, the restore option **Restore to datacenter** automatically changes to the datacenter `/ha-datacenter`. As a result, during restore, the virtual machines are unregistered inside the datacenter `/ha-datacenter`, but not inside the VirtualCenter Server datacenter from which they were backed up.

## Action

Do not change the restore destination. Changes are allowed only in case of a disaster recovery when the newly configured VirtualCenter systems, ESX Server systems, or datacenters have different names than the original ones.

## Problem

### Backup fails if a VM filesystem is still mounted

If not enough time passes between dismounting and mounting of a VM (mount point), the vcbMounter will report that it cannot create a directory, although the directory has been removed:

```
[2010-08-04 12:13:05.218 'vcbMounter' 712 error] Error: Cannot create
directory c:\mount-p\26752.3[2010-08-04 12:13:05.218 'vcbMounter' 712
error] An error occurred, cleaning up...
```

The issue appears only with interactive backups.

## Part II Microsoft Hyper-V

Data Protector offers different ways to back up Microsoft Hyper-V data online. Choose the appropriate backup and restore solution depending on desired functionality.

### Microsoft Hyper-V

You can select from among the following solutions:

- **Data Protector Virtual Environment integration**

This integration operates on the level of Microsoft Hyper-V virtual machines. The smallest object that you can back up and restore is a Microsoft Hyper-V virtual machine. In one session you can back up virtual machines from multiple Microsoft Hyper-V systems if they are configured in a cluster.

The main advantage over the Data Protector Microsoft Volume Shadow Copy Service integration is that this integration supports virtual machine migration in a Microsoft Hyper-V cluster, and that it supports the incremental backup type in applicable Microsoft Hyper-V environments.

For details, see [“Data Protector Virtual Environment integration” \(page 146\)](#).

- **Data Protector Microsoft Volume Shadow Copy Service integration**

This integration operates on the level of Microsoft Volume Shadow Copy Service writers. Since the Microsoft Hyper-V writer is only one of them, Microsoft Hyper-V data does not have the highest visibility. The smallest object that you can back up or restore is a Microsoft Hyper-V virtual machine. In one session, you can back up data from only one Microsoft Hyper-V system. The only supported backup type is full backup.

For details, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

**NOTE:** You can also back up Microsoft Hyper-V virtual machines using common Data Protector filesystem backup functionality, which operates on the file level. The smallest object that you can back up or restore this way is a file. To ensure data consistency, you must shut the virtual machines down before starting a backup session.

**Table 27 Data Protector backup solutions for Microsoft Hyper-V**

Feature		Disk Agent (filesystem backup)	Volume Shadow Copy Service integration	Virtual Environment integration
Online backup			✓	✓
Crash-consistent backup		✓	✓	✓
Application-consistent backup			✓	✓
Granularity		File	Virtual machine	Virtual machine
Backup types	Full	✓	✓	✓
	Incremental	✓		✓
	Differential	✓		
Is zero downtime backup (ZDB) supported?		✓	✓	✓
Is instant recovery supported?		✓	✓	



**Table 27 Data Protector backup solutions for Microsoft Hyper-V (continued)**

Feature	Disk Agent (filesystem backup)	Volume Shadow Copy Service integration	Virtual Environment integration
Is virtual machine migration within a Microsoft Hyper-V cluster supported?			✓
Extra licenses needed	<ul style="list-style-type: none"> <li>• Zero Downtime Backup (optional)</li> <li>• Instant Recovery (optional)</li> </ul>	<ul style="list-style-type: none"> <li>• On-line Extension</li> <li>• Zero Downtime Backup (optional)</li> <li>• Instant Recovery (optional)</li> </ul>	<ul style="list-style-type: none"> <li>• On-line Extension</li> <li>• Zero Downtime Backup (optional)</li> </ul>

---

## 3 Data Protector Virtual Environment integration

### Introduction

This chapter explains how to configure and use the Data Protector Virtual Environment integration for Microsoft Hyper-V. It describes concepts and methods you need to understand to back up and restore Microsoft Hyper-V virtual machines.

The integration supports both standalone Microsoft Hyper-V system environments (**standalone environments**) as well as environments in which Microsoft Hyper-V systems are configured in a cluster (**clustered environments**).

The Data Protector Virtual Environment integration for Microsoft Hyper-V is based on the Data Protector Volume Shadow Copy Service integration (**VSS integration**). For details on the VSS concepts, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

### Backup

During backup, virtual machines can be either powered off (**offline backup**) or actively used (**online backup**).

The following backup method is available:

- Hyper-V Image

For details, see [“Hyper-V Image backup method” \(page 152\)](#).

As Microsoft Hyper-V and VSS are involved, you should specify the following backup types for each backup session:

- **Microsoft Hyper-V backup type**

Data Protector offers interactive and scheduled backups of the following types:

- Full
- Incremental

- **VSS backup type**

The supported VSS backup types are:

- Local backup
- Transportable backup

For details on the backup types, see [“Backup types” \(page 153\)](#).

### Restore

You can restore Microsoft Hyper-V virtual machines:

- **To the default (original) location**

Using this option, you can restore virtual machines to the original location on the original or a different Microsoft Hyper-V system.

- **To a different location**

Using this option, you can restore virtual machines to a different location on the original or a different Microsoft Hyper-V system.

- **To a directory**

Using this option, you can restore virtual machine files to a directory on any client with the Data Protector Virtual Environment Integration and MS Volume Shadow Copy

Integration components installed. After such a restore you need to import virtual machines to a Microsoft Hyper-V system for virtual machines to become functional again.

The Data Protector Virtual Environment integration does not support instant recovery.

This chapter provides information specific to the Data Protector Virtual Environment integration for Microsoft Hyper-V. For limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HP Data Protector Help*.

## Integration concepts

### Supported environments

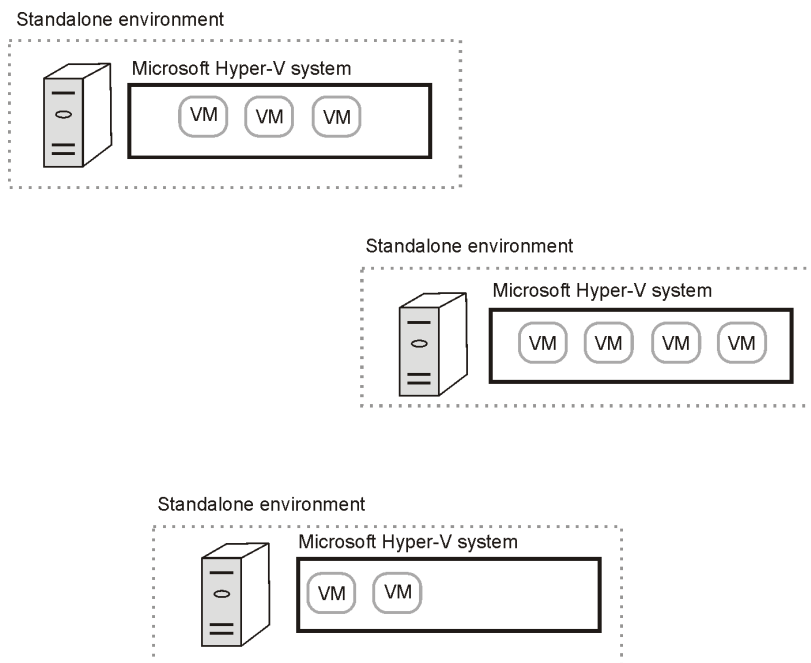
Data Protector supports both standalone and clustered Microsoft Hyper-V environments.

#### Standalone environments

In a standalone Microsoft Hyper-V environment, the Cell Manager communicates a backup request to the backup host, which then sends the request to the corresponding standalone Microsoft Hyper-V system. A **backup host** is a system in the Data Protector cell that has the Data Protector Virtual Environment Integration and MS Volume Shadow Copy Integration components installed and controls the process of backing up Microsoft Hyper-V virtual machines.

In a single session, virtual machines that reside on the same Microsoft Hyper-V system can be selected for backup. Where the data is actually backed up from depends on configuration of the Microsoft Hyper-V virtual machine replication and options selected in your Data Protector backup specification.

**Figure 80 Hyper-V standalone environments**



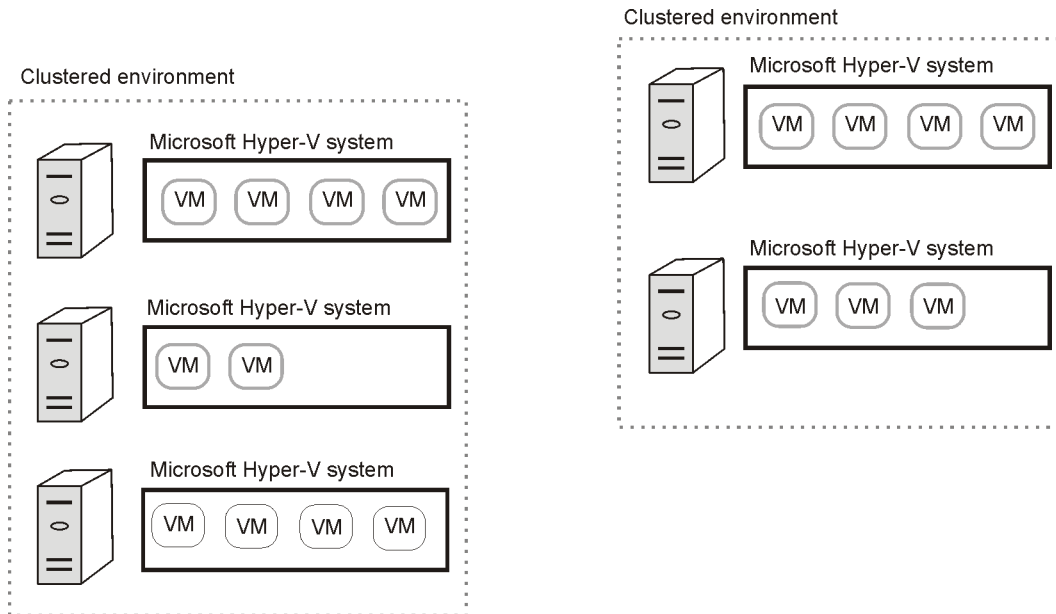
For standalone Microsoft Hyper-V systems, it is only possible to migrate a virtual machine by manually exporting it from one Microsoft Hyper-V system and then importing it into another.

#### Clustered environments

In a clustered Microsoft Hyper-V environment, the Cell Manager communicates a backup request to the backup host, which then sends the request to the Microsoft Hyper-V system on which the virtual machine to be backed up resides.

In a single session, virtual machines that reside in the same Microsoft Hyper-V cluster can be selected for backup. Where the data is actually backed up from depends on configuration of the Microsoft Hyper-V virtual machine replication and options selected in your Data Protector backup specification.

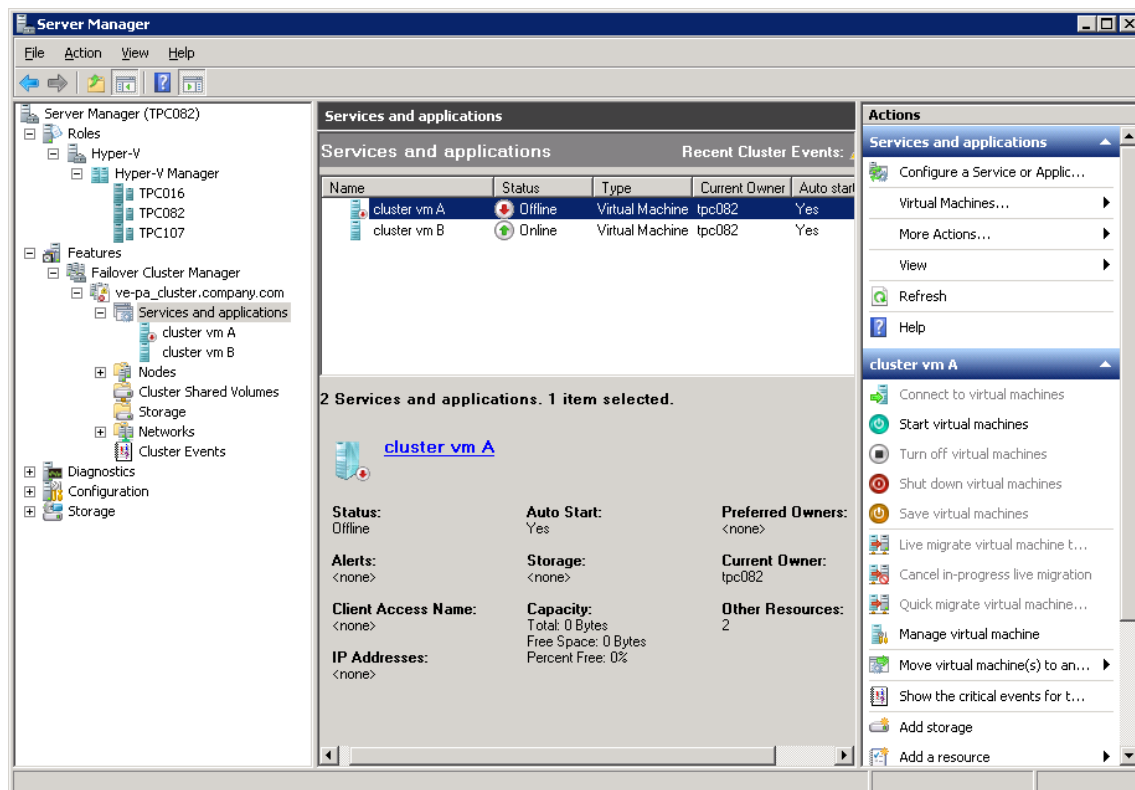
**Figure 81 Hyper-V clustered environments**



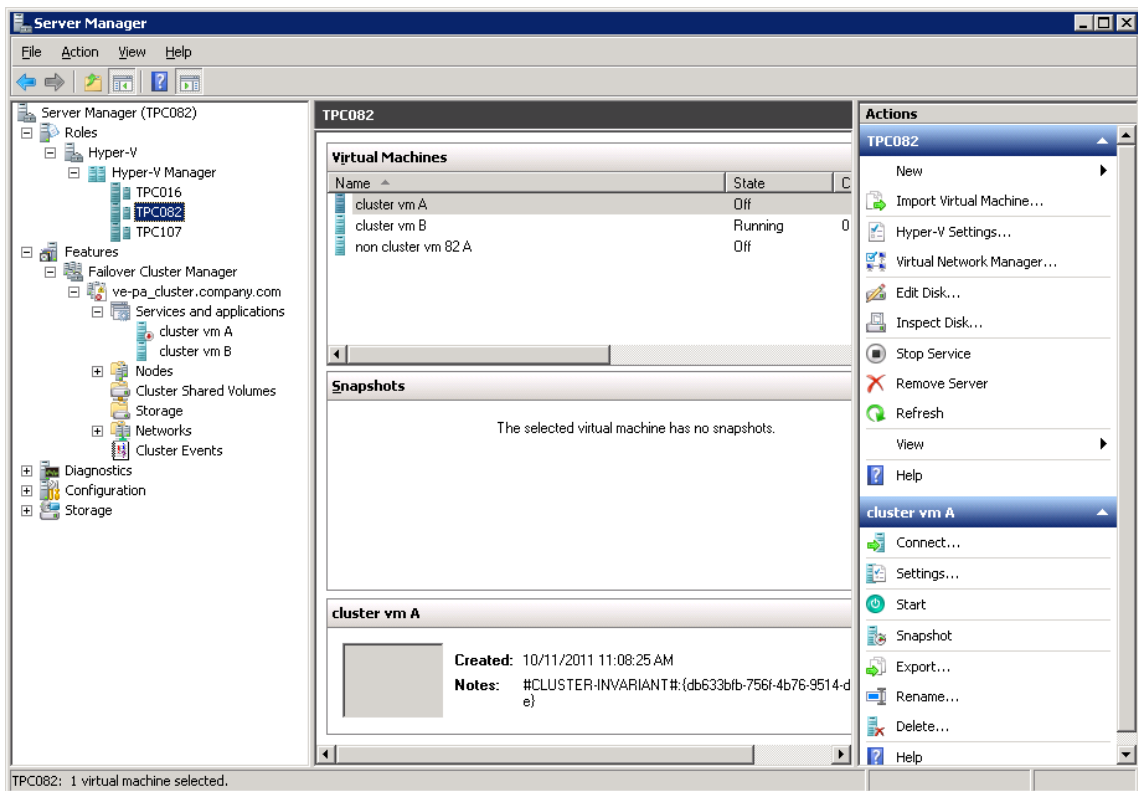
### Example

The following figures illustrate an environment consisting of non-cluster and cluster VMs. The figures show how the VMs are presented in Windows Server Manager and in the Data Protector backup specification.

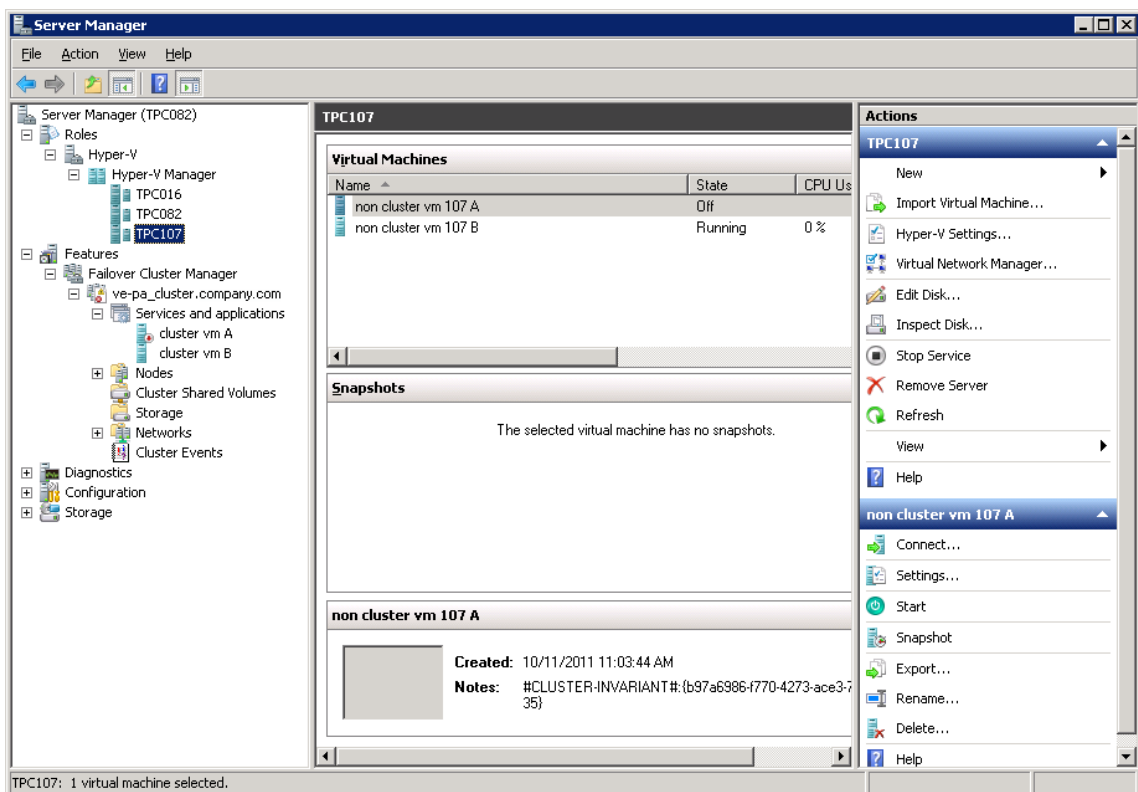
**Figure 82 Windows Server Manager — Cluster VMs**



**Figure 83 Windows Server Manager — VMs running on TPC082**

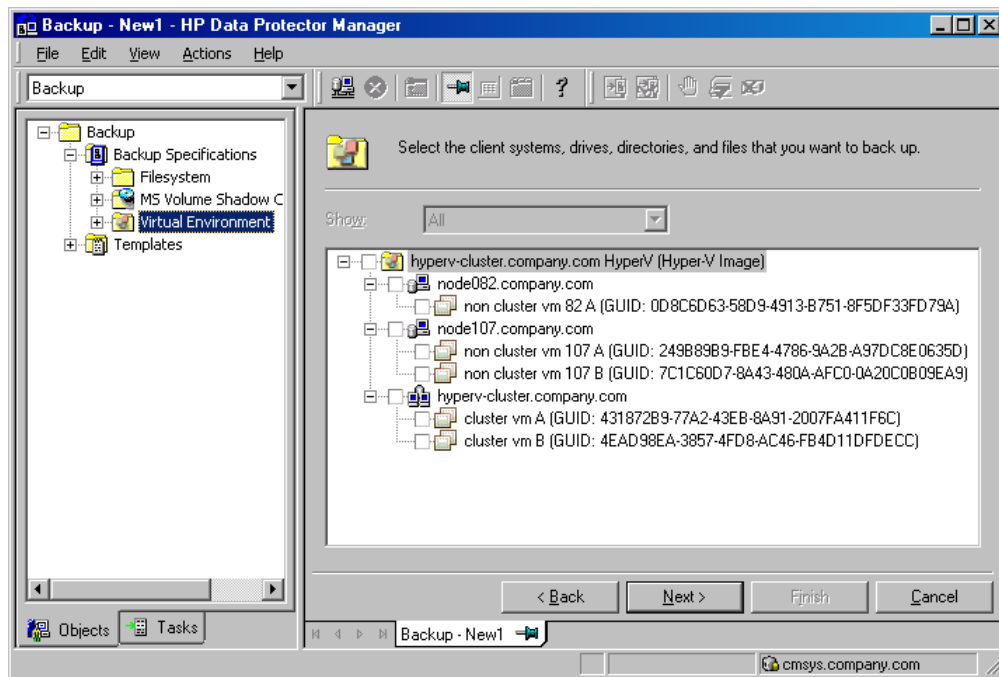


**Figure 84 Windows Server Manager — VMs running on TPC107**



In the Data Protector backup specification, both non-cluster VMs and cluster VMs are displayed for all nodes of the cluster; non-cluster VMs are listed under their physical nodes and cluster VMs under their virtual system. See [Figure 85 \(page 150\)](#).

**Figure 85 Data Protector backup specification**




### Migration of virtual machines

The Data Protector Virtual Environment integration supports migration of virtual machines between Microsoft Hyper-V systems within a clustered environment. After a virtual machine migration in such a cluster, you do not need to change the backup specification; Data Protector will find out where the virtual machine to be backed up has migrated, using the Microsoft Hyper-V WMI services, and back it up.

If the failover occurs during a backup or restore session, the session fails and has to be restarted. Similarly, if a virtual machine is in a live migration process when Data Protector attempts to back it up, the session fails.

**Figure 86 Migration of virtual machines**



## Cluster Shared Volumes


Data Protector supports Microsoft Hyper-V environments that use the Cluster Shared Volumes (CSV) feature of failover clustering. In clusters using CSV, multiple virtual machines can be configured to use the same LUN (disk) while still being able to move from one Microsoft Hyper-V system to another independently of each other.

## Hyper-V Replica

In a replication configuration, a virtual machine (the **primary VM**) is replicated to another Hyper-V server (replica server). If the primary VM gets corrupted, the replication can fail over and the virtual machine replica (**replica VM**) becomes the primary VM.

The replica VM is a snapshot of the primary VM at a given time and can be kept for backup purposes or backed up instead of the primary VM to reduce the load on the host system. The primary VM and replica VM cannot use a common shared volume even if the primary VM resides in a cluster.

**Figure 87 Virtual machine replication**



## Virtual machines on Windows file shares

Data Protector supports VMs on a Windows Server 2012 Hyper-V system, which store their data on a Windows Server 2012 file share (SMB 3.0).

## Data Protector installation components

### Data Protector Cell Manager

The Data Protector Cell Manager can be installed on a Microsoft Hyper-V virtual machine or a separate system outside the virtualization environment.

### Data Protector Virtual Environment Integration component

The Data Protector Virtual Environment Integration component should be installed on at least one client in a Data Protector Cell. This client becomes a backup host.

The Virtual Environment Integration component includes the following major parts:

- `vepa_bar.exe`, used during backup and restore operations on virtual environments.
- `vepa_util.exe`, used during browsing and query operations on virtual environments.
- `vepalib_hyperv`, a dynamic link library providing Hyper-V specific functionality for backup, restore, query and browsing tasks.

### Data Protector Disk Agent component

The Data Protector Disk Agent component must be installed on the backup host if you want to use the Browse button. This button is used to select a restore directory on a backup host.

## Data Protector MS Volume Shadow Copy Integration component

The Data Protector MS Volume Shadow Copy Integration component must be installed on all Microsoft Hyper-V systems that you plan to back up virtual machines from or restore them to. If Microsoft Hyper-V systems are configured in a cluster, it must be installed on all cluster nodes. For VSS transportable backups, the component must also be installed on the backup system. In addition, the component must also be installed on the backup host.

---

**NOTE:** The terms backup system and backup host do not refer to the same entity. A backup host is a system that has the Data Protector Virtual Environment Integration component installed, and therefore controls a backup session, while a backup system is only used in VSS transportable backup sessions; it imports shadow copies from the disk array, provides access to them, and so enables data transfer to backup media.

---

## Data Protector Media Agent component

The Data Protector General Media Agent component must be installed on the Data Protector client to which your backup device is connected. This can be any of the Microsoft Hyper-V systems, or a separate system outside the Microsoft Hyper-V virtualization environment.

## Backup concepts

The Data Protector Virtual Environment integration for Hyper-V is based on the Data Protector Microsoft Volume Shadow Copy Service Integration. By providing its Virtual Environment integration, Data Protector allows you to back up and restore virtual machines without knowing the principles of the Microsoft VSS framework.

## Hyper-V Image backup method

The Hyper-V Image backup method enables you to select Microsoft Hyper-V virtual machines as Data Protector backup objects. A Hyper-V Image backup session progresses as follows:

1. The Cell Manager establishes a connection with a backup host to send a backup request, and starts the virtual environment integration agent (`vepa_bar.exe`) on the backup host.
2. Using a custom dynamic link library (`dpvssapi.dll`, which is part of the MS Volume Shadow Copy Integration component installed on the backup host), `vepa_bar.exe` starts `vss_bar.exe` on the application system.
3. The `vss_bar.exe` agent on the application system, in turn, sends a request to the volume shadow copy service for a volume shadow copy creation. In case of a VSS transportable backup, `vss_bar.exe` also starts `vss_bar.exe` on the backup system.
4. When the volume shadow copy has been created, `vss_bar.exe` passes the volume shadow copy data to the Media Agent client, which transfers the data to backup media.
5. `vss_bar.exe` informs `vepa_bar.exe` that the backup has completed.

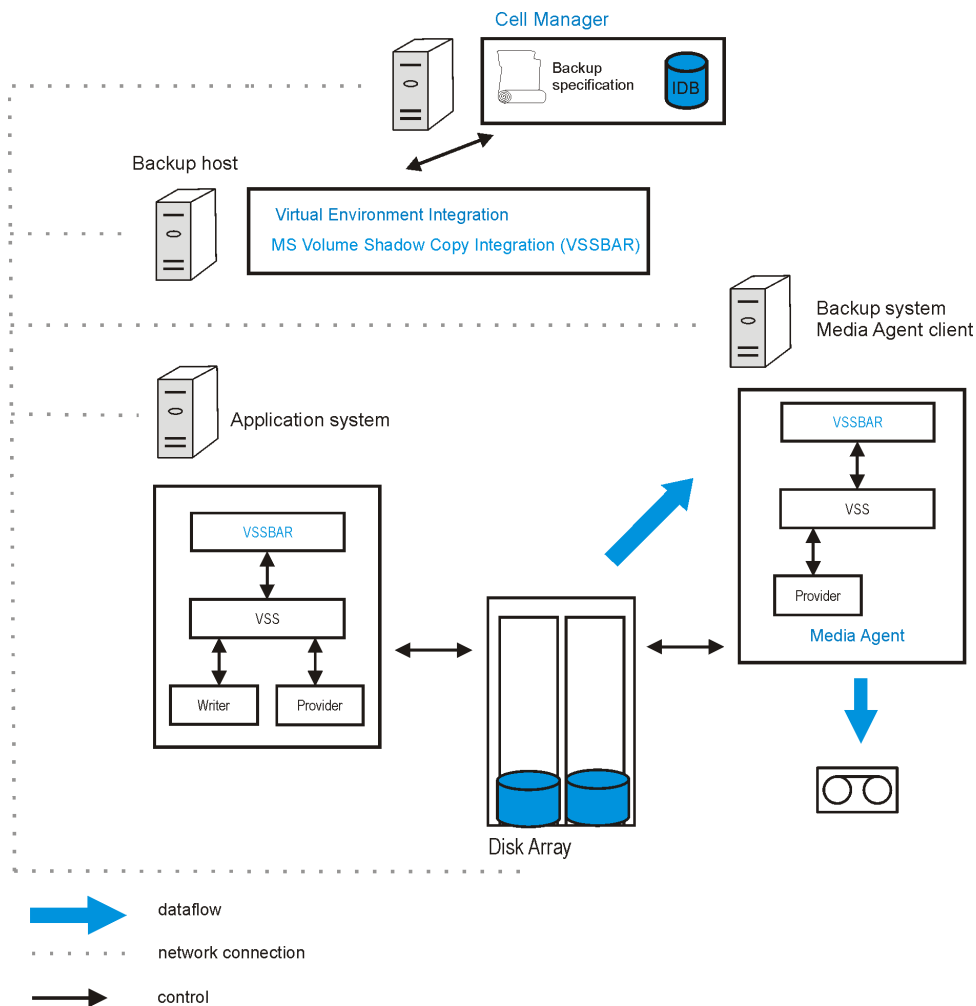
---

**NOTE:** You can use the omnirc option `OB2VEPA_HYPERV_TIMEOUT` to specify how long the `vepa_bar.exe` should wait for `vss_bar.exe` to send the "backup completed" message. By default, there is no time limit (the option is set to `INFINITE`).

---



**Figure 88 Hyper-V Image backup method (VSS transportable backup)**



In “Hyper-V Image backup method (VSS transportable backup)” (page 153), the backup system is also a Media Agent client. It therefore has the General Media Agent component installed and a device connected to it.

For more information, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*, chapter *Integration concepts*.

## Backup types

As both Microsoft Hyper-V and VSS are involved, you must specify the following backup types for each backup session:

- Microsoft Hyper-V backup type
- VSS backup type

The type of backup to be performed is specified at a backup specification level.

## Microsoft Hyper-V backup types

You can select from the following Microsoft Hyper-V backup types:

- **Full (Full)**  
A virtual machine is backed up in its entirety, including its complete virtual disks.
- **Incremental (Incr)**  
A virtual machine snapshot is first created by Microsoft Hyper-V and then backed up by Data Protector. Only changes made since the last backup of a virtual machine are included in the Data Protector backup image.  
Under specific circumstances, the incremental backup session falls back and Data Protector performs a full backup instead. Replicated virtual machines cannot be backed up incrementally. For more information, see [“Backup considerations” \(page 155\)](#).

## VSS backup types

You can select from the following VSS backup types:

- **Local backup (Local or network backup)**  
This type is used for a single-host VSS configuration. The backup is done on the same Microsoft Hyper-V system on which the VSS shadow copy is created. This can be a Data Protector zero downtime backup (using a VSS hardware provider) or standard Data Protector backup (using the software provider).

---

**NOTE:** In a Microsoft Hyper-V cluster environment, a local backup using the VSS hardware provider is not supported. This is a Microsoft product limitation.

---

- **Transportable backup (VSS transportable backup)**  
This type is used for a dual-host VSS configuration. It creates VSS shadow copies on the application system (Microsoft Hyper-V system) and presents them to the backup system, from where the backup to backup media can be performed. A VSS hardware provider is required for this type of backup and such a backup can only be a Data Protector zero downtime backup (ZDB).

---

**NOTE:** In a Microsoft Hyper-V cluster environment, the backup host cannot be a virtual machine running on one of the Microsoft Hyper-V systems. This is a Microsoft product limitation.

---

For details, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

## Quiescence

For Windows virtual machines whose guest operating system is up and running at the time of backup, Data Protector uses the VSS framework to freeze, or quiesce, the states of the applications running within a virtual machine before it is backed up. This ensures data consistency for the programs concerned (an application-consistent backup is created).

---

**NOTE:** For quiescence to function properly, ensure that the Microsoft Hyper-V integration services are installed on all Microsoft Windows VMs.

---

If the guest operating system is not a Windows operating system or if the Microsoft Hyper-V integration services are not installed, quiescence is not possible. Such virtual machines are suspended at the time of backup.

## Quiescence backup flow

1. The Microsoft Hyper-V writer triggers the integration services within the virtual machine.
2. The integration services trigger the VSS framework within the virtual machine to create a volume shadow copy (an *inner volume shadow copy* is created).

3. The VSS framework within the Microsoft Hyper-V system creates a volume shadow copy of the disk on which the virtual machine files are located (an *outer volume shadow copy* is created). In a ZDB use case, a replica storage volume is created.
4. The Microsoft Hyper-V writer performs auto-recovery to make the inner and outer volume shadow copies consistent before proceeding with the backup, as there is a time gap between the inner and outer volume shadow copy creation and some data stored on the virtual machine may change in the meantime.

### Restore chain protection

In Microsoft Hyper-V environments, to prevent data loss during restore, the Data Protector Virtual Environment integration protects the restore chain during incremental backup sessions. For example, whenever a backup snapshot of a virtual machine is created or removed outside Data Protector, or a backup session is removed from the Data Protector Internal Database, the integration detects such a change at the beginning of a subsequent incremental backup session. As a result, Data Protector automatically switches the backup type in the affected session to full. Depending on your Microsoft Hyper-V environment, the fallback extent is an entire Microsoft Hyper-V system or cluster, meaning that the backup type is switched for all virtual machines participating in the session.

You can disable restore chain protection by setting the omnirc option

OB2\_VEAGENT\_DISABLE\_RESTORE\_CHAIN\_PROTECTION to the value 1. See also [“Customizing the Data Protector behavior with omnirc options”](#) (page 167).

### Backup considerations

- Using the Data Protector Virtual Environment integration, you cannot back up Microsoft Hyper-V system configuration data.
- Due to the architectural foundation of the Data Protector Virtual Environment integration, consider the Microsoft Hyper-V writer specifics described in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

### Virtual machine storage

- Due to a Microsoft Hyper-V VSS writer limitation, the following cannot be backed up:
  - Data residing on physical disks that are attached directly to virtual machines.
  - Data residing on storage that is accessed directly by virtual machines using iSCSI initiators.For details, search for the terms “Hyper-V, Planning for Backup” in the [Microsoft TechNet Library](#).
- **Windows Server 2012 systems:** On this operating system, the following cannot be backed up:
  - Data residing on storage that is directly accessed by a virtual fibre channel HBA

## Concurrent sessions

- Backup sessions that use the same devices cannot run concurrently. If multiple sessions are started concurrently, one session waits for the other to complete.

## Cluster Shared Volumes (CSV)

- Backup sessions that try to back up virtual machines from the same CSV cannot run in parallel. If two sessions are started, the second one fails.

**Windows Server 2012 systems:** Virtual machines that are not on a CSV must be backed up in separate sessions, otherwise the session fails. If any part of the virtual machine, such as configuration or snapshot files are outside the CSV, the session will fail.

## Virtual machines on SMB file shares

Data Protector supports backup of VMs that are located on Windows filesystem shares (SMB 3.0). The following prerequisites must be met:

- The Windows role `File Server VSS Agent Service` must be installed and enabled on the Windows file server.

To check if the service is enabled, open a PowerShell window and execute:

```
Get-WindowsOptionalFeature -Online -FeatureName FileServerVSSAgent
```

To enable the service, execute:

```
Enable-WindowsOptionalFeature -Online -FeatureName FileServerVSSAgent
```

Alternatively, open the Server Manager and use the Add Roles and Features wizard to check the status of the role and, if necessary, enable the File Server VSS Agent Service.

- Add constrained delegation for all Hyper-V nodes that are using the SMB file share:  
In Active Directory Users and Computers, select the domain and select **Computers** to list all systems. Right-click the selected Hyper-V node and select **Properties**. In the Properties dialog, select **Delegation** and add the `cifs` service to the list of trusted services.
- The SMB file share must be prepared and configured to host Hyper-V VMs:
  - The VM must be configured and running on an SMB 3.0 file share. Earlier versions of SMB file shares are not supported.
  - Make sure that computer accounts of all Hyper-V hosts, the SYSTEM account, and all Hyper-V administrators managing the virtual machines have Full Control access to the share. Additionally, make sure that all computer accounts of Hyper-V cluster virtual names have Full Control access to the share.  
To avoid manually adding permissions for each additional Hyper-V host or Hyper-V administrator, you can add the Hyper-V hosts and the Hyper-V administrators to a domain security group and grant Full Control permissions to this group instead to single user accounts. You need to create such a domain security group and add all computer accounts and user accounts to this group before you configure the share. When you add permissions to the share, you now only need to add Full Control permission for this security group instead of all single accounts.
  - When you create an SMB file share in a cluster environment, you need to perform the following additional steps:
    - Create a File Server cluster role. Depending on the file server type you select, create the share either within your CSVs in the path `C:\ClusterStorage` (for the

Scale-out File Server for application data file server type) or on a cluster disk (for the File Server for general use file server type).

- If dynamic DNS registrations do not work in your environment, manually register the IP addresses of all cluster nodes for the Distributed Server Name with your DNS Server using the Server Manager.
- When you back up online VMs, single VMs which contain data on an SMB file share and the virtual disks on a local volume, are not backed up. For example, if your backup specification contains two VMs, and VM1 has data on an SMB file share and the virtual disk on a local volume, while VM2 is located completely on an SMB file share, only VM2 is backed up. If multiple VMs that reside on a mixed location (each of them is only at one location) are backed up, then the backup of the entire host is aborted. Create separate backup specifications to back up such a configuration. For example, if your backup specification contains three VMs, and VM1 has data on a local volume, VM2 on a CSV, and VM3 on an SMB file share and all VMs reside on the same system, then the entire system is not backed up. If the backup specification contains both, VMs with data on several locations and VMs with data only at one location, the VMs that have data on another location are skipped first. If the remaining VMs are on the same location, the host is backed up. If not, the backup is aborted. You can deactivate the check for mixed location by setting the omnirc variable `OB2_VEAGENT_ENABLE_SMB_MIX_LOCATIONS` to 1. However, note that although all the VMs are backed up, the complete backup on a given Hyper-V host will fail if there is one online backup performed.
- VMs on SMB file shares using Distributed File System Namespaces (DFSN) or Distributed File System Replication (DFSR) are not supported.
- If a system is both, a Hyper-V server and a File Server, the VM cannot have data on an SMB file share on the same system.

### Incremental backup

- Each Microsoft Hyper-V virtual machine needs to be prepared for incremental backup in advance, before the first incremental backup session is run for it. You can prepare a virtual machine (enable its incremental backup) in two ways:
  - By running a full backup session for it, where the Enable incremental backup of VM(s) option in the related backup specification is selected
  - By running a full backup session for it, after you have executed the `vepa_util.exe` command `--enable-incremental` command for it; for details, see the *HP Data Protector Command Line Interface Reference* or the `vepa_util.exe` man page
- An incremental backup session falls back and Data Protector performs a full backup instead under any of the following conditions:
  - A virtual machine is not prepared for incremental backup
  - Replication to another Microsoft Hyper-V system is configured for a virtual machine selected for backup
  - The reference to a previously created backup image of a virtual machine no longer exists in the Data Protector Internal Database
  - The most recent backup snapshot of a virtual machine was not created by Data Protector

- Your backup specification is configured for virtual machine replicas
- Operating system on your Microsoft Hyper-V system or cluster does not support incremental snapshots of virtual machines; for information if your Microsoft Windows operating systems supports this feature, see the operating system documentation

---

**NOTE:** Depending on your Microsoft Hyper-V environment, the fallback extent is an entire Microsoft Hyper-V system or cluster. If the fallback occurs, the backup type used in the current session changes for all participating virtual machines.

---

## Virtual machine replicas

- **Consistency of virtual machine backup images**

Replica VMs are periodic snapshots of a virtual machine. If you restore a replica VM, the restored VM is only in a crash-consistent state.

Additionally, you can configure the VM replication process to create recovery points and store them. However, these recovery snapshots are also only crash-consistent. To create application-consistent snapshots, enable the replication of incremental VSS copies.

When backing up a replica VM, Data Protector backs up the current snapshot and all stored recovery points. To make sure that each Data Protector backup contains an application-consistent recovery point, select a shorter or equal interval for the incremental VSS replication than for the Data Protector backup. For example, select the VSS incremental replication to be performed every hour and four recovery points must be retained. Schedule the Data Protector backup on every fourth hour.

- **Backing up during the initial VM replication**

Do not back up VMs during the initial VM replication. When backing up the primary VM, the backup session may fail. When backing up the replica VM, the session will also fail. However, if you select the replica VM and use the option **Back up selected VM(s)**, the session will finish without errors but back up only data which was replicated until that time.

- **Backup duration and replication intervals**

During the backup session, the virtual machine configuration is locked and as a result, the VM replication fails. The Hyper-V service attempts to synchronize the replica for a specified time. If it does not succeed, the Hyper-V service stops attempts to automatically resynchronize. You must re-enable replication after the backup session finishes.

Make sure that the backup session finishes in the specified time or change the time interval in which Hyper-V is allowed to automatically resynchronize the replicated VMs. See [“After a Data Protector backup session, the VM replication state is “Error” and health is “Critical””](#) (page 183).

- **VM replication state and health**

- If the replication connection does not function, a backup of the replica VM fails. To avoid this issue, select the option to use the primary replica if the connection does not work. See [Table 28](#) (page 170).
- After a VM failover, the replication state and health are `Error` and `FailedOverWaitingCompletion`, therefore the new replica VM is not used for backup.

After the replication is reversed, the VM replication state and health is `Replicating` and `Normal` but the information provided by Hyper-V is updated only after the next successful VM replication and backup sessions can resume only afterwards.

## Virtual machine migration

- **Windows Server 2012 systems:** During a migration, virtual machines are locked and cannot be backed up, and the other way round.

## ZDB environments

- If your virtual machine files are located on a disk array and you intend to perform ZDB sessions in which the replica storage volumes are kept after the backup session completes, create backup specifications in such a way, that you select all virtual machines that reside on the same storage volume in the same backup specification. Otherwise, the replica storage volumes created in a ZDB session also contain the backups of the virtual machines that were not selected, but these backups are not application-consistent as these virtual machines were not quiesced before the replica storage volumes were created; such virtual machine backups are only crash-consistent. If you perform a ZDB-to-tape session, only the virtual machines that were selected are backed up to backup media.
- **HP P4000 SAN Solutions:** Due to a VSS hardware provider limitation, virtual machines whose files reside on an HP P4000 SAN Solutions disk array cannot be backed up while being online.

There are two workarounds:

- Use the software provider to back up the virtual machines while they are online (standard backup).
- Back up the virtual machines when they are offline (zero downtime backup).

## Object copy considerations

- To properly use the Data Protector object copy feature with Microsoft Hyper-V virtual machine snapshots, you need to copy all Data Protector backup objects created with the Hyper-V Image backup method in the same session. Ensure you select an entire backup session when selecting Data Protector Microsoft Hyper-V objects for interactive, post-backup, or scheduled copy sessions.

To indicate this guideline, the Data Protector GUI does not list Microsoft Hyper-V backup objects in the Copy > Object copy > Interactive > Objects scope of the Object Operations context.

## Restore concepts

You can restore complete virtual machines from backups performed using the Data Protector **Hyper-V Image** method.

## Restore of virtual machines

Virtual machines backed up with the **Hyper-V Image** method can be restored:

- To the default location
- To a different location
- To a directory

### Restore to the default location

When restoring from a **Hyper-V Image** backup, by default, virtual machines are restored to their original Microsoft Hyper-V systems. If a virtual machine to be restored still exists on the original Microsoft Hyper-V system, it is deleted before the files from the backup are restored.

To keep, for example, for safety reasons, the existing virtual machine, you should export it from the system before performing the restore.

---

**NOTE:** If a virtual machine that was exported is re-imported, it is effectively a clone of the original. When you import the machine you are asked if you want to use the old virtual machine GUID. If the original virtual machine, from which the exported machine was created, still exists, keeping the original GUID causes the import to fail. Also after importing, you may have network problems unless you change the IP address of the virtual machine.

---

For the restore, you can also specify:

- Whether virtual machines should be restored to a different Microsoft Hyper-V system
- Whether the restored virtual machines should be powered on

The options in the **Options** page are, by default, set to restore virtual machines with the same names, same GUIDs, to the same Microsoft Hyper-V system (or cluster), and to the same locations.

### Restore in a cluster

When restoring VMs that are configured in a cluster, the cluster node where a restored VM will be running depends on the restore client selection and the state of the environment at the restore time. It does not depend on the state of the environment at the backup time, that is, where the VM was running when it was backed up.

---

**NOTE:** If you use shared cluster disks (not CSV), you can restore a virtual machine only to the node on which the virtual machine is currently registered, as only this node can access the shared cluster disks.

---

### Restore to a different location

To restore the virtual machine to a different location (on the same or another system), select **Restore to a target storage path**. The restore is performed using the Hyper-V writer, therefore the restored virtual machines are functional.

### Restore to a directory

When restoring to a directory (restore outside a Microsoft Hyper-V system) all the virtual machine files can be restored to a directory of your choice (for example, `C:\tmp`) on the selected restore system.

After such a restore, the virtual machines are not functional. You need to import the files to a Microsoft Hyper-V system for the virtual machines to become functional.

## Restore chain validation

When you select an incremental backup session to restore a virtual machine to the corresponding state, Data Protector automatically processes the restore chain, starting with the latest full backup image created before the selected session and then restoring all subsequent incremental backup images up to the selected session. To ensure integrity of the restored data, Data Protector checks the validity of the entire chain before starting the actual data restore. An invalid restore chain results in a failed session with no changes made to the virtual machine and its snapshots.

You can disable restore chain validation by setting the omnirc option

`OB2_VEAGENT_DISABLE_RESTORE_CHAIN_PROTECTION` to the value 1. See also [“Customizing the Data Protector behavior with omnirc options” \(page 167\)](#). Note that unprotected restore chains cannot be validated. Data Protector attempts to validate unprotected restore chains result in failed restore sessions.



## Restore considerations

- Due to the architectural foundation of the Data Protector Virtual Environment integration, consider the Microsoft Hyper-V writer specifics described in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

### Data Protector backup solutions

- Hyper-V backup objects that were created using the Data Protector Microsoft Volume Shadow Copy Service integration cannot be restored using the Data Protector Virtual Environment integration.

### Restore parallelism

- Restore sessions that use the same devices or restore to the same Microsoft Hyper-V system (or cluster) cannot run concurrently.
- Data Protector optimizes each restore session to the highest degree possible. When multiple virtual machines are selected for restore, data from their backup images is restored in parallel provided that the images were created in the same Data Protector backup session.

### Restore to a backup host

- Data Protector clients that are not imported as Hyper-V clients cannot be selected for restore. Import the system you want to use as a backup host as a Data Protector Hyper-V client. If you use the system only for restore purposes, you can omit entering the user credentials during the import. After the import, you can select the client from the Restore client drop-down list.

After the restore, you can remove the Hyper-V client and import it again as regular Data Protector client.

Note that for a backup host, you can select any client that has the Data Protector Virtual Environment Integration and MS Volume Shadow Copy Integration components installed.

### Restore to a different location

- If you restore a replica VM or a primary VM with a custom path configuration (Snapshot File Location and Smart Paging File Location are not set to the default virtual machine location) to a different location using the **Restore to a target storage path** option, the Smart Paging File Location and the Snapshot File Location paths are not updated. If you restore a replica VM that contains recovery points, only the Smart Paging File Location is not updated.

After a restore, check the Smart Paging File Location and the Snapshot File Location paths and update them if necessary.



**TIP:** To make sure that the Snapshot File Location is automatically updated after restore, HP recommends that you configure at least one recovery point for the replica VM.

## Restore of virtual machine replicas

- When selecting a VM for restore, it is not possible to determine whether a primary or replica VM was backed up. The replication state after restore however depends on whether the backed up VM is a primary or replica VM:
  - In case of a primary VM, the replication is disabled.
  - In case of a replica VM, a failover is initiated, so that the replica VM becomes a primary VM.
- **Configuring the network adapter for the replicated virtual machine**

Replica VMs are by default not connected to a network. When restored, it is without a network connection even if it is restored as a primary VM. You should manually map the virtual network switches to the NICs before starting the restored VM.

If the replica VM is connected to a virtual network switch, the virtual switch names used on both the primary and replica side should be the same, otherwise the restored VM cannot be started. To avoid startup problems, connect the network adapter to an existing virtual network switch or remove the connection.
- **Cluster environments**
  - If a primary VM with data on a cluster shared volume (CSV) is replicated to a non-cluster system, the VM machine may not be highly available after restore.

Data Protector by default restores the VMs to the original backup location. For the replica VM, this is a local volume, therefore the VM is restored to a local volume and the existing cluster resource group cannot be brought online, because it is not on a CSV. Data Protector deletes the cluster resource group and reattaches the VM to the cluster. If the restored VM data is residing on a local volume, the Storage resource is not included in the cluster resource group.
  - If the primary VM is replicated to a CSV on another Hyper-V cluster and the paths to the CSV are the same on the primary and replicated side, the VM is highly available after restore.
  - If the VMs are not located on a CSV and the storage location of the restored VM is changed, for example by restoring to a new storage location or restoring a replica VM to the primary cluster, Data Protector cannot recreate the cluster resource group and you must manually delete and recreate it.

## Virtual machines on Windows shares

- During a restore to an SMB file share, the Data Protector VSS agent needs sufficient rights to modify the ACLs of the restored VM files. Therefore you need to run the Data Protector Inet service on the Hyper-V host under a domain user account with appropriate rights. See the Help index: "Inet user impersonation".

Alternatively, add the computer account of the Hyper-V host to the Administrators group of the file server (SMB file share host).
- A virtual machine backed up from an SMB file share can be restored to a different Hyper-V host only if the Hyper-V machine account (*domain\hyperv\$*) has Full Control permissions on the SMB file share and the restored files. You can achieve this by adding the target Hyper-V host in the same domain security group as used by the original Hyper-V host.
- You can restore a virtual machine backed up from an SMB file share or local volume to a different SMB file share. In both cases Data Protector adds the Hyper-V computer domain account (*domain\hyperv\$*) to the ACLs of the restored VM files on the SMB file share with Full Control permissions.

## Configuring the integration

To configure the integration:

- Enable automatic mounting of new volumes. See “Enabling automatic mounting of new volumes on Microsoft Hyper-V systems” (page 164).
- Import and configure Microsoft Hyper-V systems. See “Importing and configuring Microsoft Hyper-V systems” (page 164).
- **Microsoft Hyper-V cluster:** Resolve cluster nodes. See “Configuring Microsoft Hyper-V clusters” (page 164).

## Prerequisites

- Ensure that you have a correctly installed and configured Microsoft Hyper-V virtualization environment.
- Ensure that you have correctly installed Data Protector.

For information on installing Data Protector, see the *HP Data Protector Installation and Licensing Guide*.

- Ensure that you have at least one client in your environment, with both the Data Protector Virtual Environment Integration and MS Volume Shadow Copy Integration components installed. Such a client is called the **backup host**. The backup host must have network access to all the Microsoft Hyper-V systems/clusters. No special configuration is required after installation.

If you intend to restore virtual machine files to a directory on a backup host, also install the Disk Agent component on the backup host. Otherwise, you will not be able to use the **Browse** button to specify the target directory (however, you will still be able to type the directory yourself).

- Ensure that each Microsoft Hyper-V system that you intend to back up from or restore to has the MS Volume Shadow Copy Integration component installed. If your Microsoft Hyper-V systems are configured in a cluster, they must be installed as cluster-aware clients. For details, see the *HP Data Protector Installation and Licensing Guide*.

### **ZDB environments:**

- If the virtual machine files are located on a disk array, ensure that you have configured and installed the corresponding VSS hardware provider.
- For VSS transportable backups, the MS Volume Shadow Copy Integration must also be installed on the backup system.

For details, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*, chapter *Configuration*.

- For supported versions, platforms, devices, and other information, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.

## Limitations

- For restore to a directory (in which case, the backup host is used as the restore client), the application client and the backup host that you use must have the same operating system version installed.
- For VSS related limitations, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.
- For other limitations and recommendations, see the *HP Data Protector Product Announcements, Software Notes, and References*.

## Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the Cell Manager communicates with the Microsoft Hyper-V systems and the backup system properly, configure and run a Data Protector filesystem backup and restore on every Microsoft Hyper-V system and backup host in your environment.

## Enabling automatic mounting of new volumes on Microsoft Hyper-V systems

To be able to perform online backups, enable automatic mounting of new volumes on all Microsoft Hyper-V systems by running the `MOUNTVOL /E` command on each system.

## Configuring Microsoft Hyper-V clusters

Before performing a backup in a Microsoft Hyper-V cluster environment, resolve every cluster node by running the following command:

```
omnidbvs -resolve -apphost HyperVNode
```

The resolution needs to be performed sequentially and, normally, only needs to be performed once. For more information, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

## Importing and configuring Microsoft Hyper-V systems

The Microsoft Hyper-V systems need to be imported into the Data Protector Cell as **Hyper-V** clients.

- ❗ **IMPORTANT:** A client cannot be imported and configured as Hyper-V client, VMware vCenter client, and VMware vCloud Director client at the same time.

---

**NOTE: Clustered environment:** If your Microsoft Hyper-V systems are configured in a cluster, you need to import all cluster nodes and the virtual server as **Hyper-V** clients.

---

### Prerequisites

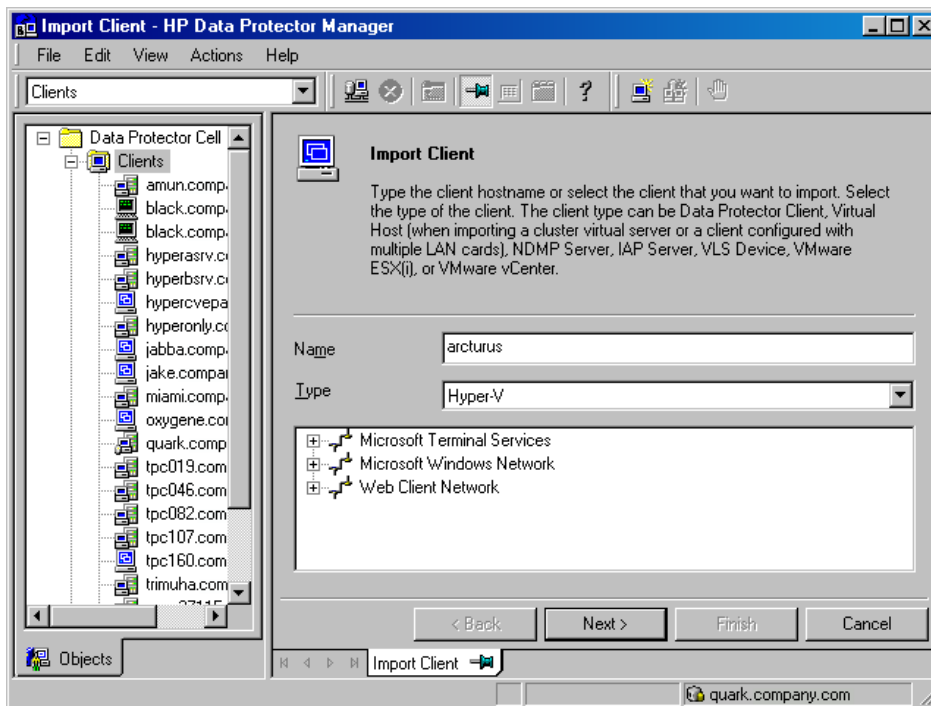
- The User Account Control on the Microsoft Hyper-V system must be configured to elevate user rights automatically.

### Procedure

To import a client into a Data Protector cell:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell**, right-click **Clients**, and click **Import Client**.
3. In the Import client page, enter the client name in the **Name** option, select the **Hyper-V** type from the **Type** drop-down list, and click **Next**.

**Figure 89 Importing a Microsoft Hyper-V client (Name and Type)**



4. Specify login credentials:

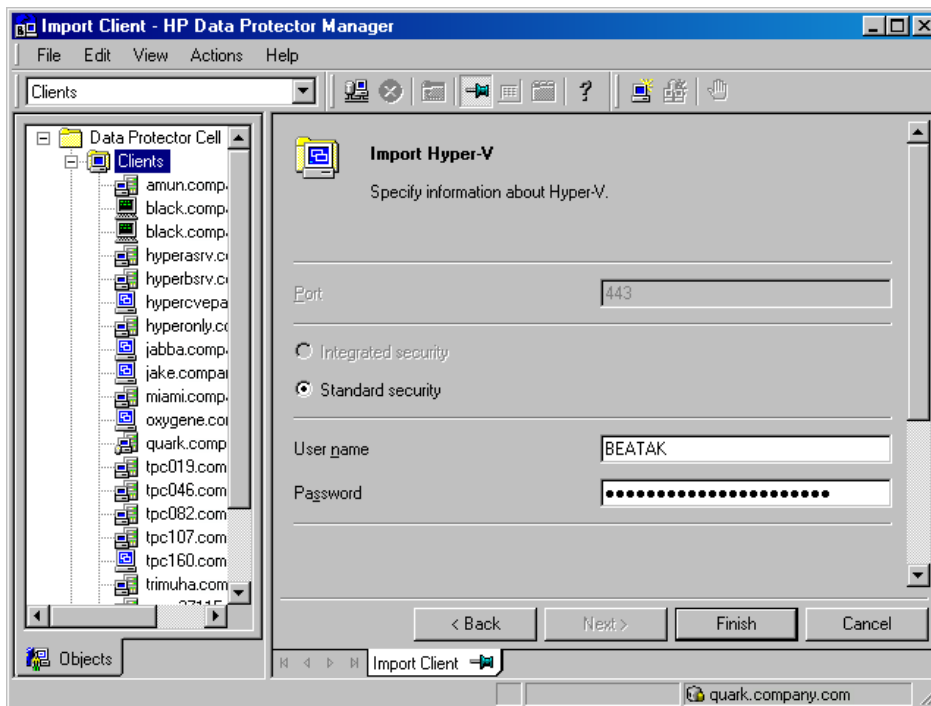
**Username** and **Password**: Specify an operating system user account that has appropriate permissions to access the Microsoft Hyper-V WMI services. You can specify the user account in the following format:

DOMAIN\Username

If the domain is not specified, it is detected automatically. In case of a Microsoft Hyper-V cluster, both the domain and the username must be specified.

Click **Finish**.

**Figure 90 Importing a Microsoft Hyper-V client (Login credentials)**



---

**NOTE:** For details of how to change the parameters later, see [“Changing the configuration of Microsoft Hyper-V systems”](#) (page 166).

---

## Changing the configuration of Microsoft Hyper-V systems

You can only change the credentials for connecting to a Microsoft Hyper-V system if you have the Data Protector Clients configuration user right. For details on the Data Protector user rights, see the *HP Data Protector Help* index: “user groups”.

To change the login credentials, use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Clients** and then select the client for which you want to change the login credentials.
3. In the Results Area, click the **Login** tab.
4. Update the credentials as needed and click **Apply**.

### Using the Data Protector CLI

1. Log on to the backup host.
2. Open a Command Prompt window and change current directory to the directory of the `vepa_util.exe` command.

For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

### 3. Run:

```
vepa_util.exe  
command  
--config  
--virtual-environment hyperv  
--host HyperVClient  
--username Username  
{--password Password | --encoded-password Password}
```

The message \*RETVAL\*0 indicates successful configuration.

For option description, see the *HP Data Protector Command Line Interface Reference* or the `vepa_util.exe` man page.

## Customizing the Data Protector behavior with omnirc options

The omnirc options are useful for troubleshooting or overriding other settings affecting the behavior of the Data Protector client. Options that apply to the Virtual Environment integration have the prefix `OB2_VEAGENT`.

For details on how to use Data Protector omnirc options, see the *HP Data Protector Help* index: “omnirc options”.

## Backup

To back up virtual machines, create a backup specification and then start a backup session. For backup concepts, see [“Backup concepts”](#) (page 152).

## Creating backup specifications


Create a backup specification using the Data Protector GUI (**Data Protector Manager**).

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Virtual Environment**, and click **Add Backup**.
3. In the Create New Backup dialog box, specify the VSS backup type.  
Click **OK**.
4. From the **Client** drop-down list, select the Microsoft Hyper-V system that you want to back up data from. The list contains all clients that were imported into the Data Protector cell as **Hyper-V** clients.

In a cluster environment, select any of the cluster nodes or the virtual server. Regardless of what part of the cluster you select, you will be able to back up all VMs residing in the cluster.

From the **Backup Host** drop-down list, select a system to be used to control the backup. The list contains all clients that have the Data Protector Virtual Environment Integration and the Data Protector MS Volume Shadow Copy Integration components installed.

Figure 91 Selecting a client system and backup host



Click **Next**.

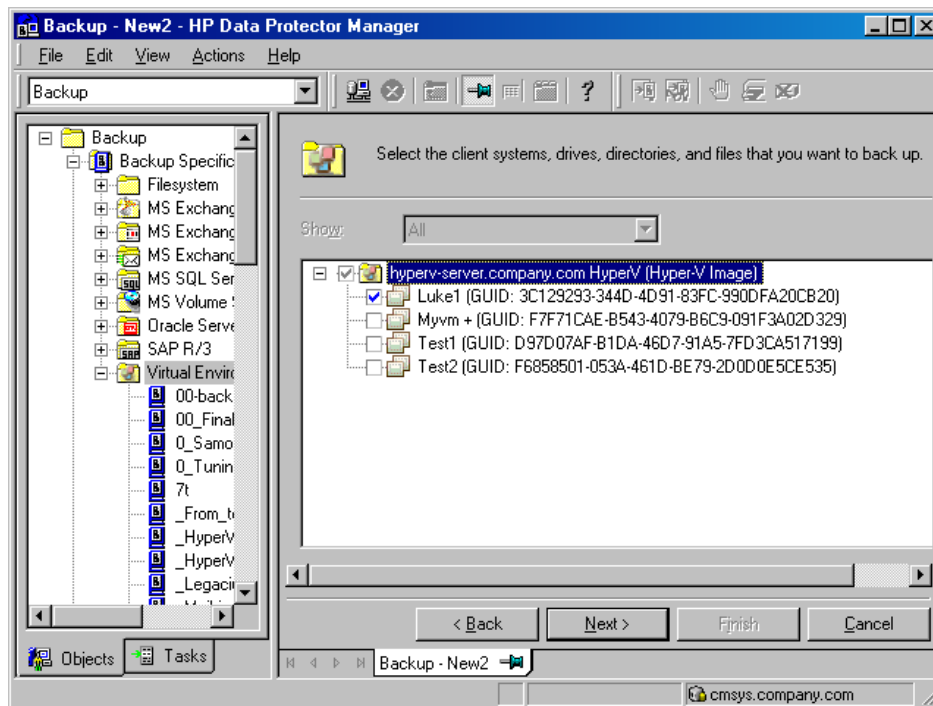
5. This step depends on the type of your environment:
  - With a non-ZDB environment, click **Next**.
  - With a ZBD environment, specify the following options:
    - For a *local or network backup*, select **Use hardware provider**.
    - For a *transportable backup*, select the **Backup system** from where the VSS shadow copies can be backed up to backup media or where you want your shadow copies to be presented and mounted after the backup. The VSS hardware provider is used automatically.
    - Specify other ZDB-specific options. For details, press **F1** or see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*. Note that the **Track the replica for instant recovery** option is not available since instant recovery is not supported.

Click **Next**.

6. Select the virtual machines you want to back up.

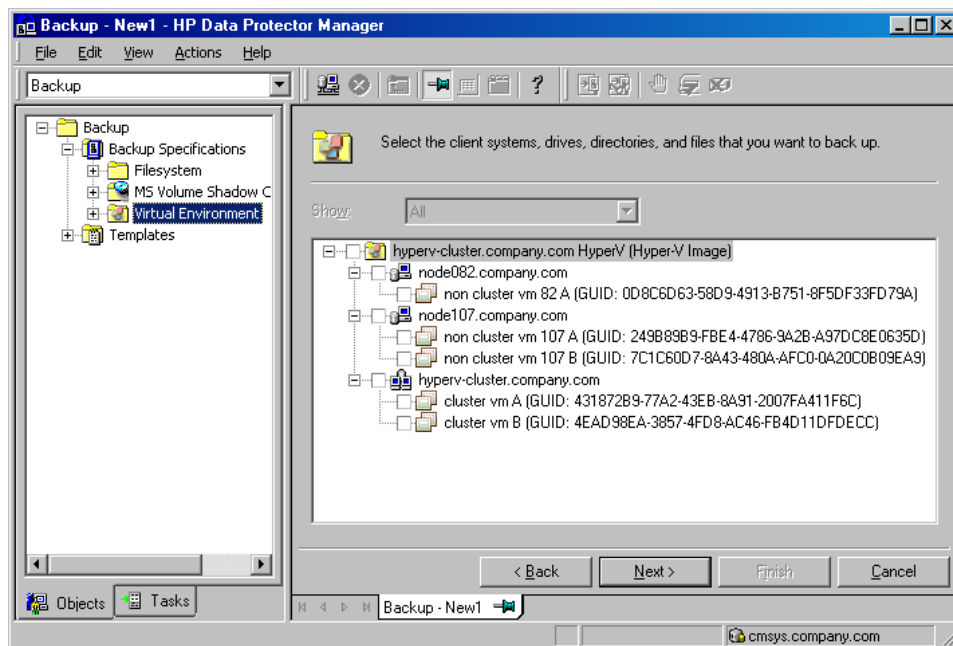


**Figure 92 Selecting the backup objects**



In a cluster, both non-cluster VMs and cluster VMs are displayed for all nodes of the cluster; non-cluster VMs are listed under their nodes and cluster VMs under their virtual server.

**Figure 93 Selecting the backup objects in a cluster**



Click **Next**.

7. Select the devices to use for the backup.

To specify device options, right-click the device and click **Properties**. Set the device concurrency, media pool, and preallocation policy. For details, press **F1** or click **Help**.


You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror**

and **Remove mirror** buttons. Select separate devices for the backup and for each mirror. For details, press **F1**.

Click **Next**.

8. Set backup options. For information on backup specification options and common application options, press **F1**. For information on application-specific options, see [“Application-specific backup options”](#) (page 170).

**Figure 94 Application-specific options**



Click **Next**.

9. Optionally, schedule the backup. See [“Scheduling backup sessions”](#) (page 172).

Click **Next**.

10. Save the backup specification, specifying a name and a backup specification group.

**Table 28 Application-specific backup options**

Option	Description
<b>Pre-exec, Post-exec</b>	<p>Invoke the specified command, script, or application on the backup host before (<b>Pre-exec</b>) or after (<b>Post-exec</b>) the backup.</p> <p>In the specified command lines, do not use double quotes. Type only the name of the executable file and ensure that the file resides in the following directory on the backup host:</p> <p><i>Data_Protector_home\bin</i></p>
<b>Enable incremental backup of VM(s)</b> (available for systems that support incremental backup of Hyper-V virtual machines)	<p>Select this option and run a <i>full</i> backup session to enable Microsoft Hyper-V incremental backup sessions based on the Data Protector backup specification you are configuring. You can then choose between a full or incremental backup of Microsoft Hyper-V virtual machines later on when scheduling or starting a subsequent Data Protector backup session. Incremental backup can only be performed if the virtual machines to be backed up are properly prepared for it. It stores the changes made to the virtual machine after the last backup of any type, and is updated after each backup session.</p> <p>This option prepares all virtual machines that are included in the backup specification. When you clear this option and run a backup session, the virtual machine configurations with regard to incremental backup readiness are left intact.</p>

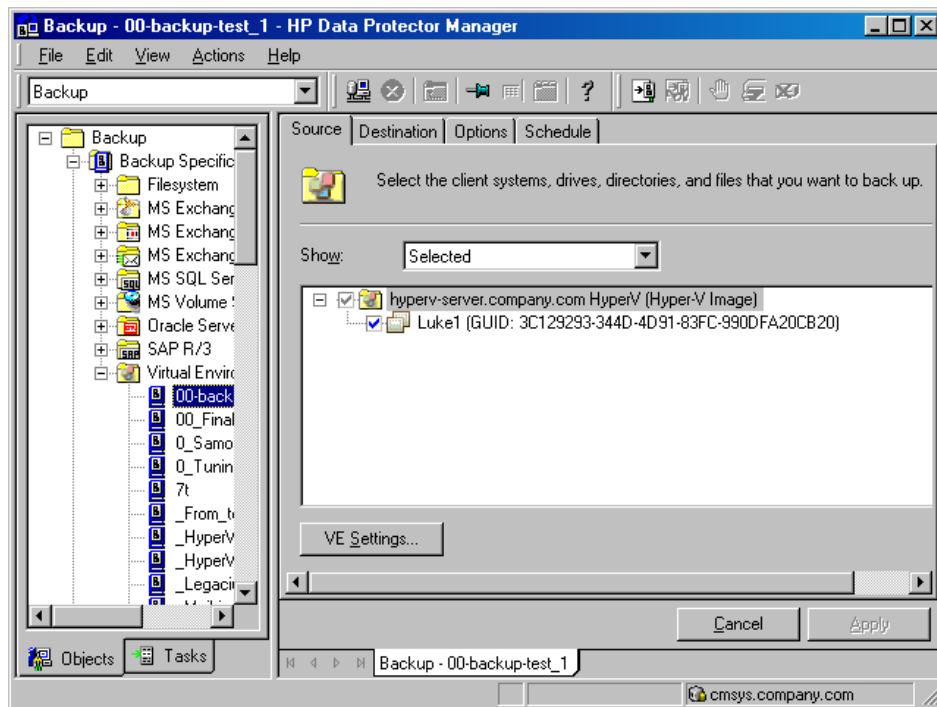
**Table 28 Application-specific backup options** *(continued)*

Option	Description
	When this option is selected, the <b>Back up selected VM(s)</b> option is automatically selected, and the <b>Back up replica VM(s)</b> and <b>Use primary VM(s) if replication link is down</b> options are made unavailable. Default: not selected.
<b>Back up selected VM(s)</b> <i>(available for systems that support virtual machine replication)</i>	Select this option to back up the selected virtual machine (VM), regardless if it is a primary or replica VM. Default: selected.
<b>Back up replica VM(s)</b> <i>(available for systems that support virtual machine replication and when <b>Enable incremental backup of VM(s)</b> is not selected)</i>	Select this option to instruct Data Protector to always back up the replica VM even if a primary VM is selected. If the replication connection is down, the backup will fail even if a primary VM is selected. Default: not selected.
<b>Use primary VM(s) if replication link is down</b> <i>(available when <b>Back up replica VM(s)</b> is selected)</i>	Select this option to instruct Data Protector to back up the primary VM instead of the replica VM if the replication connection does not function. Default: not selected. Selected if <b>Back up replica VM(s)</b> is selected.

## Modifying backup specifications

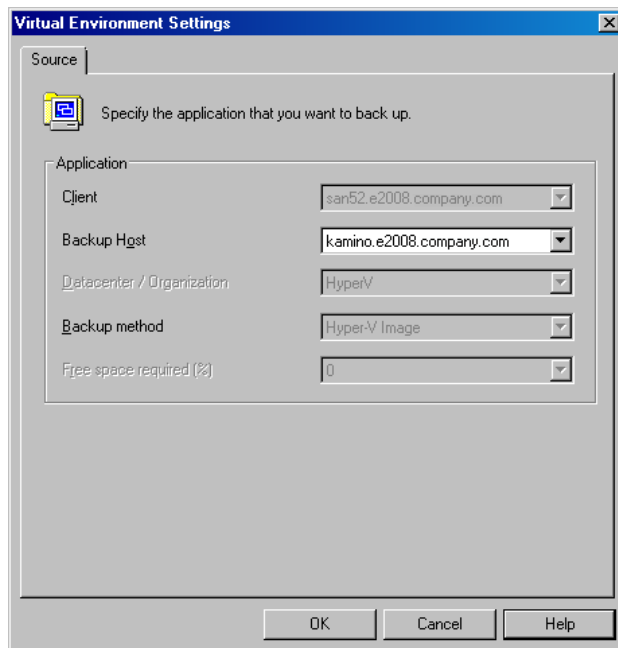
To modify your backup specification, select **Backup** in the Context List and click the backup specification name in the Scoping Pane. In the Results Area, click the appropriate tab and make the desired changes. When you are done apply them by clicking **Apply**.

**Figure 95** Modifying a backup specification



To display the virtual environment settings, in the Results Area, click **VE Settings**. Not all settings can be modified.

**Figure 96 Virtual environment settings**



---

**NOTE:** To see all virtual machines in the source page, not just those you selected, select **All** from the Show list.

---

## Scheduling backup sessions

You can schedule a backup session to start automatically at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

### Scheduling example

The following example applies to a Microsoft Hyper-V environment that supports incremental backup of virtual machines and to backup specifications in which the participating virtual machines are already prepared for Data Protector incremental backup sessions.

To schedule incremental backups at 8:00 and 13:00 and full backups at 18:00 each weekday:

1. In the Schedule property page of the Backup wizard, select the starting date in the calendar and click **Add** to open the Schedule Backup dialog box.
2. Under Recurring, select **Weekly**. Under Time options, type **8:00** or appropriately change the value in the Time box. Under Recurring options, select **Mon, Tue, Wed, Thu, and Fri**. Under Session options, select **Incr** from the Backup type list. See [“Scheduling backup sessions” \(page 173\)](#).  
Click **OK**.
3. Repeat [Step 1](#) and [Step 2](#) to schedule backups at 13:00 and 18:00, each time selecting the appropriate backup type.
4. Click **Apply** to save the changes.

**Figure 97 Scheduling backup sessions**

**Schedule Backup**

Specify the desired backup time, frequency, duration, and type.

**Recurring**

☐ None  
☐ Daily  
☒ Weekly  
☐ Monthly

**Time options**

Time: 8:00 AM  
☐ Use starting  
1/25/2011

**Recurring options**

Every 1 week(s) on

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

**Session options**

Backup type: Incr  
Network load: ☒ High ☐ Medium ☐ Low  
Backup protection: Default

OK Cancel Help

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

To start a backup, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Virtual Environment**.
3. Right-click the chosen backup specification and click **Start Backup**.
4. In the Start Backup dialog box, from the Backup type drop-down list, select either **Full** for full backup or **Incr** for incremental backup.
5. From the Network load drop-down list, select the desired network load for the session.
6. Click **OK**.

At the end of a successful backup session, the message `Session completed successfully` is displayed.

## Using the Data Protector CLI

1. Log on to a Data Protector client where the Data Protector User Interface component is installed.
2. Open a Command Prompt window and change current directory to the directory of the `omnib` command.

For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

### 3. Execute:

```
omnib -veagent_list BackupSpecificationName [-barmode  
VirtualEnvironmentMode] [LIST_OPTIONS]
```

where *VirtualEnvironmentMode* can be either *full* for full backup or *incr* for incremental backup.

If the option *-barmode* is not specified, Data Protector attempts to start a full backup.

For description of the option group *LIST\_OPTIONS*, see the *HP Data Protector Command Line Interface Reference* or the *omnib* man page.

#### Examples

To start a full backup using the backup specification *MyVirtualMachines*, run:

```
omnib -veagent_list MyVirtualMachines -barmode full
```

## Restore

You can restore Microsoft Hyper-V virtual machines using the Data Protector GUI or CLI. For restore concepts, see [“Restore concepts” \(page 159\)](#).

## Limitations

- **Windows Server 2008 systems:** If you perform a restore to a target storage path:
  - The option Power-on VM after restore is not supported. You must manually delete the saved states before you can power on the VM.
  - If the restored VM was part of a cluster and the cluster resource of the VM still existed at the begin of the backup, you must manually recreate the cluster resource after the restore.
- If you perform a restore to a different x64 platform, for example, if the backup was performed on an AMD system and you restore the VM to an Intel system, the system cannot be powered on. You must manually delete the saved states before you can power on the VM.
- Data Protector adds the original VM paths to the specified target storage locations. If the resulting file path is longer than 255 characters Hyper-V cannot access the files although the files are properly restored.

## Finding information for restore

You can retrieve information about backup sessions (such as information on the backup media used and the messages reported during the backup) from the Data Protector Internal Database (IDB).

For retrieval, use the Data Protector GUI or CLI.

## Using the Data Protector GUI

1. In the Context List, click **Internal Database**.

2. In the Scoping pane, expand **Objects** or **Sessions**.  
If you expand **Objects**, backup objects are sorted according to the Microsoft Hyper-V systems for which they were created.

---

**NOTE:** The backup object name contains the virtual machine GUID.

For example, the backup object for the database with the GUID  
4844CA0C-E952-48D9-AE04-C68DDE08F1BR is:

`/%2FHyperV/6/4844CA0C-E952-48D9-AE04-C68DDE08F1BR [VEAgent]`

---

If you expand **Sessions**, backup objects are sorted according to the sessions in which they were created. For example, backup objects created in the session 2011/02/7-7 are listed under 2011/02/7-7.

To view details on a backup object, right-click it and click **Properties**.



**TIP:** To view messages reported during the session, click the **Messages** tab.

---

## Using the Data Protector CLI

1. Log on to a Data Protector client where the Data Protector User Interface component is installed.
2. Open a Command Prompt window and change current directory to the directory of the `omnidb` command.

For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

3. To get a list of backup objects created in the backup session *SessionID*, run:

```
omnidb -session SessionID
```

4. To get details on the backup object *BackupObjectName*, run:

```
omnidb -veagent BackupObjectName -session SessionID -catalog
```

Here is one example of a backup object name:

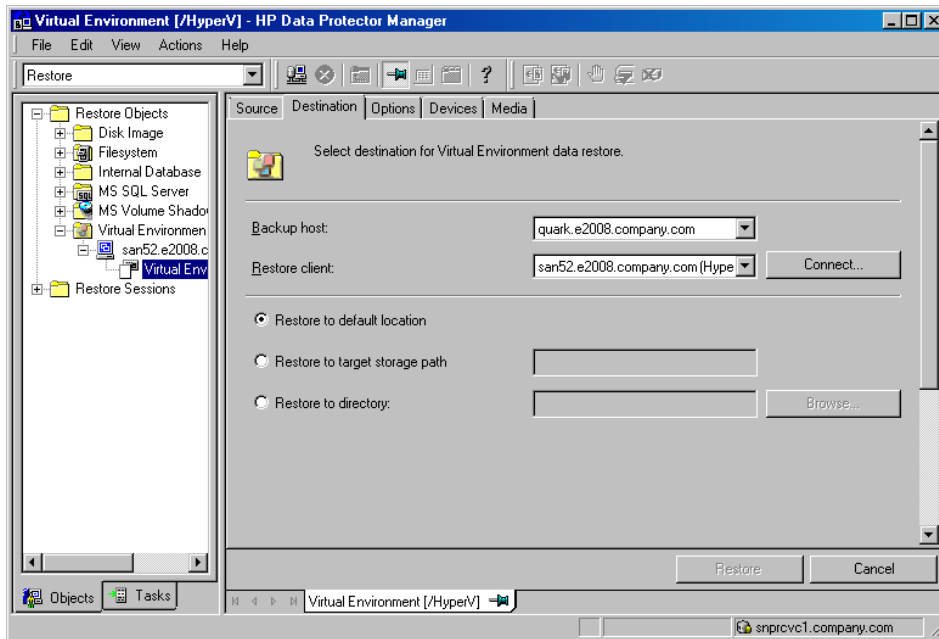
```
quark.company.com:/%2FHyperV/6/4844CA0C-E952-48D9-AE04-C68DDE08F1BR  
[VEAgent]
```

For details, see the *HP Data Protector Command Line Interface Reference* or the `omnidb` man page.

## Restoring using the Data Protector GUI

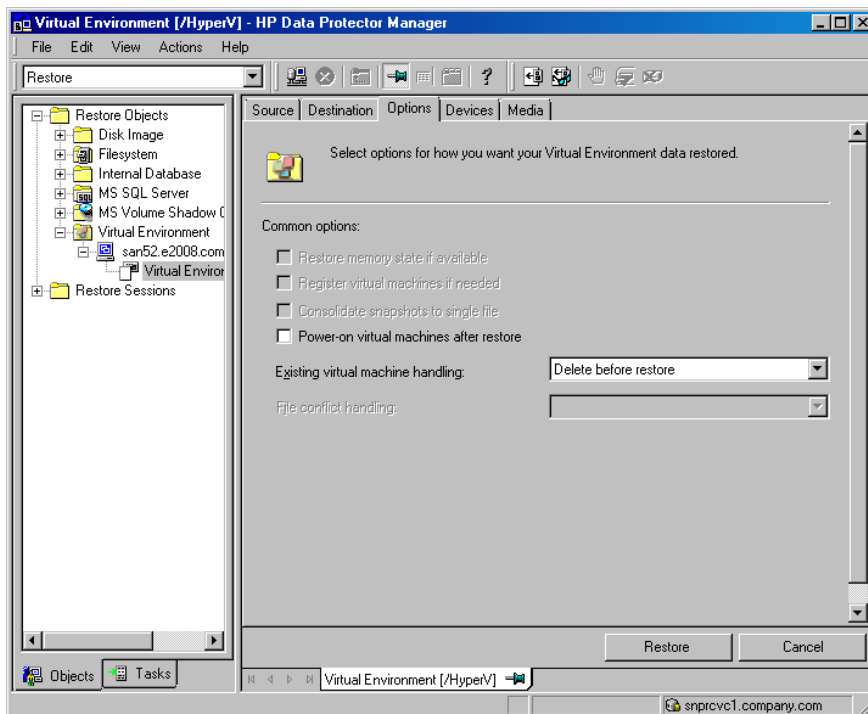
1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Virtual Environment**, expand the client from which you want to restore, and click **Virtual Environment [HyperV]**.
3. In the Source page, select the virtual machines you want to restore.
4. In the Destination page, specify the restore destination. For details, see [“Restore destination” \(page 177\)](#).

Figure 98 Restore destination



5. In the Options page, specify the restore options. For details, see “Restore options” (page 177).

Figure 99 Restore options



6. In the Devices page, specify which devices to use for restore.  
For details, see the *HP Data Protector Help* index: “restore, selecting devices for”.
7. Click **Restore**.
8. In the Start Restore Session dialog box, click **Next**.



- Specify the report level and the network load.

Click **Finish** to start the restore.

At the end of a successful restore session, the message `Session completed successfully` is displayed.

**NOTE:** When you restore a virtual machine to a directory and data from its incremental backup images is also restored during the session, to access folders and files in the restored virtual machine, you have to merge the virtual machine snapshots first. See [“Merging virtual machine snapshots manually”](#) (page 179).

**Table 29 Restore destination**

GUI option omnir command option	Description
<b>Backup host</b> -barhost	Select the Data Protector client that should control the restore. The drop-down list contains all clients that have the Data Protector Virtual Environment Integration and MS Volume Shadow Copy Integration components installed.
<b>Restore client</b> -apphost, -destination	Select the Microsoft Hyper-V client to which the selected virtual machines should be restored to. By default, the client from which the virtual machines were backed up is selected.
<b>Restore to default location</b>	Select this option to restore virtual machines to the original location on the selected restore client.
<b>Restore to target storage path</b> -targetstoragepath	Select this option to restore virtual machines to a different location on the selected restore client.  All VMs are restored to a single volume, even if the VMs were originally on different volumes. For each VM, the original path is appended to the target storage path.
<b>Restore to directory</b> -directory	Select this option to restore virtual-machine files to a directory on selected restore client. You can use the Browse button to find the target directory.

**Table 30 Restore options**

GUI option omnir command option	Description
<b>Power on virtual machines after restore</b> -poweron	Select this option to power the virtual machines on once they are restored. This option is not available when restoring to a directory.
<b>Existing virtual machine handling</b>	Specifies Data Protector's behavior when restoring existing virtual machines.
	<b>Delete before restore</b> -deletebefore  Select this option to delete an existing virtual machine before it is restored, and then restore it from new. Default (GUI): selected.
	<b>Skip restore</b> -skip  Select this option to skip the restore of an existing virtual machine. When restoring multiple virtual machines, selecting this option enables you to restore only the virtual machines that do not exist at restore time.

## Restoring using the Data Protector CLI

- Log on to a Data Protector client where the Data Protector User Interface component is installed.

2. Open a Command Prompt window and change current directory to the directory of the `omnir` command.

For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

3. Run:

```
omnir -veagent
      -virtual-environment hyperv
      -barhost BackupHost
      -apphost OriginalHyperVClient
      [-session BackupID]
      -vm GUID [-vm GUID...]
      [-destination DifferentHyperVClient]
      [-targetstoragepath TargetStoragePathOfAllHyper-V-VMs]
      -directory RestoreDirectory]
      [-poweron]
```

For a description of all the options, see the *HP Data Protector Command Line Interface Reference* or the `omnir` man page.

- 
- ① **IMPORTANT:** A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you have to specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.

The `omnir` syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

---

**NOTE:** When you restore a virtual machine to a directory and data from its incremental backup images is also restored during the session, to access folders and files in the restored virtual machine, you have to merge the virtual machine snapshots first. See [“Merging virtual machine snapshots manually”](#) (page 179).

---

#### Example: Restoring virtual machines to a Microsoft Hyper-V system

Suppose you want to restore the virtual machines VM1 with the GUID 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C and VM2 with the GUID 54C22930-E3B9-43AA-AFCD-1E90BB99F130. At the time of backup, the virtual machines were running on the Microsoft Hyper-V system `hyperv1.company.com`. The virtual machines have been backed up with the Hyper-V Image backup method.

To restore the virtual machines to the Microsoft Hyper-V system `hyperv2.company.com` to the default location, using the backup session 2011/01/11-1, and to power the newly restored virtual machines on, run:

```
omnir -veagent -virtual-environment hyperv -barhost
backuphost.company.com -apphost hyperv1.company.com -session 2011/1/11-1
-vm 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C -vm
54C22930-E3B9-43AA-AFCD-1E90BB99F130 -destination hyperv2.company.com
-poweron
```

#### Example: Restoring virtual machines to a different location

Suppose you want to restore the virtual machines VM1 with the GUID 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C and VM2 with the GUID 54C22930-E3B9-43AA-AFCD-1E90BB99F130. At the time of backup, the virtual machines

were running on the Microsoft Hyper-V system `hyperv1.company.com`. The virtual machines were backed up with the Hyper-V Image backup method.

To restore the virtual machines to the Microsoft Hyper-V system `hyperv2.company.com` to the location `c:\machines`, using the backup session `2011/01/11-1`, and to power the newly restored virtual machines on, run:

```
omnir -veagent -virtual-environment hyperv -barhost  
backuphost.company.com -apphost hyperv1.company.com -session 2011/1/11-1  
-vm 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C -vm  
54C22930-E3B9-43AA-AFCD-1E90BB99F130 -destination hyperv2.company.com  
-targetstoragepath c:\machines -poweron
```

### Example: Restoring virtual machines outside a Microsoft Hyper-V system

Suppose the virtual machines VM1 with the GUID `62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C` and VM2 with the GUID `54C22930-E3B9-43AA-AFCD-1E90BB99F130` that were backed up in the session `2011/02/12-5` from the Microsoft Hyper-V system `hyperv.company.com`, using the Hyper-V Image backup method. To restore the virtual machines outside the Microsoft Hyper-V system, to the directory `C:\tmp` on the backup host `backuphost.company.com`, run:

```
omnir -veagent -virtual-environment hyperv -barhost  
backuphost.company.com -apphost hyperv.company.com -session 2011/2/12-5  
-vm 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C -vm  
54C22930-E3B9-43AA-AFCD-1E90BB99F130 -directory c:\tmp
```

## Merging virtual machine snapshots manually

When a Microsoft Hyper-V virtual machine is restored to a directory and a restore chain is involved, Data Protector does not automatically merge the incremental backup images into their base backup image. Instead, Data Protector restores each backup image to a subdirectory of the selected target directory. To get the volumes of the virtual machine disks to the state as they were at the time the last incremental backup session was run, and gain access to the contained folders and files, you have to manually merge the backup images.

Follow the steps:

1. In the subdirectories of the target restore directory on the backup host, enumerate all virtual hard disk files (files of the type `*.vhd`, `*.vhdx`, `*.avhd`, or `*.avhdx`) whose filenames do not contain the string `ChildVhd`. Copy all such files to a single directory.
2. Launch Hyper-V Manager. In the Actions pane, click **Edit Disk**.
3. In the Edit Virtual Hard Disk Wizard, click **Browse**.
4. In the Open dialog box, browse to the directory with the copied virtual hard disk files.
5. Select the oldest differential virtual hard disk file (file of the type `*.avhd` or `*.avhdx`) and confirm your selection.
6. In the Edit Virtual Hard Disk Wizard, select **Merge** and click **Next >**.
7. Select **To the parent virtual hard disk**, click **Next >**, and confirm your action by clicking **OK**.
8. Repeat steps [Step 5](#) to [Step 7](#) until all virtual hard disk snapshots are merged.
9. Right-click the virtual hard disk and select **Mount** to mount its volume to a drive letter or a directory.
10. Copy the folders and files you need.
11. Using the administrative tool Computer Management, dismount the virtual hard disk volume by right-clicking it and selecting **Detach VHD**.

## Restore of cluster virtual machines

After a restore of a cluster virtual machine that was, for example, deleted, so no cluster resource for that virtual machine is available any more, the restored virtual machine becomes a local virtual

machine running on a node. The reason is missing cluster resource configuration that cannot be restored by Data Protector. The virtual machine requires manual configuration to become a cluster virtual machine again, which is performed as follows:

1. Make sure the virtual machine is powered off.
2. In the Server Manager, expand **Features, Failover Cluster Manager**, and the virtual server, right-click **Services and applications**, and click **Configure a Service or Application**.
3. In the Select Service or Application page of the High Availability Wizard, select **Virtual Machine** and click **Next**.
4. In the Select Virtual Machine page, select the virtual machine you want to configure and click **Next**.
5. In the Summary page, review your planned changes and click **Finish**.

## Restoring a replicated virtual machine

### Re-enabling the replication

After restoring a VM (primary or replica VM), you must manually re-enable the replication:

1. Disable replication for the VM on the primary server. The replication state there will be **Failed**.
2. Enable replication for the VM on the primary server and in the Enable Replication wizard go to the page Choose Initial Replication Method and select **Use an existing virtual machine on the Replica server as the initial copy**.

### Reverting a restored replica VM to an application-consistent recovery point

After the restore of a replica VM, the replication state of the VM is **FailedOverWaitingCompletion**. To revert the VM to a selected recovery snapshot:

1. Use the Hyper-V Manager or the Failover Cluster Manager to cancel the failover. If the latest VM state is the desired restore state, then the VM replication can be directly removed.
2. Initiate a new failover. In the Failover window, select the appropriate recovery point and click **Fail Over**. The VM is reverted to the selected recovery point and automatically started.

## Restoring using another device

You can restore using a different device than used for the backup. For details, see the *HP Data Protector Help* index: "restore, selecting devices for".

## Monitoring sessions

You can monitor currently running sessions from any Data Protector client with the **User Interface** component installed, from the Data Protector GUI. When you run a backup or restore session, a monitoring window shows the session progress. Closing the GUI does not affect the session. You can also monitor sessions using the Monitor context of the GUI.

For details on monitoring, see the *HP Data Protector Help* index: "viewing currently running sessions".

## Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Virtual Environment integration.

For troubleshooting information for the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the *HP Data Protector Help* index: "patches".
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

## Checks and verifications

If your browsing operation or a backup or restore session failed:

- Examine system errors reported in the `debug.log` located in the directory:  
`Data_Protector_home\log`.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HP Data Protector Help*.

## Problems

### Problem

#### Backup session for a Microsoft Hyper-V virtual machine fails

When backing up a Microsoft Hyper-V virtual machine, the session ends unexpectedly with an error similar to the following:

```
[Major] From: OB2BAR_VSSBAR@computer.company.com "MSVSSW"  
Time: 2/1/2011 11:29:03 AM [145:575]  
Writer 'Microsoft Hyper-V VSS Writer' failed to prepare files  
for backup:  
    Reported state:      VSS_WS_FAILED_AT_POST_SNAPSHOT  
    Expected state:      VSS_WS_WAITING_FOR_BACKUP_COMPLETE  
    Failure code:        VSS_E_WRITERERROR_NONRETRYABLE
```

The following are possible causes:

- Automatic mounting of new volumes is disabled on the Microsoft Hyper-V system. For details, see <http://support.microsoft.com/kb/2004712>.
- There are issues inside the virtual machine, such as insufficient storage space for shadow copies, usage of a non-NTFS filesystem, and so on.

### Action

Check if automatic mounting of new volumes is enabled on the Microsoft Hyper-V system. For example:

```
diskpart.exe  
Microsoft DiskPart version 6.1.7600  
Copyright (C) 1999-2008 Microsoft Corporation.  
On computer: TPC021
```

```
DISKPART> automount  
Automatic mounting of new volumes enabled.
```

If automatic mounting is disabled, enable it by running the following command:

```
MOUNTVOL /E
```

If the issue persists even though automatic mounting is enabled, check the application log files inside the virtual machine to determine the cause.

## Problem

### In a CSV environment, a backup session fails

When backing up virtual machines with files located on Cluster Shared Volumes, the session fails with an error similar to the following:

```
[Major] From: OB2BAR_VSSBAR@tpc049.company.com "HyperV"   Time: 2/18/2011 3:53:09 PM
Cannot perform backup of:
/Microsoft Hyper-V VSS Writer/Virtual Machines/Backup Using Child Partition Snapshot\vmw39192',
which contains data in:
    C:\ClusterStorage\Volume3\vmw39192\Virtual Machines\1FC08961-08B8-4AC5-BDE8-AF4E2AAA07E8.xml
    C:\ClusterStorage\Volume3\vmw39192\Virtual Machines\1FC08961-08B8-4AC5-BDE8-AF4E2AAA07E8\*
    C:\ClusterStorage\Volume3\vmw39192\vmw39192.vhd
[Major] From: OB2BAR_VSSBAR@tpc049.company.com "HyperV"   Time: 2/18/2011 3:53:09 PM
There is no data to be backed up.
```

This problem appears if two sessions attempt to back up virtual machines from the same CSV at the same time. Since the first session locks the CSV, the second session cannot access the CSV and consequently fails.

## Action

Run the sessions one after the other.

## Problem

### Browse operation or a backup or restore session fails due to a wrong password

When the application system and the backup host are the same system, the configuration check of the username and password is not performed. Actually, in such a setup, the configuration check always appears to succeed even when the credentials are invalid, and the subsequent browse, backup, and restore operations will function. However, when a non-local connection is established, for example, when a different backup host is used, the operations will fail due to a wrong password.

## Action

Configure backup host that is not the application system at the same time.

## Problem

### The Data Protector Inet service configuration is missing

During a restore session, a message similar to the following is reported in the session output:

```
[Warning] From: INET_thread_vepa_bar.exe@tpc040.company.com
"tpc040.company.com"   Time: 2/18/2011 5:24:05 PM
    No data for user ADMINISTRATOR@DOMAIN in Inet's impersonation
configuration.
```

## Action

None. You can safely ignore the message.

## Problem

### Data Protector cannot restore a file

During a restore session, a message similar to the following is displayed:

```
[Major] From: OB2BAR_VSSBAR_COMP@tpc021.company.com
"HyperV"   Time: 3/3/2011 2:05:35 PM
Cannot restore file
'C:\ProgramData\Microsoft\Windows\Hyper-V\Virtual Machines
\8F74F8EE-E93A-4CD3-998F-3324784ED932.xml'.
Failed with error: '[33] The process cannot access the file because
another process has locked a portion of the file.'
```

```
[Major] From: OB2BAR_VSSBAR_COMP@tpc021.company.com "HyperV"
Time: 3/3/2011 2:05:35 PM
[145:221]   Restore of component
```

```
'/Microsoft Hyper-V VSS Writer/Virtual Machines/  
Backup Using Child Partition Snapshot\VM_name' failed.
```

### Action

Restart the restore without cleaning up the files that have already been restored.

### Problem

#### **During restore of data on CSV, data is sent through the LAN instead of the SAN**

Due to a Microsoft limitation, during restore of virtual machine files on Cluster Shared Volumes (CSV), in some cases data is sent through the LAN instead of the SAN, which may increase the restore time. This happens when a node other than the coordinator node is used for data transfer.

### Action

Before starting a restore, make sure the coordinator node is used for the restore. Perform either of the following:

- In Windows Server Manager, check which node is currently the coordinator node, and select that node for the restore in Data Protector.
- In Windows Server Manager, manually switch coordination to the node you wish to use for the restore.

### Problem

#### **Warning event ID 5605 is logged in the Windows Application Event Log**

The warning appears when the Virtual Environment integration for Microsoft Hyper-V connects to the cluster namespace. The session completes successfully nevertheless.

### Action

None. You can safely ignore the warning. For details, see the Microsoft KB article 2590230 at [support.microsoft.com/kb/2590230](https://support.microsoft.com/kb/2590230).

### Problem

#### **After a backup session, the VM replication state is "Error" and health is "Critical"**

During a Data Protector backup session the Hyper-V service locks the virtual machine to prevent inconsistencies. On the primary Hyper-V server it logs an error and a warning in the Windows Event Log:

```
Hyper-V failed to replicate changes for virtual machine 'vmw38078'  
(Virtual Machine ID FBCDBFDF-ACBD-4A01-BA8D-A4A9CF597651). Hyper-V will  
retry replication after 5 minute(s).
```

On the replication server side, the Hyper-V service logs multiple errors in the Windows Event Log:

```
Hyper-V failed to apply replication logs for 'vmw38078': Operation  
aborted (0x80004004). (Virtual Machine ID  
FBCDBFDF-ACBD-4A01-BA8D-A4A9CF597651)
```

By default, the Hyper-V service retries the replication every five minutes for a maximum of one hour. If it is not possible to perform a successful replication within this time frame because the backup of the VM takes longer than that, the replication remains in the `Error` state and with health status set to `Critical` and the Hyper-V service no longer attempts to automatically resynchronize the VMs.

### Action

Change the time interval in which Hyper-V is allowed to automatically resynchronize replica VMs. For example, to change the interval to five hours, execute the following PowerShell command:

```
Set-VMReplication -VMName VMName -AutoResynchronizeEnabled $true  
-AutoResynchronizeIntervalStart 00:00:00 -AutoResynchronizeIntervalEnd  
05:00:00
```

## Problem

### Backup session fails when triggered after a restore session that involves a restore chain

After you start a Data Protector session that restores data from incremental backup images, a subsequent backup session fails after reporting an error similar to the following in the session output:

```
[Major] From: OB2BAR_VSSBAR@computer.company.com "HyperV" Time: 5/14/2012 10:50:18 AM  
[145:575] Writer 'Microsoft Hyper-V VSS Writer' failed to prepare files for backup:  
Reported state: VSS_WS_FAILED_AT_POST_SNAPSHOT  
Expected state: VSS_WS_WAITING_FOR_BACKUP_COMPLETE  
Failure code: 0x80042336
```

When you inspect the operating system's event log, there is an event of the Windows Logs category with ID 10145 logged in it.

The root cause of the failure may be Data Protector's prevention of inconsistent data from being backed up. When Data Protector processes a Microsoft Hyper-V restore chain, thus one or more incremental backup images, data from each such image is merged into the base backup image. The automatic merge process is performed by Microsoft Hyper-V and is asynchronous to the Data Protector restore session. It may continue to run in the background after the session finishes. When a backup session for the involved virtual machines is invoked before merging is complete, Data Protector detects it and aborts the process to avoid inconsistencies in the backup image.

## Action

Restart the problematic backup session at a later time or adjust the session schedule accordingly.

## Problem

### Restore session fails when restoring to an SMB share with insufficient privileges

When restoring to an SMB file share Data Protector needs sufficient privileges to change the owner of files on the File Server. If the Data Protector agent is not allowed not modify the owner, then an error is reported similar to the following:

```
[Major] From: OB2BAR_VSSBAR_COMP@hyperv.example.com "HyperV" Time:  
5/13/2013 3:23:09 PM
```

```
Cannot restore file '\\fileserver.example.com\share4vms\local_vm\Virtual  
Machines\A9321AEE-EC69-4F79-BB34-59BA97D83CAC.xml'.
```

```
Failed with error: '[1307] This security ID may not be assigned as the  
owner of this object.'
```

## Action

Run the Data Protector Inet service on the Hyper-V host under a domain user with sufficient privileges, or add the computer account of the Hyper-V host to the Administrators group of the File Server.



---

## Part III Citrix XenServer

Data Protector offers a script solution to back up Citrix XenServer data online.

### Citrix XenServer

- ***Data Protector Citrix XenServer script solution***

This solution is based on the standard Data Protector filesystem backup functionality.

See [“Data Protector Citrix XenServer script solution” \(page 186\)](#).

---

**NOTE:** You can also back up Citrix XenServer virtual machines using the common Data Protector filesystem backup functionality, which operates on the file level. Consequently, the smallest object that you can back up or restore is a file. However, to ensure data consistency, you must take virtual machines offline before starting a backup session.

---

---

## 4 Data Protector Citrix XenServer script solution

### Introduction

This chapter describes a script based solution that, when used together with standard Data Protector functionality, enables online and offline backup and restore of **Virtual Machines** hosted on a **Citrix XenServer**. For the supported version of Citrix XenServer for your version of Data Protector, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.

### Integration concepts

#### Types of backup

Using this integration, you can perform both online and offline backups. In both cases, only full backups of virtual machines are available.

#### Online backup

Online backups are performed without stopping the virtual machines concerned, using the snapshot functionality available on the Citrix XenServer. By creating a snapshot of a virtual machine, a copy of the whole virtual machine is created, using the capabilities of the underlying storage repository and is saved as a template virtual machine. This template can subsequently be used to restore the disks of the original virtual machine, or to create new virtual machines as exact copies of the original at the time of the snapshot.

- **online.** The user selects which virtual machine, or virtual machines, Data Protector should back up without disturbing their power state(s).
- **allOnline.** Data Protector backs up all virtual machines present in the XenServer without disturbing their power states.

---

**NOTE:** The types of storage repositories for which online backup is supported is dependent on the version of XenServer that you are using. See “[Special considerations](#)” (page 199).

---

Online backup is available for Virtual Hard Disk (VHD), Netapp and EqualLogic storage repositories.

#### Offline backup

Offline backups are performed after first shutting down the virtual machines concerned. The complete virtual machine is then saved as an `.xva` file using the Citrix export functionality. This file can subsequently be imported into any XenServer to create a new machine equivalent to the one exported.

- **offline.** The user selects which virtual machine, or virtual machines, should be suspended and backed up.
- **allOffline.** All the virtual machines present in the XenServer are sequentially suspended and backed up.

#### Disaster recovery

For disaster recovery, you need a backup of the disks that a virtual machine uses plus the metadata for the virtual machine. With this integration, this is only possible with offline backup.

## Backup processes

### Online backup

The process used by the integration script for an online backup is as follows:

1. Check that snapshot operation is allowed for all the disks of the virtual machine.
2. Snapshot the virtual machine.
3. Export the snapshot to the temporary backup folder as an `.xva` file.
4. Delete the snapshot.

After the script has finished, the contents of the temporary backup folder are backed up to the backup device using standard Data Protector functionality.

### Offline backup

The process used by the integration script for an offline backup is as follows:

1. If the virtual machine to be backed up is running, stop it.
2. Export the virtual machine to the temporary backup folder as an `.xva` file.
3. If the virtual machine was running before the backup, restart it.

After the script has finished, the contents of the temporary backup folder are backed up to the backup device using standard Data Protector functionality.

---

**NOTE:** The file type of the exported file for offline backup is the same as that for online backup, but it is possible to differentiate between offline and online backups from the filename. See [“Specifying a restore” \(page 197\)](#).

---

## Types of restore

The type of restore that can be performed is dependent on the type of backup that was performed: Restore from online backup requires a different process to restore from offline backup.

The following restore possibilities are available:

- **Restore from online backup.** In a similar way to online backup, this has two options available:
  - **online.** This mode allows a user to restore the disks of selected virtual machines that were backed up in a session using backup mode `online` or `allOnline`.
  - **allOnline.** This mode allows a user to restore the disks of all the virtual machines backed up in a session using backup mode `online` or `allOnline`.
- **Restore from offline backup.** In a similar way to offline backup, this has two options available:
  - **offline.** This mode allows a user to restore selected virtual machines that were backed up in a session using backup mode `offline` or `allOffline`.
  - **allOffline.** This mode allows a user to restore all virtual machines that were backed up in a session using backup mode `offline` or `allOffline`.

---

**NOTE:** The disks or virtual machines available for restore can be influenced by selections made within the Data Protector GUI. See [Figure 105 \(page 198\)](#).

---

## Restore processes

### Restore from online backup

First, Data Protector restores the selected contents from the relevant backup session to the temporary backup folder.

Then, the integration script uses the following process to perform the restore from online backup:

1. Ensure that the virtual machine is present and stopped.
2. Delete the old disks of the virtual machine.
3. Import the `.xva` file containing the snapshot from the temporary backup directory.
4. Restore the disks of the virtual machine using the disks from the snapshot.

## Restore from offline backup

First, Data Protector restores the selected contents from the relevant backup session to the temporary backup folder.

Then, the integration script uses the following process to perform the restore from offline backup:

1. If present, ensure that the virtual machine is stopped.
2. Delete the old virtual machine.
3. Import the `.xva` file from the temporary backup directory, activating the restore option.

The import operation of XenServer converts the `.xva` file into a virtual machine and sets the virtual machine in the same power state as it was during the export. When the restore option is set to true, the MAC of the old virtual machine is preserved

---

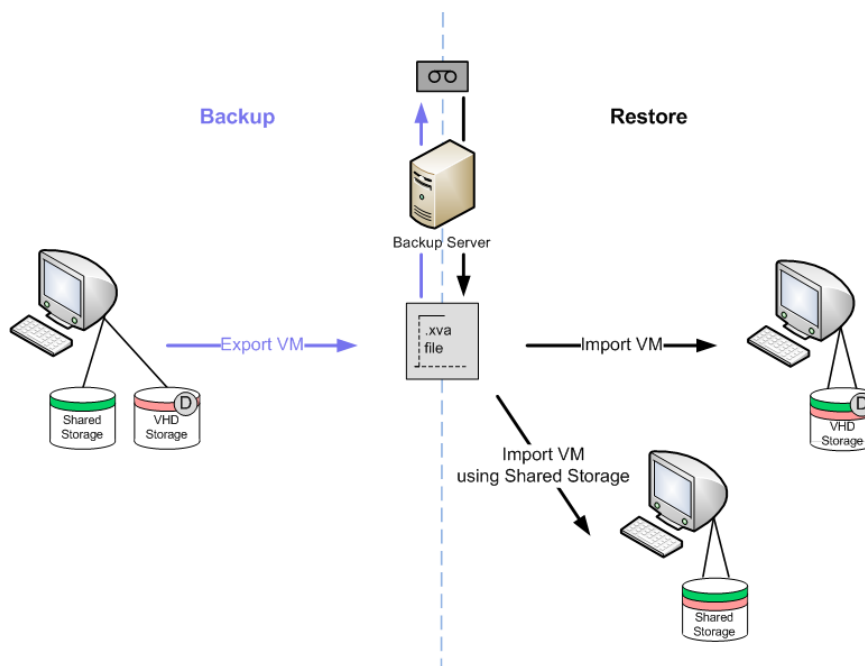
**NOTE:** The file type of the backup file for offline backup is the same as that for online backup, but it is possible to differentiate between offline and online backups from the filename. See “Specifying a restore” (page 197).

---

## Restore considerations

When a restore is performed, the original storage repositories used by the virtual machine are not known, so the user must specify in which storage repository the restore should be performed. If no storage repository is specified, the virtual machine will be restored in the default storage repository. Consider Figure 100 (page 188).

**Figure 100 Example backup and restore**



On the left side is a virtual machine with two disks: One disk uses the shared storage repository and the other the default storage repository.

When this virtual machine is backed up, an `.xva` file is produced.

When the virtual machine is subsequently restored, two options are available:

- A storage repository is not specified. In this case, both disks are restored into the default storage repository.
- A storage repository is specified. In this case, both disks are restored into the specified storage repository. For example, if the Shared Storage repository is specified, both disks are restored into the Shared Storage repository.

This means that the virtual machine can be restored successfully, but it is only possible to restore the disks to a single storage repository, regardless of the original configuration.

## Main integration components

To function, this integration requires the following components:


- Supported Citrix XenServer
- HP Data Protector client(s) with:
  - Disk Agent
  - Media Agent
- Integration scripts

A simple example layout for these components (with both Disk Agent and Media Agent on the same Data Protector client) is shown in [Figure 101 \(page 189\)](#).

No scripts are installed on the Citrix XenServer, but its export facility is used by the integration scripts. The integration scripts themselves are all installed on the Data Protector client with the Disk Agent installed.

During backup, virtual machines, or virtual machine disks, are exported to the Disk Agent host and backed up using Data Protector. The device used for the backup can be any device supported by Data Protector for file-system backup.

**Figure 101 Example layout**



## Installation of the integration

This section gives detailed instructions for the installation of the XenServer integration.

### Prerequisites

The machine on which the integration is to be installed must be a Data Protector client (shown as DP Client in [Figure 101 \(page 189\)](#)) with the Data Protector Disk Agent installed and must be a member of a Data Protector Cell. For further details, see the *HP Data Protector Installation and Licensing Guide*.

The Data Protector Disk Agent must be installed first, since other software must be installed in specific Data Protector directories.

The following software must then be installed on the Disk Agent host:

- Python version 2.5.x or 2.6.x (available from <http://www.python.org/download>). This must be installed first. (Also, see the note regarding Pycurl below.)

Install the software in the following directories:

**Windows systems:** `Data_Protector_home\bin`

**Linux systems:** `/opt/omni/bin`

- Curl with libcurl version 7.18.x or 7.19.x (available from <http://curl.haxx.se/download.html>). Note that the package installed must contain libcurl.
- Pycurl (available from <http://pycurl.sourceforge.net/>).

If you prefer there are some compiled versions available from:

**Windows systems:** <http://pycurl.sourceforge.net/download/>

**Linux systems:** <http://rpm.pbone.net/index.php3>

---

**NOTE:** The installers may require a Python version earlier than the latest available, so it is best to check the situation before installing Python. For example, the fact that the `pycurl-ssl-7.18.2.win32-py2.5.exe` contains the string `py2.5` implies that python 2.5 is required.

---

- `xenAPI.py` for the version of XenServer that you are using (available from <http://community.citrix.com/display/xs/Download+SDKs>)



**TIP:** Note the full installation directory of python. You will need this later when calling the Data Protector Pre-exec and Post-exec scripts. For this documentation it is assumed to be:

**Windows systems:** `Data_Protector_home\bin\Python25\`

**Linux systems:** `/opt/omni/bin/Python25/`

---

## Installation of the integration scripts

The following scripts are supplied for the integration:

- `DPxen_config.py`
- `DPxen_operations.py`
- `DPxen_backup.py`
- `DPxen_restore.py`
- `DPxen_postbackup.py`

They are available as follows:

- For Data Protector 8.00, on the Data Protector 8.00 installation DVD-ROM in the directory:

**Windows systems:** `Windows_other\Xen_sup`

**Linux systems:** `Xen_sup`

- For other versions of Data Protector for which the integration is supported, as an installation package that can be downloaded:
  1. Go to <http://www.hp.com/go/dataprotector>.
  2. Click **SUPPORT & DRIVERS** at the top of the webpage.
  3. On the Support & Drivers webpage, click **Drivers & Software**.
  4. In the text box, enter Data Protector and click **SEARCH**.
  5. In the search results list, click the Data Protector version with which you will be using the integration.

6. Select **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**
7. In the Software – Solutions section, click **Data Protector Release <Version> / Xen Server integration packet** and download the installation package from the subsequent page.

•

When you have obtained the scripts, install them as follows:

1. If you have downloaded the package, unzip `xenServer_backup_solution.zip`.
2. Copy all of the scripts to your preferred location on the Data Protector Disk Agent host. For security reasons, the following locations are recommended:

**Windows systems:** `Data_Protector_home\bin`

**Linux systems:** `/opt/omni/bin`

Make a note of this location: You will need it later.

3. If you do not already have the `xenAPI.py` file:
  - a. Go to <http://community.citrix.com/display/xs/Download+SDKs>.
  - b. Go to the download section for the version of XenServer that you are using.
  - c. Download the `XenAPI.py`.
4. Copy `XenAPI.py` to the same location as the integration scripts.
5. Create a folder with read and write permissions, for temporary use during backup. Create the folder as follows:

**Windows systems:** `C:\tmp\backup\`

**Linux systems:** `/tmp/backup/`

- 
- ① **IMPORTANT:** It is the contents of this directory that will actually be backed up to the backup device by Data Protector. It must therefore have enough space available to cope with your largest likely backup. For instance, if you want to be able to back up all the virtual machines on your XenServer in one operation, there must be enough space for all of them. The folder will be referred to as the **temporary backup folder** in these instructions.
- 

## Integration script functions

The integration components, supplied in the `xenServer_backup_solution.zip` file, and their purposes are described below:

- **DPxen\_config.py.** This is the default configuration file for backup and restore operations using the integration. It is the only file that should be changed by the user, to match particular work scenarios. For details of setting up its contents, see “[Updating configuration script DPxen\\_config.py for backup](#)” (page 192) and “[Updating configuration script DPxen\\_config.py for restore](#)” (page 195).

For convenience, you can create your own versions of this file, with your own names, for instance `myBackupConfig` and `myRestoreConfig`, and specify these when calling the main scripts `DPxen_backup.py` and `DPxen_restore.py`.

- **DPxen\_operations.py.** This contains same auxiliary functions that are used during backup and restore, for example, functions to import or export a virtual machine, to find a virtual machine by name, and so on.
- **DPxen\_backup.py.** This controls obtaining the files to back up from the XenServer, according to the parameters specified in the configuration file, and writing them to the temporary backup folder ready for backup by Data Protector. If snapshots are created during the backup process it also deletes them afterwards.

This script can be called with the name of the configuration file as an argument, which allows you to use your own configuration file. If called without an argument, the default configuration file, `DPxen_config.py` is used.

- **DPxen\_restore.py.** This restores the files saved during the backup in the XenServer. It also deletes all the auxiliary elements created during a restore in the XenServer. The files restored by Data Protector to the temporary backup folder are not deleted by this script. This script also accepts as an argument the name of the configuration file to use.
- **DPxen\_postbackup.py.** This can be used to delete the contents of the temporary backup folder. It should be run after completion of a backup session as a cleanup operation. Alternatively, it could be used to delete the contents of the temporary backup folder before a restore session.

## Backup using the integration

To produce a backup using the integration, you need to perform the following operations:

1. Update the configuration script `DPxen_config.py`, or your own equivalent.
2. Ensure that the temporary backup directory is empty.
3. Set up a Data Protector backup specification.
4. Run the backup specification (or schedule it in the specification).

These operations are described in this section.

## Updating configuration script `DPxen_config.py` for backup

In the configuration file `DPxen_config.py`, or a configuration file with a name of your own, you can specify the following parameters required to perform the backup:

### *Mode*

The type of backup to be used and whether all, or only specified virtual machines are to be backed up. Valid options are: `online`, `allOnline`, `offline`, and `allOffline`. For explanations of these options, see [“Integration concepts” \(page 186\)](#).

### *url*

The URL of the XenServer containing the virtual machines to be backed up.

### *username*

The username of a user with the required access rights to the XenServer.

### *password*

The password of the user specified in *username*.

### *vmName*

A list containing the names of the virtual machines to be backed up. If you specify `online` or `offline` for the *Mode* parameter, you must specify the names of the individual virtual machines that you want to back up. The format of the list is:

```
["name1", "name2", ...]
```

The names can be obtained from the command line interface on the XenServer (running the `xe vm-list` command) or from the XenCenter interface, as shown in [Figure 102 \(page 193\)](#).

If you specify `allOnline` or `allOffline` for the *Mode* parameter, this parameter will be ignored.

### *debug*

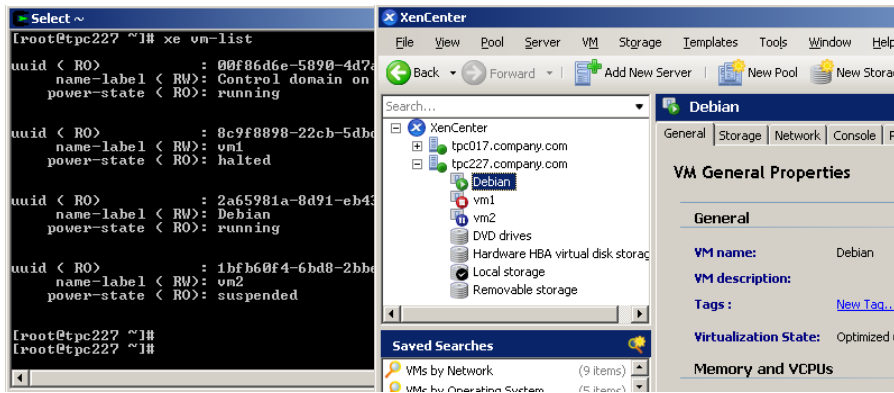
A Boolean value that should be set to `True` if you want to create a debug file during the backup.

### *sr*

This parameter is used for restore only. It is ignored during backup operations and can be left out for backup, if preferred.



Figure 102 Obtaining the names of the virtual machines



Once you have all the information for those parameters you can proceed to change the `DPxen_config.py` file or create a new configuration file.

### Example backup configurations

- Suspend and backup the Debian virtual machine, without creating a debug file.

```
Mode = "offline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian"]
debug = False
sr = ""
```
- Suspend and backup all the virtual machines present in the XenServer and create a debug file. The debug file will be created in `/tmp/` or `C:/tmp/`.

```
Mode = "allOffline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = []
debug = True
sr = ""
```
- Backup the virtual machines Debian and vm2 in hot mode, creating a debug file.

```
Mode = "online"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian", "vm2"]
debug = True
sr = ""
```
- Backup all the virtual machines present in the XenServer, without changing their power state. When the `Mode` selected is `allOnline` or `allOffline`, the value of the `vmName` parameter is ignored so its value is not important. In this example a debug file will not be created.

```
Mode = "allOnline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = []
debug = False
sr = ""
```

## Creating a backup specification

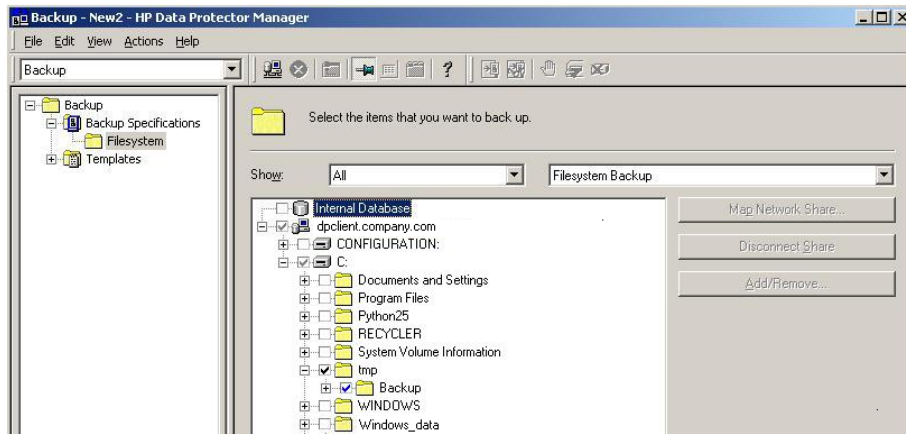
In the Data Protector GUI, create a filesystem backup specification:

1. Select the **Backup** context and expand **Backup Specifications**.
2. Right click on **Filesystem** and select **Add Backup**.
3. Select a **Blank Filesystem Backup** template and click **OK**.
4. Select the item to be backed up. In this case, you must select the temporary backup folder that was created during the installation process:

**Windows systems:** C:\tmp\backup

**Linux systems:** /tmp/backup

**Figure 103 Temporary backup folder selection**



5. Click **Next**.
6. Select the backup device on which the backup will be performed and click **Next**.
7. In Backup Specification Options, click **Advanced**.
8. In the Backup Options window, make the following entries in the Pre-exec and Post-exec sections:

**Pre-exec:** `PythonDirectory\python.exe ScriptPath\DPxen_backup.py`  
[ConfigFile]

**On client:** myDPClient.com

**Post-exec:** `PythonDirectory\python.exe ScriptPath\DPxen_postbackup.py`

**On client:** myDPClient.com

where:

- *PythonDirectory* is the name of the directory containing the python interpreter, for example, Python25. Python must be installed in the correct location, as described in “Prerequisites” (page 189).

---

**NOTE:** If the integration scripts have been installed on the Cell Manager, the *full* path must be given. For example, `Data_Protector_home\bin\Python25`.

---

- *ScriptPath* is the full path of the integration scripts, for example, `Data_Protector_home\bin`
- *ConfigFile* is an optional parameter for specifying a configuration file with a name of your own. If this is omitted, the default configuration file `DPxen_config.py` is used.

9. Click **OK** and then **Next**.

10. Schedule the backup in the normal way for a Data Protector backup. Select a backup type of **Full**. Even though incremental backups are shown as available, only full backups are supported for this integration.
11. Click **Next**
12. After reviewing the specification, click **Next**. Note that if you click **Properties** on the review panel, the options tab contains object options, not backup options, so your pre-exec and post-exec entries are not displayed.
13. Click **Save as...**, enter a name for the specification and click **OK**

## Restore using the integration

To restore a backup using the integration, you need to perform the following operations:

1. Update the configuration file `DPxen_config.py`, or your own equivalent.
2. Ensure that the temporary backup directory is empty.
3. Specify a Data Protector restore.
4. Perform the restore.

These operations are described in this section.

## Updating configuration script `DPxen_config.py` for restore

In a similar way to with backup, the configuration file `DPxen_config.py`, or a configuration file with a name of your own, must be used to specify the parameters required to perform a restore:

*Mode*

: The type of restore to be performed and whether all, or only specified virtual machines are restored. Valid options are: `online`, `allOnline`, `offline`, and `allOffline`. For explanations of these options, see ["Integration concepts" \(page 186\)](#).

*url*

The URL of the XenServer to which virtual machines are to be restored.

*username*

The username of a user with the required access rights to the XenServer.

*password*

The password of the user specified in *Username*.

*vmName*

A list containing the names of the virtual machines to be restored. If you specify `online` or `offline` for the *Mode* parameter, you must specify the names of the individual virtual machines that you want to restore. The format of the list is:

```
["name1", "name2", ..]
```

If you specify `allOnline` or `allOffline` for the *Mode* parameter, the *vmName* parameter is ignored.

*debug*

A Boolean value that should be set to `True` if you want to create a debug file during the restore.

*sr*

The name of the storage repository to which you want to restore the virtual machine(s) specified in the *vmName* parameter. This does not have to be the same storage repository as the one in use when the virtual machine was backed up. If the specified value is blank (`sr = ""`) the restore is performed to the default storage repository.

---

❗ **IMPORTANT:** You must have a default storage repository defined if you leave this parameter blank for a restore, otherwise the restore will fail.

---

## Example restore configurations

- Restore the Debian virtual machine, which was backed up using offline mode, into the default storage repository. Create a debug file.

```
Mode = "offline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian"]
debug = False
sr = ""
```

- Restore all the files backed up in the selected session into the default storage repository. The backup was performed in offline mode. Create a debug file.

```
Mode = "allOffline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = []
debug = True
sr = ""
```

- Restore the virtual machines Debian and vm2, which were backed up in online mode, into the default storage repository and create a debug file.

```
Mode = "online"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian", "vm2"]
debug = True
sr = ""
```

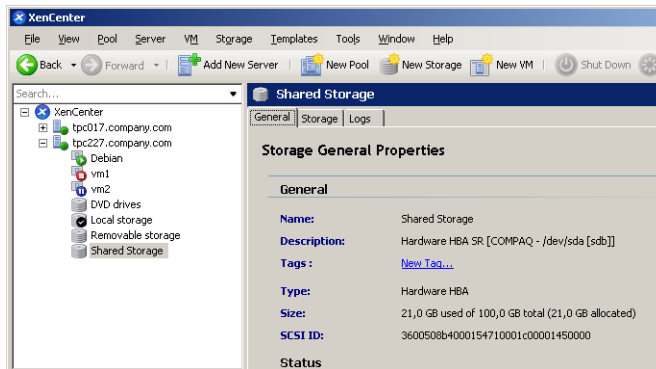
- Restore all the virtual machines present in the XenServer, without changing their power state. Create a debug file.

```
Mode = "allOnline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = []
debug = False
sr = ""
```

- Restore the Debian and vm2 virtual machines, which were backed up using the offline mode, into storage repository "Shared Storage" (see [Figure 104 \(page 197\)](#)). Create a debug file.

```
Mode = "offline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian", "vm2"]
debug = True
sr = "Shared Storage"
```

**Figure 104 Available storage repositories**



- Restore the Debian virtual machine, which was originally backed up using the `online` mode, into the "Removable storage" storage repository. Do not create a debug file.

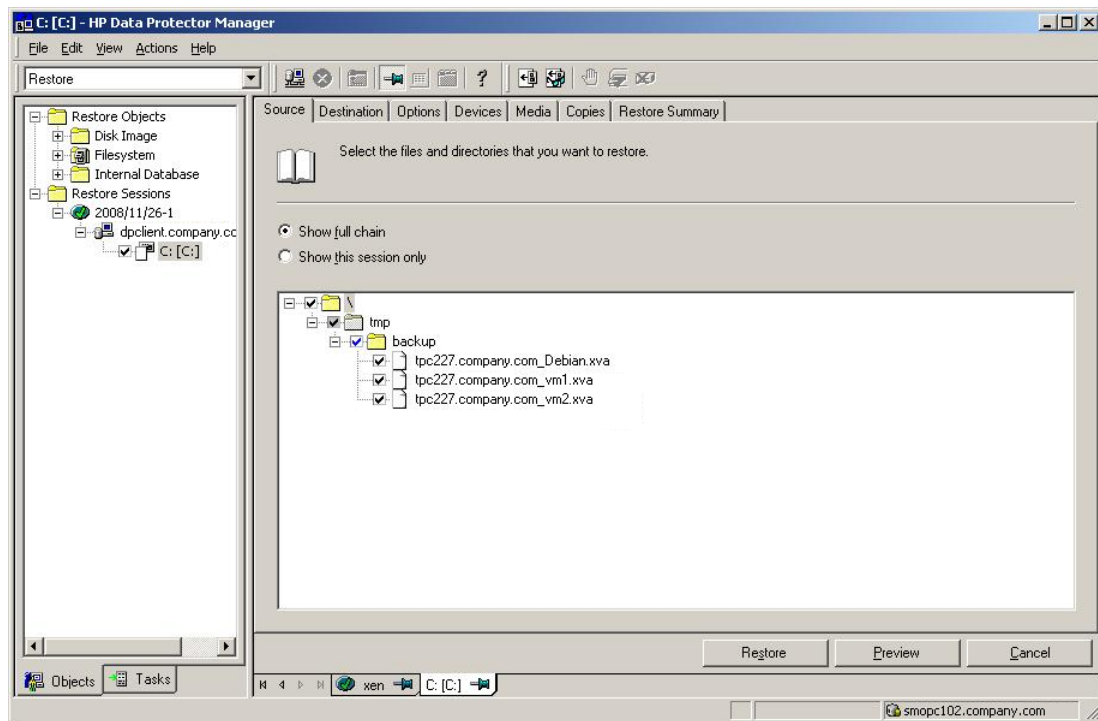
```
Mode = "online"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian"]
debug = False
sr = "Removable storage"
```

## Specifying a restore

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Sessions** to display clients and then the object backed up on those clients (The mount point containing the temporary backup directory you created during the installation process).
3. Click the object to open the property pages.
4. In the Source tab, select the temporary backup directory to view the files available for restore. Note that the filename formats indicate the mode of the backups in the folder, as follows:
  - `offline` or `allOffline`: `servername_vmName.xva`
  - `online` or `allOnline`: `servername_vmName_Snapshot.xva`

Check the file format, to decide the mode of restore required, (it must be equivalent to the type of backup). See [Figure 105 \(page 198\)](#).

**Figure 105 Selection of files to restore**



- ❗ **IMPORTANT:** The Source tab shows the contents of the temporary backup folder when the backup session was performed. It does not show the current contents of the folder. It is important to ensure that the temporary backup folder is empty before performing a restore.

During a restore, the files selected in this tab are first restored to the temporary backup folder. Only files selected here are subsequently available for the restore of virtual machines using the `DPxen_restore.py` script (provided the temporary backup folder was empty beforehand).

5. In the Options tab, in the **Post-exec** field, specify the `DPxen_restore.py` script as follows:

```
PythonPath\python.exe ScriptPath\DPxen_restore.py [ConfigFile]
```

where:

- *PythonPath* is the *full* path of the python interpreter, for example, `Data_Protector_home\bin\Python25`
- *ScriptPath* is the full path of the integration scripts, for example, `Data_Protector_home\bin`
- *ConfigFile* is an optional parameter for specifying a configuration file with a name of your own. If this is omitted, the default configuration file `DPxen_config.py` is used.

**NOTE:** The restore mode specified in the configuration file must be suitable for the filename format of the files in the folder. See the previous point.

6. Click **Restore** and check that the conditions for the restore, summarized in the **Start Restore Session** panel, are correct.
7. When you are sure that the restore is correctly specified, click **Finish** to start the restore.

## Notes on restore

- It is possible to perform the restore without specifying a post-exec. In this case, the .xva file is restored to the temporary backup folder only. You can then run the script independently.
- After the restore has completed, the files in the temporary backup folder are not deleted automatically. If you don't require them for any other purpose, it is advisable to delete them.
- When you restore a virtual machine, its UUID changes.

## Special considerations

- The online backup performs a snapshot of the virtual machine, but this operation is not always allowed. For XenServer 5.0, at the time of writing, snapshots were supported with VHD, Netapp, or EqualLogic storage repositories only. Any attempt to perform an online backup of a virtual machine that uses a different storage repository will fail and the error message, `At least one sr used by vmName does not support snapshot`, will be displayed in the monitor. If you are performing a backup of multiple virtual machines, only the virtual machines that allow snapshot operation will be backed up.

For the latest snapshot support with your XenServer version, see the XenServer documentation.

If you want to convert your default storage repository to file-based VHD-on-EXT3, follow the instructions given in the Citrix forum at <http://support.citrix.com/article/ctx116324>.

- When you perform an online backup, only the disks of the virtual machines are available for restore, not the metadata. To be able to restore the disks and the metadata, you need to perform an offline backup. The best practice is to perform an offline backup after the metadata of a virtual machine has changed, for example after installing a new application. The rest of the time it is sufficient to perform online backups.

If only the disks of a virtual machine have been corrupted you only need to restore your latest online backup. If the metadata is corrupted you need to first restore your latest offline backup and then your latest online backup, if newer.

- Before performing a backup, ensure that there is enough space on the Disk Agent host. If there is not enough space, the backup will fail.

## Further information

For further information, see <http://www.hp.com/go/dataprotector>.

---

# Glossary

## A

<b>access rights</b>	See user rights.
<b>ACSLs</b>	( <i>StorageTek specific term</i> ) The Automated Cartridge System Library Server (ACSLs) software that manages the Automated Cartridge System (ACS).
<b>Active Directory</b>	( <i>Windows specific term</i> ) The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.
<b>AES 256-bit encryption</b>	Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.
<b>AML</b>	( <i>ADIC/GRAU specific term</i> ) Automated Mixed-Media library.
<b>AMU</b>	( <i>ADIC/GRAU specific term</i> ) Archive Management Unit.
<b>application agent</b>	A component needed on a client to back up or restore online database integrations. See also Disk Agent.
<b>application system</b>	( <i>ZDB specific term</i> ) A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume.
<b>archive logging</b>	( <i>Lotus Domino Server specific term</i> ) Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.
<b>archived log files</b>	( <i>Data Protector specific term</i> ) Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online and offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.
<b>archived redo log</b>	( <i>Oracle specific term</i> ) Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: <ul style="list-style-type: none"><li>• ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A “hot” backup can be performed only when the database is running in this mode.</li><li>• NOARCHIVELOG - The filled online redo log files are not archived.</li></ul> See also online redo log.
<b>ASR set</b>	A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory <code>Data_Protector_program_data\Config\Server\dr\asr</code> (Windows systems) or <code>/etc/opt/omni/server/dr/asr</code> (UNIX systems) as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.
<b>audit logs</b>	Data files to which auditing information is stored.
<b>audit report</b>	User-readable output of auditing information created from data stored in audit log files.
<b>auditing information</b>	Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.
<b>autochanger</b>	See library.
<b>autoloader</b>	See library.



<b>Automatic Storage Management (ASM)</b>	<i>(Oracle specific term)</i> A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.
<b>auxiliary disk</b>	A bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.
<b>B</b>	
<b>BACKINT</b>	<i>(SAP R/3 specific term)</i> SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.
<b>backup API</b>	The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.
<b>backup chain</b>	See restore chain.
<b>backup device</b>	A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.
<b>backup generation</b>	One backup generation includes one full backup and all incremental backups until the next full backup.
<b>backup ID</b>	An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.
<b>backup object</b>	<p>A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image.</p> <p>A backup object is defined by:</p> <ul style="list-style-type: none"> <li>• Client name: Hostname of the Data Protector client where the backup object resides.</li> <li>• Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is located (drive on Windows systems and mount point on UNIX systems). For integration objects — backup stream identification, indicating the backed up database/application items.</li> <li>• Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration objects — displays the integration type (for example, SAP or Lotus).</li> <li>• Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".</li> </ul>
<b>backup owner</b>	Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.
<b>backup session</b>	A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type. The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification, full backup, and incremental backup.
<b>backup set</b>	A complete set of integration objects associated with a backup.
<b>backup set</b>	<i>(Oracle specific term)</i> A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.
<b>backup specification</b>	A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

<b>backup system</b>	<i>(ZDB specific term)</i> A system connected to a disk array together with one or multiple application systems. The backup system is typically connected to a disk array to create target volumes (a replica) and is used for mounting the target volumes (the replica). See also application system, target volume, and replica.
<b>backup types</b>	See incremental backup, differential backup, transaction backup, full backup, and delta backup.
<b>backup view</b>	Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.
<b>BC</b>	<i>(EMC Symmetrix specific term)</i> Business Continuity are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices. See also BCV.
<b>BC Process</b>	<i>(EMC Symmetrix specific term)</i> A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuity Volumes to protect data on EMC Symmetrix standard devices. See also BCV.
<b>BCV</b>	<i>(EMC Symmetrix specific term)</i> Business Continuity Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuity operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected. See also BC and BC Process.
<b>Boolean operators</b>	The Boolean operators for the full text search functionality of the Data Protector Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.
<b>boot volume/disk/partition</b>	A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.
<b>BRARCHIVE</b>	<i>(SAP R/3 specific term)</i> An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. See also BRBACKUP and BRRESTORE.
<b>BRBACKUP</b>	<i>(SAP R/3 specific term)</i> An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files. See also BRARCHIVE and BRRESTORE.
<b>BRRESTORE</b>	<i>(SAP R/3 specific term)</i> An SAP R/3 tool that can be used to restore files of the following type: <ul style="list-style-type: none"> <li>• Database data files, control files, and online redo log files saved with BRBACKUP</li> <li>• Redo log files archived with BRARCHIVE</li> <li>• Non-database files saved with BRBACKUP</li> </ul> You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup. See also BRBACKUP and BRARCHIVE.
<b>BSM</b>	The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.
<b>C</b>	
<b>CAP</b>	<i>(StorageTek specific term)</i> Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

<b>Catalog Database (CDB)</b>	A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.
<b>catalog protection</b>	Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.
<b>CDB</b>	See Catalog Database (CDB).
<b>CDF file</b>	(UNIX systems specific term) A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.
<b>cell</b>	A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN or SAN. Central control is available to administer the backup and restore policies and tasks.
<b>Cell Manager</b>	The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.
<b>centralized licensing</b>	Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs. See also MoM.
<b>Centralized Media Management Database (CMMDB)</b>	See CMMDB.
<b>Certificate Server</b>	A Windows Certificate Server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.
<b>Change Journal</b>	(Windows specific term) A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.
<b>Change Log Provider</b>	(Windows specific term) A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.
<b>channel</b>	<p>(Oracle specific term) An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:</p> <ul style="list-style-type: none"> <li>• type 'disk'</li> <li>• type 'sbt_tape'</li> </ul> <p>If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.</p>
<b>circular logging</b>	(Microsoft Exchange Server and Lotus Domino Server specific term) Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.
<b>client backup</b>	<p>A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification:</p> <ul style="list-style-type: none"> <li>• If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up.</li> <li>• If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the</li> </ul>

backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the backup specification was created are not backed up.

<b>client or client system</b>	Any system configured with any Data Protector functionality and configured in a cell.
<b>cluster continuous replication</b>	<p>(<i>Microsoft Exchange Server specific term</i>) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.</p> <p>A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.</p> <p>See also Exchange Replication Service and local continuous replication.</p>
<b>cluster-aware application</b>	It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).
<b>CMD script for Informix Server</b>	( <i>Informix Server specific term</i> ) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.
<b>CMMDB</b>	<p>The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended</p> <p>See also MoM.</p>
<b>COM+ Class Registration Database</b>	( <i>Windows specific term</i> ) The COM+ Class Registration Database and the Windows Registry store application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.
<b>command device</b>	( <i>HP P9000 XP Disk Array Family specific term</i> ) A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.
<b>command-line interface (CLI)</b>	A set of commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.
<b>concurrency</b>	See Disk Agent concurrency.
<b>container</b>	( <i>HP P6000 EVA Disk Array Family specific term</i> ) Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.
<b>control file</b>	( <i>Oracle and SAP R/3 specific term</i> ) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.
<b>copy set</b>	<p>(<i>HP P6000 EVA Disk Array Family specific term</i>) A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA.</p> <p>See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.</p>
<b>CRS</b>	The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account <code>root</code> .
<b>CSM</b>	The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

## D

<b>data file</b>	<i>(Oracle and SAP R/3 specific term)</i> A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.
<b>data protection</b>	Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions. <i>See also</i> catalog protection.
<b>data replication (DR) group</b>	<i>(HP P6000 EVA Disk Array Family specific term)</i> A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log. <i>See also</i> copy set.
<b>data stream</b>	Sequence of data transferred over the communication channel.
<b>Data_Protector_home</b>	A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is <code>%ProgramFiles%\OmniBack</code> , but the path can be changed in the Data Protector Setup Wizard at installation time. <i>See also</i> Data_Protector_program_data.
<b>Data_Protector_program_data</b>	A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012. Its default path is <code>%ProgramData%\OmniBack</code> , but the path can be changed in the Data Protector Setup Wizard at installation time. <i>See also</i> Data_Protector_home.
<b>database library</b>	A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.
<b>database parallelism</b>	More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.
<b>database server</b>	A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.
<b>Dbobject</b>	<i>(Informix Server specific term)</i> An Informix Server physical database object. It can be a blob space, db space, or logical log file.
<b>DC directory</b>	A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. <i>See also</i> Detail Catalog Binary Files (DCBF) and Internal Database (IDB).
<b>DCBF</b>	<i>See</i> Detail Catalog Binary Files (DCBF).
<b>delta backup</b>	A delta backup is a backup containing all the changes made to the database from the last backup of any type. <i>See also</i> backup types.
<b>Detail Catalog Binary Files (DCBF)</b>	A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. <i>See also</i> DC directory and Internal Database (IDB).
<b>device</b>	A physical unit which contains either just a drive or a more complex unit such as a library.
<b>device chain</b>	A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.
<b>device group</b>	<i>(EMC Symmetrix specific term)</i> A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.
<b>device streaming</b>	A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written

	to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.
<b>DHCP server</b>	A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.
<b>differential backup</b>	An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. <i>See also</i> incremental backup.
<b>differential backup</b>	<i>(Microsoft SQL Server specific term)</i> A database backup that records only the data changes made to the database after the last full database backup. <i>See also</i> backup types.
<b>differential database backup</b>	A differential database backup records only those data changes made to the database after the last full database backup.
<b>directory junction</b>	<i>(Windows specific term)</i> Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.
<b>disaster recovery</b>	A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.
<b>disaster recovery operating system</b>	<i>See</i> DR OS.
<b>Disk Agent</b>	A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.
<b>Disk Agent concurrency</b>	The number of Disk Agents that are allowed to send data to one Media Agent concurrently.
<b>disk group</b>	<i>(Veritas Volume Manager specific term)</i> The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.
<b>disk image backup</b>	A high-speed backup where Data Protector backs up files as bitmap images. A disk image backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.
<b>disk quota</b>	A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.
<b>disk staging</b>	The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).
<b>distributed file media format</b>	A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. <i>See also</i> virtual full backup.
<b>Distributed File System (DFS)</b>	A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.
<b>DMZ</b>	The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.
<b>DNS server</b>	In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

<b>domain controller</b>	A server in a network that is responsible for user security and verifying passwords within a group of other servers.
<b>DR image</b>	Data required for temporary disaster recovery operating system (DR OS) installation and configuration.
<b>DR OS</b>	An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.
<b>drive</b>	A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.
<b>drive index</b>	A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.
<b>drive-based encryption</b>	Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the metadata that is written to the medium.
<b>E</b>	
<b>EMC Symmetrix Agent</b>	A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.
<b>emergency boot file</b>	(Informix Server specific term) The Informix Server configuration file <code>ixbar.server_id</code> that resides in the directory <code>INFORMIXDIR/etc</code> (on Windows systems) or <code>INFORMIXDIR\etc</code> (on UNIX systems). <code>INFORMIXDIR</code> is the Informix Server home directory and <code>server_id</code> is the value of the <code>SERVENUM</code> configuration parameter. Each line of the emergency boot file corresponds to one backup object.
<b>encrypted control communication</b>	Data Protector secure communication between the clients in the Data Protector cell is based on Secure Socket Layer (SSL) that uses SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.
<b>encryption key</b>	A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.
<b>encryption KeyID-StoreID</b>	Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. <code>KeyID</code> identifies the key within the keystore. <code>StoreID</code> identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may several <code>StoreIDs</code> used on the same Cell Manager.
<b>enhanced incremental backup</b>	Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.
<b>enterprise backup environment</b>	Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept. <i>See also MoM.</i>
<b>Event Log (Data Protector Event Log)</b>	A central repository of all Data Protector-related notifications. By default, all notifications are sent to the Event Log. The events are logged on the Cell Manager into the file <code>Data_Protector_program_data\log\server\Ob2EventLog.txt</code> (Windows systems),



or `/var/opt/omni/server/log/Ob2EventLog.txt` (UNIX systems). The Event Log is accessible only to users of the Data Protector Admin user group and to users who are granted the Data Protector Reporting and notifications user rights. You can view or delete all events in the Event Log.

<b>Event Logs</b>	( <i>Windows specific term</i> ) Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.
<b>Exchange Replication Service</b>	( <i>Microsoft Exchange Server specific term</i> ) The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.
<b>exchanger</b>	Also referred to as SCSI Exchanger. See also library.
<b>exporting media</b>	A process that removes all data about backup sessions, such as systems, objects, and filenames, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. See also importing media.
<b>Extensible Storage Engine (ESE)</b>	( <i>Microsoft Exchange Server specific term</i> ) A database technology used as a storage system for information exchange in Microsoft Exchange Server.

## F

<b>failover</b>	Transferring of the most important cluster data, called group (on Windows systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.
<b>failover</b>	( <i>HP P6000 EVA Disk Array Family specific term</i> ) An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations. See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.
<b>FC bridge</b>	See Fibre Channel bridge.
<b>Fibre Channel</b>	An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.
<b>Fibre Channel bridge</b>	A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.
<b>file depot</b>	A file containing the data from a backup to a file library device.
<b>file jukebox device</b>	A device residing on disk consisting of multiple slots used to store file media.
<b>file library device</b>	A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.
<b>File Replication Service (FRS)</b>	A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.
<b>file tree walk</b>	( <i>Windows specific term</i> ) The process of traversing a filesystem to determine which objects have been created, modified, or deleted.
<b>file version</b>	The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.
<b>filesystem</b>	The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.
<b>first-level mirror</b>	( <i>HP P9000 XP Disk Array Family specific term</i> ) A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level



mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used.

See also primary volume and mirror unit (MU) number.

<b>flash recovery area</b>	<p>(Oracle specific term) A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files).</p> <p>See also recovery files.</p>
<b>formatting</b>	<p>A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.</p>
<b>free pool</b>	<p>An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.</p>
<b>full backup</b>	<p>A backup in which all selected objects are backed up, whether or not they have been recently modified.</p> <p>See also backup types.</p>
<b>full database backup</b>	<p>A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.</p>
<b>full mailbox backup</b>	<p>A full mailbox backup is a backup of the entire mailbox content.</p>
<b>full ZDB</b>	<p>A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.</p> <p>See also incremental ZDB.</p>

## G

<b>global options</b>	<p>A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager</p>
<b>group</b>	<p>(Microsoft Cluster Server specific term) A collection of resources (for example disk volumes, application services, IP names, and addresses) that are needed to run a specific cluster-aware applications.</p>
<b>GUI</b>	<p>A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.</p>

## H

<b>hard recovery</b>	<p>(Microsoft Exchange Server specific term) A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.</p>
<b>heartbeat</b>	<p>A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.</p>
<b>Hierarchical Storage Management (HSM)</b>	<p>A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.</p>
<b>Holidays file</b>	<p>A file that contains information about holidays. You can set different holidays by editing the Holidays file on the Cell Manager in the directory <code>Data_Protector_program_data\Config\Server\holidays</code> (Windows systems), or <code>/etc/opt/omni/server/Holidays</code> (UNIX systems).</p>
<b>hosting system</b>	<p>A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.</p>
<b>HP Business Copy (BC) P6000 EVA</b>	<p>(HP P6000 EVA Disk Array Family specific term) A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware.</p>

See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

**HP Business Copy (BC) P9000 XP**

(*HP P9000 XP Disk Array Family specific term*) An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system.

See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit, application system, and backup system.

**HP Command View (CV) EVA**

(*HP P6000 EVA Disk Array Family specific term*) The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, and mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.

See also HP P6000 / HP 3PAR SMI-S Agent and HP SMI-S P6000 EVA Array provider.

**HP Continuous Access (CA) P9000 XP**

(*HP P9000 XP Disk Array Family specific term*) An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs).

See also HP Business Copy (BC) P9000 XP, Main Control Unit, and LDEV.

**HP Continuous Access + Business Copy (CA+BC) P6000 EVA**

(*HP P6000 EVA Disk Array Family specific term*) An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array.

See also HP Business Copy (BC) P6000 EVA, replica, and source volume.

**HP P6000 / HP 3PAR SMI-S Agent**

A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the HP P6000 / HP 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface.

See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

**HP P9000 XP Agent**

A Data Protector component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It uses RAID Manager Library for communication with a P9000 XP Array storage system.

See also RAID Manager Library.

**HP SMI-S P6000 EVA Array provider**

An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses. See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

**ICDA**

(*EMC Symmetrix specific term*) EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

**IDB**

See Internal Database (IDB).

**IDB recovery file**

A file that maintains information about completed IDB backup sessions and the backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

<b>importing media</b>	A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media. <i>See also</i> exporting media.
<b>incremental (re)-establish</b>	<i>(EMC Symmetrix specific term)</i> A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
<b>incremental backup</b>	A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length. <i>See also</i> backup types.
<b>incremental backup</b>	<i>(Microsoft Exchange Server specific term)</i> A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up. <i>See also</i> backup types.
<b>incremental mailbox backup</b>	An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.
<b>incremental restore</b>	<i>(EMC Symmetrix specific term)</i> A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.
<b>incremental ZDB</b>	A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape. <i>See also</i> full ZDB.
<b>incremental1 mailbox backup</b>	An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.
<b>Inet</b>	A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.
<b>Information Store</b>	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. <i>See also</i> Key Management Service and Site Replication Service.
<b>Informix Server initializing</b>	<i>(Informix Server specific term)</i> Refers to Informix Dynamic Server. <i>See</i> formatting.
<b>Installation Server</b>	A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.
<b>instant recovery</b>	<i>(ZDB specific term)</i> A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery. <i>See also</i> replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

<b>integration object</b>	A backup object of a Data Protector integration, such as Oracle or SAP DB.
<b>Internal Database (IDB)</b>	An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It stores its data in an embedded database and a collection of proprietary data files which reside on the Cell Manager. <i>See also</i> DC directory and Detail Catalog Binary Files (DBCF).
<b>Internet Information Services (IIS)</b>	<i>(Windows specific term)</i> Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).
<b>ISQL</b>	<i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.

## J

<b>jukebox</b>	<i>See</i> library.
<b>jukebox device</b>	A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the "file jukebox device".

## K

<b>Key Management Service</b>	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that provides encryption functionality for enhanced security. <i>See also</i> Information Store and Site Replication Service.
<b>keychain</b>	A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.
<b>keystore</b>	All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).
<b>KMS</b>	Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

## L

<b>LBO</b>	<i>(EMC Symmetrix specific term)</i> A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.
<b>LDEV</b>	<i>(HP P9000 XP Disk Array Family specific term)</i> A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. <i>See also</i> HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.
<b>library</b>	Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.
<b>lights-out operation or unattended operation</b>	A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.
<b>LISTENER.ORA</b>	<i>(Oracle specific term)</i> An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.
<b>load balancing</b>	By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used

	for each object in the backup specification. Data Protector will access the devices in the specified order.
<b>local and remote recovery</b>	Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.
<b>local continuous replication</b>	<p>(<i>Microsoft Exchange Server specific term</i>) Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.</p> <p>An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal.</p> <p>A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.</p> <p>See also cluster continuous replication and Exchange Replication Service.</p>
<b>lock name</b>	You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.
<b>log_full shell script</b>	<p>(<i>Informix Server UNIX systems specific term</i>) A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server <code>ALARMPROGRAM</code> configuration parameter defaults to the <code>INFORMIXDIR/etc/log_full.sh</code>, where <code>INFORMIXDIR</code> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the <code>ALARMPROGRAM</code> configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code>.</p>
<b>logging level</b>	An option that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth and the convenience of browsing data for restore.
<b>logical-log files</b>	This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.
<b>login ID</b>	( <i>Microsoft SQL Server specific term</i> ) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table <code>syslogin</code> .
<b>login information to the Oracle Target Database</b>	<p>(<i>Oracle and SAP R/3 specific term</i>) The format of the login information is <code>user_name/password@service</code>, where:</p> <ul style="list-style-type: none"> <li><code>user_name</code> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle <code>SYSDBA</code> or <code>SYSOPER</code> rights.</li> <li><code>password</code> must be the same as the password specified in the Oracle password file (<code>orapwd</code>), which is used for authentication of users performing database administration.</li> <li><code>service</code> is the name used to identify an SQL*Net server process for the target database.</li> </ul>
<b>login information to the Recovery Catalog Database</b>	( <i>Oracle specific term</i> ) The format of the login information to the Recovery (Oracle) Catalog Database is <code>user_name/password@service</code> , where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the

Oracle target database. In this case, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

## **Lotus C API**

(*Lotus Domino Server specific term*) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

## **LVM**

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

## **M**

### **Magic Packet**

See Wake ONLAN.

### **mailbox**

(*Microsoft Exchange Server specific term*) The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

### **mailbox store**

(*Microsoft Exchange Server specific term*) A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

### **Main Control Unit (MCU)**

(*HP P9000 XP Disk Array Family specific term*) An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP or HP CA+BC P9000 XP configuration and acts as a master device.

See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

### **maintenance mode**

An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the Data Protector installation.

### **make\_net\_recovery**

`make_net_recovery` is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX `bootsys` command or interactively specified on the boot console.

### **make\_tape\_recovery**

`make_tape_recovery` is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

### **Manager-of-Managers (MoM)**

See MoM.

### **MAPI**

(*Microsoft Exchange Server specific term*) The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

### **MCU**

See Main Control Unit (MCU).

### **Media Agent**

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

### **media allocation policy**

Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

<b>media condition</b>	The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.
<b>media condition factors</b>	The user-assigned age threshold and overwrite threshold used to determine the state of a medium.
<b>media label</b>	A user-defined identifier used to describe a medium.
<b>media location</b>	A user-defined physical location of a medium, such as "building 4" or "off-site storage".
<b>media management session</b>	A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.
<b>media pool</b>	A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.
<b>media set</b>	The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.
<b>media type</b>	The physical type of media, such as DDS or DLT.
<b>media usage policy</b>	The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.
<b>medium ID</b>	A unique identifier assigned to a medium by Data Protector.
<b>merging</b>	This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. <i>See also</i> overwrite.
<b>Microsoft Exchange Server</b>	A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.
<b>Microsoft Management Console (MMC)</b>	<i>(Windows specific term)</i> An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.
<b>Microsoft SQL Server</b>	A database management system designed to meet the requirements of distributed "client-server" computing.
<b>Microsoft Volume Shadow Copy Service (VSS)</b>	A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets. <i>See also</i> shadow copy, shadow copy provider, replica, and writer.
<b>mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</b>	<i>See</i> target volume.
<b>mirror rotation (HP P9000 XP Disk Array Family specific term)</b>	<i>See</i> replica set rotation.
<b>mirror unit (MU) number</b>	<i>(HP P9000 XP Disk Array Family specific term)</i> A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family. <i>See also</i> first-level mirror.
<b>mirrorclone</b>	<i>(HP P6000 EVA Disk Array Family specific term)</i> A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.



<b>MMD</b>	The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.
<b>MMDB</b>	The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells. See also CMMDB and Catalog Database (CDB).
<b>MoM</b>	Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.
<b>mount point</b>	The access point in a directory structure for a disk or logical volume, for example /opt or d: . On UNIX systems, the mount points are displayed using the bdf or df command.
<b>mount request</b>	A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.
<b>MSM</b>	The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.
<b>multisnapping</b>	(HP P6000 EVA Disk Array Family specific term) Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot. See also snapshot.
○	
<b>OBDR capable device</b>	A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.
<b>obdrindex.dat</b>	See IDB recovery file.
<b>object</b>	See backup object.
<b>object consolidation</b>	The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.
<b>object consolidation session</b>	A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.
<b>object copy</b>	A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.
<b>object copy session</b>	A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.
<b>object copying</b>	The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.
<b>object ID</b>	(Windows specific term) The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.
<b>object mirror</b>	A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.
<b>object mirroring</b>	The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.
<b>object verification</b>	The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.



<b>object verification session</b>	A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.
<b>offline backup</b>	A backup during which an application database cannot be used by the application. In an offline backup session, the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the time period of the data replication process. For instance, for backup to tape, until streaming of data to the tape is finished. Normal database operation is resumed before potential post-backup operations are started. See also zero downtime backup (ZDB) and online backup.
<b>offline recovery</b>	Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Cell Manager can only be recovered offline.
<b>offline redo log</b>	See archived redo log.
<b>ON-Bar</b>	<i>(Informix Server specific term)</i> A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: <ul style="list-style-type: none"> <li>• the onbar command</li> <li>• Data Protector as the backup solution</li> <li>• the XBSA interface</li> <li>• ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.</li> </ul>
<b>ONCONFIG</b>	<i>(Informix Server specific term)</i> An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file in the directory <code>INFORMIXDIR/etc</code> (on Windows systems or <code>INFORMIXDIR/etc/</code> (on UNIX systems).
<b>online backup</b>	A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period of the data replication process. For instance, for backup to tape, until streaming of data to tape is finished. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. Normal database operation is resumed before potential post-backup operations are started.  In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. See also zero downtime backup (ZDB) and offline backup.
<b>online recovery</b>	A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.
<b>online redo log</b>	<i>(Oracle specific term)</i> Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused. See also archived redo log.
<b>Oracle Data Guard</b>	<i>(Oracle specific term)</i> Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.
<b>Oracle instance</b>	<i>(Oracle specific term)</i> Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.
<b>ORACLE_SID</b>	<i>(Oracle specific term)</i> A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <code>ORACLE_SID</code> . The <code>ORACLE_SID</code> is included in the <code>CONNECT DATA</code>

parts of the connect descriptor in a `TNSNAMES.ORA` file and in the definition of the TNS listener in the `LISTENER.ORA` file.

**original system**

The system configuration backed up by Data Protector before a computer disaster hits the system.

**overwrite**

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.

See also merging.

**ownership**

Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.

If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is `root:sys` unless the above conditions are true.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

**P**

**P1S file**

P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into the directory

`Data_Protector_program_data\Config\Server\dr\p1s` (Windows systems), or  
`/etc/opt/omni/server/dr/p1s` (UNIX systems) with the filename `recovery.p1s`.

**package**

(MC/ServiceGuard and Veritas Cluster specific term) A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

**pair status**

(HP P9000 XP Disk Array Family specific term) The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family.

Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP P9000 XP Agent:

- PAIR – The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.
- SUSPENDED – The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.
- COPY – The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

**parallel restore**

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same objects using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

**parallelism**

The concept of reading multiple data streams from an online database.

<b>phase 0 of disaster recovery</b>	Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.
<b>phase 1 of disaster recovery</b>	Installation and configuration of DR OS, establishing previous storage structure.
<b>phase 2 of disaster recovery</b>	Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.
<b>phase 3 of disaster recovery</b>	Restoration of user and application data.
<b>physical device</b>	A physical unit that contains either a drive or a more complex unit such as a library.
<b>post-exec</b>	A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also pre-exec.
<b>pre- and post-exec commands</b>	Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems.
<b>pre-exec</b>	A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also post-exec.
<b>prealloc list</b>	A subset of media in a media pool that specifies the order in which media are used for backup.
<b>primary volume (P-VOL)</b>	(HP P9000 XP Disk Array Family specific term) An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU). See also secondary volume (S-VOL) and Main Control Unit (MCU).
<b>protection</b>	See data protection and also catalog protection.
<b>public folder store</b>	(Microsoft Exchange Server specific term) The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.
<b>public/private backed up data</b>	When configuring a backup, you can select whether the backed up data will be: <ul style="list-style-type: none"> <li>• public, that is visible (and accessible for restore) to all Data Protector users</li> <li>• private, that is, visible (and accessible for restore) only to the owner of the backup and administrators</li> </ul>

## R

<b>RAID</b>	Redundant Array of Independent Disks.
<b>RAID Manager Library</b>	(HP P9000 XP Disk Array Family specific term) A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands. See also HP P9000 XP Agent.
<b>RAID Manager P9000 XP</b>	(HP P9000 XP Disk Array Family specific term) A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.
<b>rawdisk backup</b>	See disk image backup.
<b>RCU</b>	See Remote Control Unit (RCU).
<b>RDBMS</b>	Relational Database Management System.

<b>RDF1/RDF2</b>	<i>(EMC Symmetrix specific term)</i> A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.
<b>Recovery Catalog</b>	<p><i>(Oracle specific term)</i> A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:</p> <ul style="list-style-type: none"> <li>• The physical schema of the Oracle target database</li> <li>• Data file and archived log backup sets</li> <li>• Data file copies</li> <li>• Archived redo logs</li> <li>• Stored scripts</li> </ul>
<b>Recovery Catalog Database</b>	<i>(Oracle specific term)</i> An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.
<b>recovery files</b>	<p><i>(Oracle specific term)</i> Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces.</p> <p>See also flash recovery area.</p>
<b>Recovery Manager (RMAN)</b>	<i>(Oracle specific term)</i> An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.
<b>RecoveryInfo</b>	When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.
<b>recycle or unprotect</b>	A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.
<b>redo log</b>	<i>(Oracle specific term)</i> Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.
<b>Remote Control Unit (RCU)</b>	<i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.
<b>Removable Storage Management Database</b>	<i>(Windows specific term)</i> A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.
<b>reparse point</b>	<i>(Windows specific term)</i> A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.
<b>replica</b>	<i>(ZDB specific term)</i> An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on a UNIX system, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition is replicated.

	See also snapshot, snapshot creation, split mirror, and split mirror creation.
<b>replica set</b>	(ZDB specific term) A group of replicas, all created using the same backup specification. See also replica and replica set rotation.
<b>replica set rotation</b>	(ZDB specific term) The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set. See also replica and replica set.
<b>restore chain</b>	Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.
<b>restore session</b>	A process that copies data from backup media to a client.
<b>resync mode</b>	(HP P9000 XP Disk Array Family VSS provider specific term) One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL. See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), mirror unit (MU) number, and replica set rotation.
<b>RMAN (Oracle specific term)</b>	See Recovery Manager.
<b>RSM</b>	The Data Protector Restore Session Manager controls restore and object verification sessions. This process always runs on the Cell Manager system.
<b>RSM</b>	(Windows specific term) Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.
<b>S</b>	
<b>SAPDBA</b>	(SAP R/3 specific term) An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.
<b>scanning</b>	A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.
<b>Scheduler</b>	A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.
<b>secondary volume (S-VOL)</b>	(HP P9000 XP Disk Array Family specific term) An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration. See also primary volume (P-VOL) and Main Control Unit (MCU).
<b>session</b>	See backup session, media management session, and restore session.
<b>session ID</b>	An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.
<b>session key</b>	This environment variable for the pre-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat, and omniabort commands.
<b>shadow copy</b>	(Microsoft VSS specific term) A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original

	<p>volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.</p> <p>See also Microsoft Volume Shadow Copy Service and replica.</p>
<b>shadow copy provider</b>	<p>(Microsoft VSS specific term) An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays). See also shadow copy.</p>
<b>shadow copy set</b>	<p>(Microsoft VSS specific term) A collection of shadow copies created at the same point in time. See also shadow copy and replica set.</p>
<b>shared disks</b>	<p>A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.</p>
<b>Site Replication Service</b>	<p>(Microsoft Exchange Server specific term) The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.</p> <p>See also Information Store and Key Management Service.</p>
<b>slot</b>	<p>A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.</p>
<b>SMB</b>	<p>See split mirror backup.</p>
<b>SMBF</b>	<p>The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.</p>
<b>SMI-S Agent (SMISA)</b>	<p>See HP P6000 / HP 3PAR SMI-S Agent.</p>
<b>snapshot</b>	<p>(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term) A type of target volumes created using a specific replication technology. Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume.</p> <p>See also replica and snapshot creation.</p>
<b>snapshot backup</b>	<p>See ZDB to tape, ZDB to disk, and ZDB to disk+tape.</p>
<b>snapshot creation</b>	<p>(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term) A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation.</p> <p>See also snapshot.</p>
<b>source (R1) device</b>	<p>(EMC Symmetrix specific term) An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.</p> <p>See also target (R2) device.</p>
<b>source volume</b>	<p>(ZDB specific term) A storage volume containing data to be replicated.</p>
<b>sparse file</b>	<p>A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.</p>
<b>split mirror</b>	<p>(EMC Symmetrix Disk Array and HP P9000 XP Disk Array Family specific term) A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes.</p> <p>See also replica and split mirror creation.</p>

<b>split mirror backup</b> (EMC Symmetrix specific term)	See ZDB to tape.
<b>split mirror backup</b> (HP P9000 XP Disk Array Family specific term)	See ZDB to tape, ZDB to disk, and ZDB to disk+tape.
<b>split mirror creation</b>	(EMC Symmetrix and HP P9000 XP Disk Array Family specific term) A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes. See also split mirror.
<b>split mirror restore</b>	(EMC Symmetrix and HP P9000 XP Disk Array Family specific term) A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method. See also ZDB to tape, ZDB to disk+tape, and replica.
<b>sqlhosts file or registry</b>	(Informix Server specific term) An Informix Server connectivity information file (on UNIX systems) or registry (on Windows systems) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.
<b>SRD file</b>	(disaster recovery specific term) A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows or Linux system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster. See also target system.
<b>SRDF</b>	(EMC Symmetrix specific term) The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.
<b>SSE Agent (SSEA)</b>	See HP P9000 XP Agent.
<b>sst.conf file</b>	The file <code>/usr/kernel/drv/sst.conf</code> is required on each Data Protector Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.
<b>st.conf file</b>	The file <code>/kernel/drv/st.conf</code> is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.
<b>stackers</b>	Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.
<b>standalone file device</b>	A file device is a file in a specified directory to which you back up data.
<b>Storage Group</b>	(Microsoft Exchange Server specific term) A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.
<b>storage volume</b>	(ZDB specific term) An object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, filesystems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.
<b>StorageTek ACS library</b>	(StorageTek specific term) Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.
<b>switchover</b>	See failover.



<b>Sybase Backup Server API</b>	( <i>Sybase specific term</i> ) An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.
<b>Sybase SQL Server</b>	( <i>Sybase specific term</i> ) The server in the Sybase “client-server” architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.
<b>SYMA</b>	See EMC Symmetrix Agent.
<b>synthetic backup</b>	A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.
<b>synthetic full backup</b>	The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.
<b>System Backup to Tape</b>	( <i>Oracle specific term</i> ) An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.
<b>system databases</b>	( <i>Sybase specific term</i> ) The four system databases on a newly installed Sybase SQL Server are the: <ul style="list-style-type: none"> <li>• master database (master)</li> <li>• temporary database (tempdb)</li> <li>• system procedure database (sybsystemprocs)</li> <li>• model database (model).</li> </ul>
<b>System Recovery Data file</b>	See SRD file.
<b>System State</b>	( <i>Windows specific term</i> ) The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.
<b>system volume/disk/partition</b>	A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.
<b>SysVol</b>	( <i>Windows specific term</i> ) A shared directory that stores the server copy of the domain’s public files, which are replicated among all domain controllers in the domain.
<b>T</b>	
<b>tablespace</b>	A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.
<b>tapeless backup (ZDB specific term)</b>	See ZDB to disk.
<b>target (R2) device</b>	( <i>EMC Symmetrix specific term</i> ) An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. See also source (R1) device.
<b>target database</b>	( <i>Oracle specific term</i> ) In RMAN, the target database is the database that you are backing up or restoring.
<b>target system</b>	( <i>disaster recovery specific term</i> ) A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.



<b>target volume</b>	(ZDB specific term) A storage volume to which data is replicated.
<b>Terminal Services</b>	(Windows specific term) Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.
<b>thread</b>	(Microsoft SQL Server specific term) An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.
<b>TimeFinder</b>	(EMC Symmetrix specific term) A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).
<b>TLU</b>	Tape Library Unit.
<b>TNSNAMES.ORA</b>	(Oracle and SAP R/3 specific term) A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.
<b>transaction</b>	A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.
<b>transaction backup</b>	Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.
<b>transaction backup</b>	(Sybase and SQL specific term) A backup of the transaction log providing a record of changes made since the last full or transaction backup.
<b>transaction log backup</b>	Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.
<b>transaction log files</b>	Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.
<b>transaction log table</b>	(Sybase specific term) A system table in which all changes to the database are automatically recorded.
<b>transportable snapshot</b>	(Microsoft VSS specific term) A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed. See also Microsoft Volume Shadow Copy Service (VSS).

## U

<b>unattended operation</b>	See lights-out operation.
<b>user account (Data Protector user account)</b>	You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.
<b>User Account Control (UAC)</b>	A security component in Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.
<b>user disk quotas</b>	NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.
<b>user group</b>	Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.
<b>user profile</b>	(Windows specific term) Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

<b>user rights</b>	User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.
<b>user_restrictions file</b>	A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than <i>admin</i> and <i>operator</i> .
<b>V</b>	
<b>vaulting media</b>	The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.
<b>verify</b>	A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.
<b>Virtual Controller Software (VCS)</b>	<i>(HP P6000 EVA Disk Array Family specific term)</i> The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. <i>See also</i> HP Command View (CV) EVA.
<b>Virtual Device Interface</b>	<i>(Microsoft SQL Server specific term)</i> This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.
<b>virtual disk</b>	<i>(HP P6000 EVA Disk Array Family specific term)</i> A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. <i>See also</i> source volume and target volume.
<b>virtual full backup</b>	An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.
<b>Virtual Library System (VLS)</b>	A disk-based data storage device hosting one or more virtual tape libraries (VTLs).
<b>virtual server</b>	A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.
<b>virtual tape</b>	<i>(VLS specific term)</i> An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. <i>See also</i> Virtual Library System (VLS) and virtual tape library (VTL).
<b>virtual tape library (VTL)</b>	<i>(VLS specific term)</i> An emulated tape library that provides the functionality of traditional tape-based storage. <i>See also</i> Virtual Library System (VLS).
<b>VMware management client</b>	<i>(VMware (Legacy) integration specific term)</i> The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).
<b>volser</b>	<i>(ADIC and STK specific term)</i> A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.
<b>volume group</b>	A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.
<b>volume mountpoint</b>	<i>(Windows specific term)</i> An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

<b>Volume Shadow Copy Service</b>	See Microsoft Volume Shadow Copy Service (VSS).
<b>VSS</b>	See Microsoft Volume Shadow Copy Service (VSS).
<b>VSS compliant mode</b>	<i>(HP P9000 XP Disk Array Family VSS provider specific term)</i> One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching the disks. See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.
<b>VxFS</b>	Veritas Journal Filesystem.
<b>VxVM (Veritas Volume Manager)</b>	A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

## W

<b>Wake ONLAN</b>	Remote power-up support for systems running in power-save mode from some other system on the same LAN.
<b>Web reporting</b>	The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.
<b>wildcard character</b>	A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.
<b>Windows configuration backup</b>	Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.
<b>Windows Registry</b>	A centralized database used by Windows to store configuration information for the operating system and the installed applications.
<b>WINS server</b>	A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.
<b>writer</b>	<i>(Microsoft VSS specific term)</i> A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

## X

<b>XBSA interface</b>	<i>(Informix Server specific term)</i> ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).
-----------------------	--

## Z

<b>ZDB</b>	See zero downtime backup (ZDB).
<b>ZDB database</b>	<i>(ZDB specific term)</i> A part of the IDB, storing ZDB-related information such as source volumes, replicas, and security information. The ZDB database is used in zero downtime backup, instant recovery, and split mirror restore sessions. See also zero downtime backup (ZDB).
<b>ZDB to disk</b>	<i>(ZDB specific term)</i> A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process. See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

<b>ZDB to disk+tape</b>	<p>(<i>ZDB specific term</i>) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore.</p> <p>See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.</p>
<b>ZDB to tape</b>	<p>(<i>ZDB specific term</i>) A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used.</p> <p>See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.</p>
<b>zero downtime backup (ZDB)</b>	<p>A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.</p> <p>See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.</p>

# Index

## A

### architecture

- Microsoft Hyper-V, Virtual Environment integration, , 147
- Virtual Environment integration, 21
- VMware (Legacy) integration, 90

audience, 10

## B

### backing up Microsoft Hyper-V

- backup methods, 152
- backup options, 171
- backup specification, modifying, 171
- backup specifications, creating, 167
- scheduling backups, 172
- scheduling backups, example, 172
- starting backups, 173

### backing up Microsoft Hyper-V, Virtual Environment integration, 167–174

### backing up Virtual Environment

- backup methods, 26, 28

### backing up VMware

- backup options, 61
- backup specification, modifying, 61
- backup specifications, creating, 56
- previewing backups, 65
- scheduling backups, 64
- scheduling backups, example, 64
- starting backups, 65

### backing up VMware (Legacy), 116–122

- backup methods, 95
- backup options, 119
- backup specification, modifying, 119
- backup specifications, creating, 116
- backup types, 105
- differential backups, 105
- full backups, 105
- incremental backups, 105
- previewing backups, 120
- scheduling backups, 119
- scheduling backups, example, 119
- starting backups, 121

### backing up VMware, Virtual Environment integration, 56–67

- differential backups, 29
- full backups, 29
- incremental backups, 29

### backup

- XenServer integration, 192

### backup methods

- Microsoft Hyper-V, Virtual Environment integration, , 152
- Virtual Environment integration, 26, 28
- VMware (Legacy) integration, 95

### backup options

Virtual Environment integration, 61, 171

VMware (Legacy) integration, 119

### backup sessions, scheduling

- Virtual Environment integration, 64, 172
- VMware (Legacy) integration, 119

### backup specification

- XenServer integration, 194

### backup specifications, creating

- Virtual Environment integration, 56, 167
- VMware (Legacy) integration, 116

### backup specifications, modifying

- Virtual Environment integration, 61, 171
- VMware (Legacy) integration, 119

### backup types

- VMware (Legacy) integration, 105

## C

### checking configuration

- Virtual Environment integration, 48
- VMware (Legacy) integration, 115

### concepts

- Virtual Environment integration, 25
- VMware (Legacy) integration, 89
- XenServer integration, 186

### configuring Virtual Environment integration, 41–56

### configuring Virtual Environment integration, Microsoft Hyper-V, 163–167

### configuring VMware

- checking configuration, 48
- configuring VMware (Legacy), 108–116
- checking configuration, 115

### conventions

- document, 15

### creating backup specifications

- Virtual Environment integration, 56, 167
- VMware (Legacy) integration, 116

## D

### differential backups

- VMware (Legacy) integration, 105
- VMware, Virtual Environment integration, 29

### document

- conventions, 15
- related documentation, 10

### documentation

- HP website, 10
- providing feedback, 17

## E

### examples, Virtual Environment integration

- scheduling backups, 64, 172
- starting interactive backups, 66, 174

### examples, VMware (Legacy) integration

- scheduling backups, 119
- starting interactive backups, 121

## F

### full backups

- VMware (Legacy) integration, 105
- VMware, Virtual Environment integration, 29

## H

### help

- obtaining, 16

### HP

- technical support, 16

## I

### incremental backups

- VMware (Legacy) integration, 105
- VMware, Virtual Environment integration, 29

### installing XenServer integration, 189

### interactive backups

- Virtual Environment integration, 65, 173
- VMware (Legacy) integration, 121

### introduction

- Microsoft Hyper-V, Virtual Environment integration, 146
- Virtual Environment integration, 20
- VMware (Legacy) integration, 89
- VSS integration, 152

## M

### main integration components

- XenServer integration, 189

### Microsoft Hyper-V backup

- backup methods, 152
- backup options, 171
- backup specification, modifying, 171
- backup specifications, creating, 167
- scheduling backups, 172
- scheduling backups, example, 172
- starting backups, 173

### Microsoft Hyper-V backup, Virtual Environment integration, 167–174

### Microsoft Hyper-V restore

- finding information, 174
- merging virtual machine snapshots, 179
- restore options, 177
- using another device, 180
- using CLI, 177
- using GUI, 175

### Microsoft Hyper-V restore, Virtual Environment integration, 174

### Microsoft Hyper-V troubleshooting, Virtual Environment integration, 180–184

### Microsoft Hyper-V, Virtual Environment integration

- introduction, 146

### modifying backup specifications

- Virtual Environment integration, 61, 171
- VMware (Legacy) integration, 119

### monitoring sessions

- Virtual Environment integration, 180
- VMware (Legacy) integration, 140
- VMware, Virtual Environment integration, 84

## O

### offline backup

- XenServer integration, 186

### online backup

- XenServer integration, 186

### online backups

- Microsoft Hyper-V, Virtual Environment integration, 146
- Virtual Environment integration, 20
- VMware (Legacy) integration, 89

## P

### previewing backups

- Virtual Environment integration, 65
- VMware (Legacy) integration, 120

## R

### related documentation, 10

### restore

- XenServer integration, 195

### restore destination

- Virtual Environment integration, 74

### restore from offline backup

- XenServer integration, 187

### restore from online backup

- XenServer integration, 187

### restore options

- Virtual Environment integration, 75, 177
- VMware (Legacy) integration, 127–128

### restore processes

- XenServer integration, 187

### restoring Microsoft Hyper-V

- finding information, 174
- merging virtual machine snapshots, 179
- restore options, 177
- using CLI, 177
- using GUI, 175

### restoring Microsoft Hyper-V Server

- using another device, 180

### restoring Microsoft Hyper-V, Virtual Environment integration, 174

### restoring VMware

- finding information, 68
- restore destination, 74
- restore options, 75
- using another device, 83
- using CLI, 76
- using GUI, 69

### restoring VMware (Legacy), 122–140

- finding information, 122
- restore options, 127–128
- using another device, 140
- using CLI, 128
- using GUI, 123

### restoring VMware, Virtual Environment integration, 67–84

## S

### scheduling backups

- Virtual Environment integration, 64, 172

- VMware (Legacy) integration, 119
- starting backups
  - Virtual Environment integration, 65, 173
  - VMware (Legacy) integration, 121
- Subscriber's Choice, HP, 16

## T

- technical support
  - HP, 16
  - service locator website, 17
- troubleshooting Virtual Environment integration, Microsoft Hyper-V, 180–184
- troubleshooting Virtual Environment integration, VMware, 84–88
- troubleshooting VMware (Legacy), 141

## V

- VEPA
  - Virtual Environment integration, 24
- Virtual Environment backup
  - backup methods, 26, 28
- Virtual Environment integration
  - architecture, 21
  - backup, 56
  - concepts, 25
  - configuration, 41–56
  - introduction, 20
  - monitoring sessions, 180
  - restore, 67–84
- Virtual Environment integration configuration, 41–56
- Virtual Environment integration configuration, Microsoft Hyper-V, 163–167
- Virtual Environment integration troubleshooting, VMware, 84–88
- Virtual Environment integration, Microsoft Hyper-V
  - architecture, 147
  - backup, 167–174
  - configuration, 163–167
  - restore, 174
  - troubleshooting, 180–184
- Virtual Environment integration, VMware
  - troubleshooting, 84–88
- VMware (Legacy) backup, 116–122
  - backup methods, 95
  - backup options, 119
  - backup specification, modifying, 119
  - backup specifications, creating, 116
  - backup types, 105
  - differential backups, 105
  - full backups, 105
  - incremental backups, 105
  - previewing backups, 120
  - scheduling backups, 119
  - scheduling backups, example, 119
  - starting backups, 121
- VMware (Legacy) configuration, 108–116
  - checking configuration, 115
- VMware (Legacy) integration
  - architecture, 90

- backup, 116–122
  - concepts, 89
  - configuration, 108–116
  - introduction, 89
  - monitoring sessions, 140
  - restore, 122–140
  - troubleshooting, 141
- VMware (Legacy) restore, 122–140
  - finding information, 122
  - restore options, 127–128
  - using another device, 140
  - using CLI, 128
  - using GUI, 123
- VMware (Legacy) troubleshooting, 141
- VMware backup
  - backup options, 61
  - backup specification, modifying, 61
  - backup specifications, creating, 56
  - previewing backups, 65
  - scheduling backups, 64
  - scheduling backups, example, 64
  - starting backups, 65
- VMware backup, Virtual Environment integration, 56–67
  - differential backups, 29
  - full backups, 29
  - incremental backups, 29
- VMware configuration
  - checking configuration, 48
- VMware management client
  - VMware (Legacy) integration, 109
- VMware restore
  - finding information, 68
  - restore destination, 74
  - restore options, 75
  - using another device, 83
  - using CLI, 76
  - using GUI, 69
- VMware restore, Virtual Environment integration, 67–84
- VMware, Virtual Environment integration
  - monitoring sessions, 84
- Volume Shadow Copy Service
  - VSS integration, 152
- VSS integration
  - introduction, 152

## W

- websites
  - HP, 17
  - HP Subscriber's Choice for Business, 16
  - product manuals, 10

## X

- XenServer integration
  - backup, 192
  - backup specification, 194
  - concepts, 186
  - installation, 189
  - integration scripts, 191
  - main integration components, 189

- offline backup, [186](#)
- online backup, [186](#)
- restore, [195](#)
- restore from offline backup, [187](#)
- restore from online backup, [187](#)
- restore processes, [187](#)
- specifying restore, [197](#)
- updating DPxen\_config.py for backup, [192](#)
- updating DPxen\_config.py for restore, [195](#)