HP Data Protector 6.20 Integration Guide for Virtualization Environments

HP Part Number: N/A
Published: December 2011
Edition: Fourth

© Copyright 2011 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.

Contents

Publication history....9
About this guide....10
Intended audience....10
Documentation set....10
Guides....10
Online Help....12
Documentation map....13
Abbreviations....13
Map....13
Integrations....14
Document conventions and symbols....15
Data Protector graphical user interface....15
General information....16
HP technical support....16
Subscription service....16
HP websites....16
Documentation feedback....17
I VMware....18
1 Data Protector Virtual Environment integration....20
Introduction....20
Integration concepts....20
Supported environments....20
vCenter environment....21
Migration of virtual machines....22
vCenter Server system in a cluster....23
Standalone ESX/ESXi Server environment....23
vCloud Director environment....24
Data Protector components....27
Data Protector Cell Manager....27
Data Protector Virtual Environment Integration component....27
Data Protector Disk Agent component....28
Data Protector Media Agents....28
Backup concepts....28
What is backed up?....28
Virtual machines....28
Virtual machine templates....29
vStorage Image backup method....29
Snapshot management....30
Backup types....31
Snapshot handling mode....32
Changed Block Tracking....38
Quiescence....39
Disk space requirements....40
Backup parallelism....40
Backup considerations....40
Restore concepts....41
Restore to a datacenter....41
Restore of individual virtual machine disks....41
Restore to a directory....42
Restore chain....42

Restore considerations....42
Performance....43
Configuring the integration....43
Prerequisites....43
Before you begin....43
Importing and configuring VMware clients....43
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....48
Using the Data Protector GUI....49
Using the Data Protector CLI....52
Customizing the Data Protector behavior with omnirc variables....54
Backup....54
Creating backup specifications....54
Modifying backup specifications....57
Scheduling backup specifications....58
Scheduling example....59
Previewing backup sessions....59
Using the Data Protector GUI....59
Using the Data Protector CLI....59
What happens during the preview?....60
Starting backup sessions....60
Using the Data Protector GUI....60
Using the Data Protector CLI....60
Preparing for disaster recovery....61
Restore....61
Finding information for restore....61
Using the Data Protector GUI....61
Using the Data Protector CLI....62
Restoring using the Data Protector GUI....62
Restoring using the Data Protector CLI....66
Recovering virtual machines manually....67
Recovering virtual machines after restore to a directory....67
Recovering virtual machines after restore to a datacenter....71
Restoring using another device....72
Cleaning up a datastore after a failed restore....72
Disaster recovery....72
Monitoring sessions....72
Troubleshooting....72
Before you begin....73
Checks and verifications....73
Problems....73
2 Data Protector VMware (Legacy) integration....77
Introduction....77
Integration concepts....77
Supported environments....78
VirtualCenter environment....78
Migration of virtual machines....79
VirtualCenter Server system in a cluster....80
Standalone ESX Server environment....80

ESXi Server systems....80
Backup....80
Restore....80
Limitations in VMware vSphere environments....81
Data Protector installation components....81
Data Protector Cell Manager....81
Data Protector VMware Integration (Legacy) component....81
Data Protector Media Agents....81
Backup concepts....81
What is backed up?....82
Virtual machines....82
Filesystems....82
Backup methods....83
Snapshot method....83
Suspend method....89
VCBimage method....90
VCBfile method....91
Backup types....93
Considerations....93
Restore concepts....94
Restore of virtual machines....94
Restore to datacenter....94
Restore to default directory....95
Restore of filesystems....95
Restore considerations....95
Configuring the integration....96
Prerequisites....96
Before you begin....96
Configuring VMware users....97
Configuring clusters....97
ESX Server systems in a cluster....97
Configuring backup proxy systems....97
Configuring VMware management clients....97
Using the Data Protector GUI....98
Using the Data Protector CLI....99
Configuring virtual machines....100
Using the Data Protector GUI....100
Using the Data Protector CLI....102
Checking the configuration of VMware clients....103
Using the Data Protector GUI....103
Using the Data Protector CLI....103
Backup....104
Creating backup specifications....104
Modifying backup specifications....107
Scheduling backup specifications....107
Scheduling example....107
Previewing backup sessions....108
Using the Data Protector GUI....108
Using the Data Protector CLI....108
What happens during the preview?....108
Starting backup sessions....109
Using the Data Protector GUI....109
Using the Data Protector CLI....109
Preparing for disaster recovery....110
Restore....110

Finding information for restore....110
Using the Data Protector GUI....110
Using the Data Protector CLI....111
Restoring using the Data Protector GUI....111
Restoring using the Data Protector CLI....116
Recovering virtual machines manually....118
Recovery after restoring outside a datacenter....118
Recovery after restoring to a datacenter....123
Linking virtual machine disk files....123
Consolidating virtual machine disks....125
Importing virtual machine disks....125
Restoring using another device....128
Disaster recovery....128
Monitoring sessions....128
Troubleshooting....128
Before you begin....128
Checks and verifications....129
Problems....129
II Microsoft Hyper-V....131
3 Data Protector Virtual Environment integration....132
Introduction....132
Integration concepts....133
Supported environments....133
Standalone environments....133
Clustered environments....133
Migration of virtual machines....136
Cluster Shared Volumes....137
Data Protector installation components....137
Data Protector Cell Manager....137
Data Protector Virtual Environment Integration component....137
Data Protector Disk Agent component....137
Data Protector MS Volume Shadow Copy Integration component....137
Data Protector Media Agents....137
Backup concepts....138
Hyper-V Image backup method....138
Backup types....139
Microsoft Hyper-V backup types....140
VSS backup types....140
Quiescence....140
Backup considerations....141
Restore concepts....141
Restore of virtual machines....141
Restore to the default location....142
Restore to a directory 142
Restore considerations....142
Configuring the integration....143
Prerequisites....143
Limitations....144
Before you begin....144
Enabling automatic mounting on Microsoft Hyper-V systems....144
Configuring Microsoft Hyper-V clusters....144
Importing and configuring Microsoft Hyper-V systems....144
Changing the configuration of Microsoft Hyper-V systems....146
Using the Data Protector GUI....146

Using the Data Protector CLI....146
Customizing the Data Protector behavior with omnirc variables....147
Backup....147
Creating backup specifications....147
Modifying backup specifications....150
Scheduling backup specifications....151
Scheduling example....152
Starting backup sessions....152
Using the Data Protector GUI....152
Using the Data Protector CLI....152
Restore....153
Finding information for restore....153
Using the Data Protector GUI....153
Using the Data Protector CLI....153
Restoring using the Data Protector GUI....154
Restoring using the Data Protector CLI....155
Restore of cluster-aware virtual machines....156
Restoring using another device....156
Restore options....156
Monitoring sessions....157
Troubleshooting....157
Before you begin....157
Checks and verifications....157
Problems....158
III Citrix XenServer....161
4 Data Protector Citrix XenServer script solution....162
Introduction....162
Integration concepts....162
Types of backup....162
Online backup....162
Offline backup....162
Disaster recovery....162
Backup processes....163
Online backup....163
Offline backup....163
Types of restore....163
Restore processes....163
Restore from online backup....163
Restore from offline backup....164
Restore considerations....164
Main integration components....165
Installation of the integration....165
Prerequisites....165
Installation of the integration scripts....166
Integration script functions....167
Backup using the integration....168
Updating configuration script DPxen_config.py for backup....168
Example backup configurations....169
Creating a backup specification....170
Restore using the integration....171
Updating configuration script DPxen_config.py for restore....171
Example restore configurations....172
Specifying a restore....173
Notes on restore....174

Special considerations.....175
Further information.....175
Glossary.....176
Index.....206

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	March 2011	Data Protector Release 6.20
N/A	March 2011 (second edition)	Data Protector Release 6.20
N/A	December 2011	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183
N/A	December 2011 (fourth edition)	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183

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

Other documents and online Help provide related information.

Guides

Data Protector guides are available in the electronic PDF format. Install the PDF files during the Data Protector setup procedure by selecting the `English Documentation (Guides, Help)` component on Windows or the `OB2-DOCS` component on UNIX. Once installed, the guides reside in the *`Data_Protector_home`*`\docs` directory on Windows and in the `/opt/omni/doc/C` directory on UNIX.

You can find these documents from the Manuals page of the HP Information Management Digital Hub website:

http://www.hp.com/go/imhub

In the Storage section, click **Storage Software** and then select your product.

- *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 online 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 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 or 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, Network Node Manager, and Network Data Management Protocol Server*

    This guide describes the integrations of Data Protector with Sybase Server, HP Network Node Manager, and Network Data Management Protocol Server.

  - *HP Data Protector Integration Guide for Virtualization Environments*

    This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure and VMware vSphere, Microsoft Hyper-V, and Citrix XenServer.

  - *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 HP Operations Manager for UNIX*

  This guide describes how to monitor and manage the health and performance of the Data Protector environment with HP Operations Manager and HP Service Navigator on UNIX.

- *HP Data Protector Integration Guide for HP Operations Manager for Windows*

  This guide describes how to monitor and manage the health and performance of the Data Protector environment with HP Operations Manager on Windows.

- *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 P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP P4000 SAN Solutions, 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 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 Media Operations User Guide*

  This guide provides information for network administrators responsible for maintaining and backing up systems on the tracking and management of offline storage media. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.

- *HP Data Protector Product Announcements, Software Notes, and References*

  This guide gives a description of new features of HP Data Protector 6.20. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.

- *HP Data Protector Product Announcements, Software Notes, and References for Integrations to HP Operations Manager*

  This guide fulfills a similar function for the HP Operations Manager integration.

- *HP Data Protector Media Operations Product Announcements, Software Notes, and References*

  This guide fulfills a similar function for Media Operations.

- *HP Data Protector Command Line Interface Reference*

  This guide describes the Data Protector command-line interface, command options and their usage as well as providing some basic command-line examples.

Online Help

Data Protector provides Help topics and context-sensitive (F1) Help for Windows and UNIX platforms.

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

- ***Windows:*** Open `DP_help.chm`.
- ***UNIX:*** Unpack the zipped tar file `DP_help.tar.gz`, and access the online Help system through `DP_help.htm`.

Documentation map

Abbreviations

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

Abbreviation	Guide
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE-SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE-VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Online Help
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-O/S	Integration Guide for Oracle and SAP
IG-OMU	Integration Guide for HP Operations Manager for UNIX
IG-OMW	Integration Guide for HP Operations Manager for Windows
IG-Var	Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol Server
IG-VirtEnv	Integration Guide for Virtualization Environments
IG-VSS	Integration Guide for Microsoft Volume Shadow Copy Service
Install	Installation and Licensing Guide
MO GS	Media Operations Getting Started Guide
MO RN	Media Operations Product Announcements, Software Notes, and References
MO UG	Media Operations User Guide
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concept	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.

	Help	GS	Concepts	Install	Trouble	DR	PA	Integration Guides: MS	O/S	IBM	Var	VSS	VirtEnv	OMU	OMW	ZDB: Concept	Admin	IG	GRE: SPS	VMware	MO: GS	User	PA	CLI
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	X				
Disaster recovery	X		X			X																		
Installation/ upgrade	X	X		X			X							X	X						X	X		
Instant recovery	X		X													X	X	X						
Licensing	X			X			X															X		
Limitations	X				X		X	X	X	X	X	X	X					X					X	
New features	X						X																X	
Planning strategy	X		X													X								
Procedures/ tasks	X			X	X	X		X	X	X	X	X	X	X	X		X	X	X	X		X		
Recommendations			X				X									X							X	
Requirements				X			X	X	X	X	X	X	X	X	X						X	X	X	
Restore	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	X				

Integrations

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

Software application	Guides
HP Network Node Manager (NNM)	IG-Var
HP Operations Manager	IG-OMU, IG-OMW
IBM DB2 UDB	IG-IBM
Informix Server	IG-IBM
Lotus Notes/Domino Server	IG-IBM
Media Operations	MO User
Microsoft Exchange Server	IG-MS, ZDB IG
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

Software application	Guides
Sybase Server	IG-Var
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 Concept, ZDB Admin, IG-VSS
HP P6000 EVA Disk Array Family	all ZDB, IG-VSS
HP P9000 XP Disk Array Family	all ZDB, IG-VSS

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: "Document conventions" (page 15)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` 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 cross-platform (Windows and UNIX) graphical user interface. You can use the original Data Protector GUI (Windows only) or the Data Protector Java GUI. For information about the Data Protector graphical user interface, see the online Help.

Figure 1 Data Protector graphical user interface

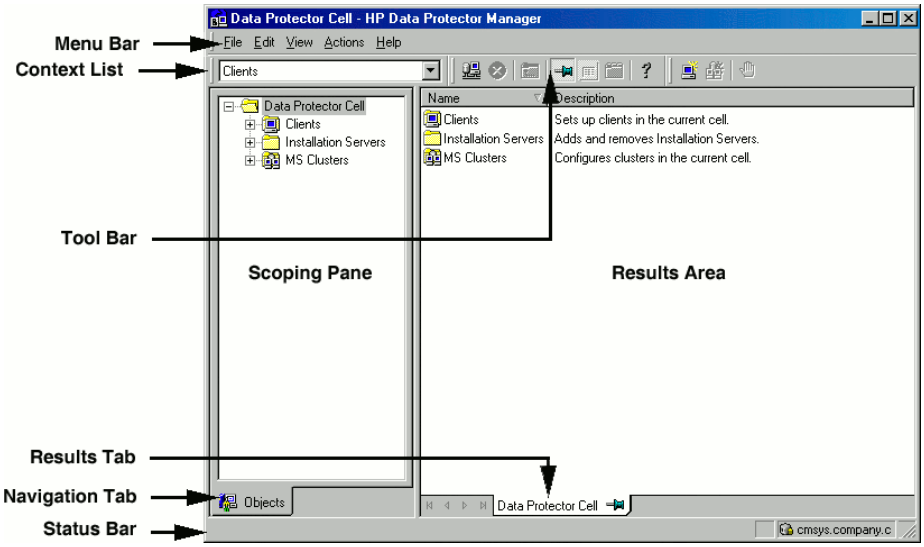

General information

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://www.hp.com/go/imhub
- 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 to **DP.DocFeedback@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 77).

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 77).

NOTE: You can also back up VMware virtual machine data 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.

Table 3 Data Protector backup solutions for VMware

Feature		Filesystem backup	VMware (Legacy) integration	Virtual Environment integration
Online backup			✓[1]	✓
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?			✓[2]	
Is Changed Block Tracking supported?				✓
Where does the Data Protector integration component need to be installed?		ESX(i) Server system	• ESX(i) Server system • vCenter Server system • Backup proxy	Backup host[3]
Extra licenses needed		optionally ZDB and IR	Online Extension	Online Extension

[1] In a VMware vSphere environment, online backup of virtual machines using the Data Protector VMware (Legacy) integration "Snapshot method" (page 83) is not supported.

[2] The VCB software is needed to back up virtual machines that reside on an ESXi Server system.

[3] 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. Data Protector integrates with VMware vSphere, including ESX Server, ESXi Server systems, and vCenter Server systems, to back up and restore the following **VMware objects**:

- Virtual machines
- Virtual machine disks
- Virtual machine templates

Backup

The following backup method is available:

- vStorage Image

This is a snapshot-based method that can be used to back up virtual machines while powered off (**offline backup**) or actively used (**online backup**).

For details on the backup method, see "vStorage Image backup method" (page 29).

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

- Full
- Incremental
- Differential

For details on the backup types see, "Backup types" (page 31).

Restore

Virtual machines can be restored:

- **To a datacenter**

  This datacenter can be a standalone ESX(i) Server system or any of the datacenters created and managed by a vCenter Server system.

- **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. For limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the online Help.

Integration concepts

Supported environments

Data Protector supports environments where ESX Server and/or ESXi Server systems are managed 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.

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

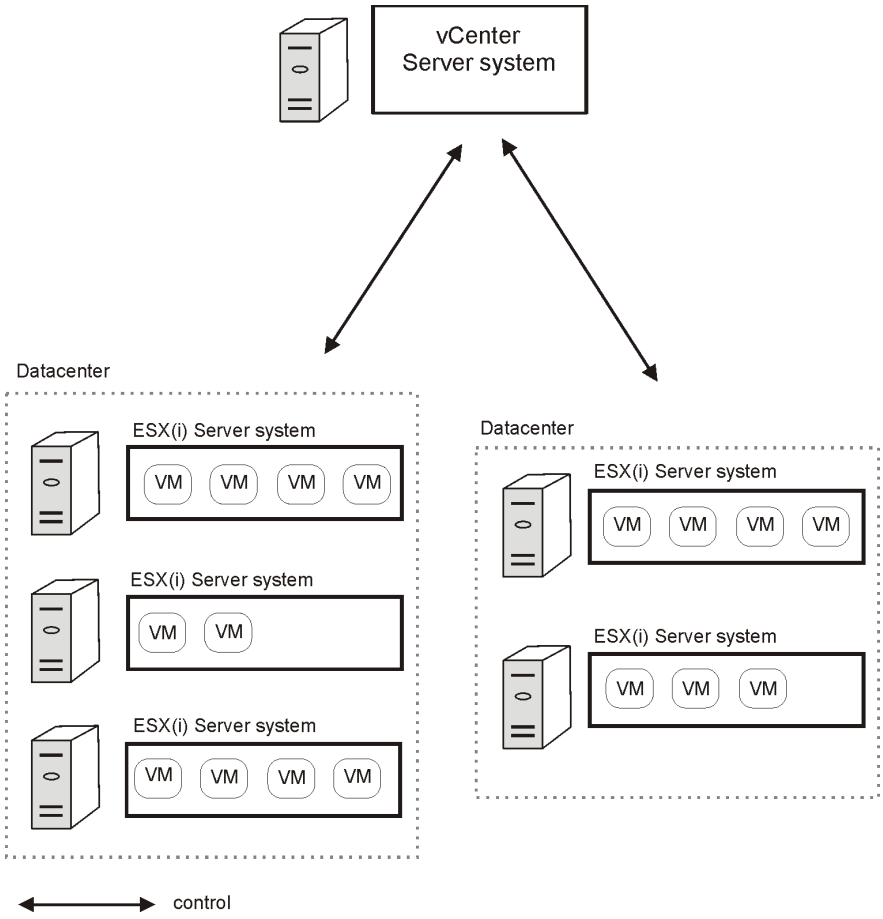

ESX Server or ESXi Server system	VMware platform capable of hosting multiple virtual machines.
VM	Virtual machine. Virtualized x86 or x64 PC environment, in which a guest operating system and associated application software can run.
Datacenter	An organizational unit that consists of one or more ESX Server 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 supporting 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

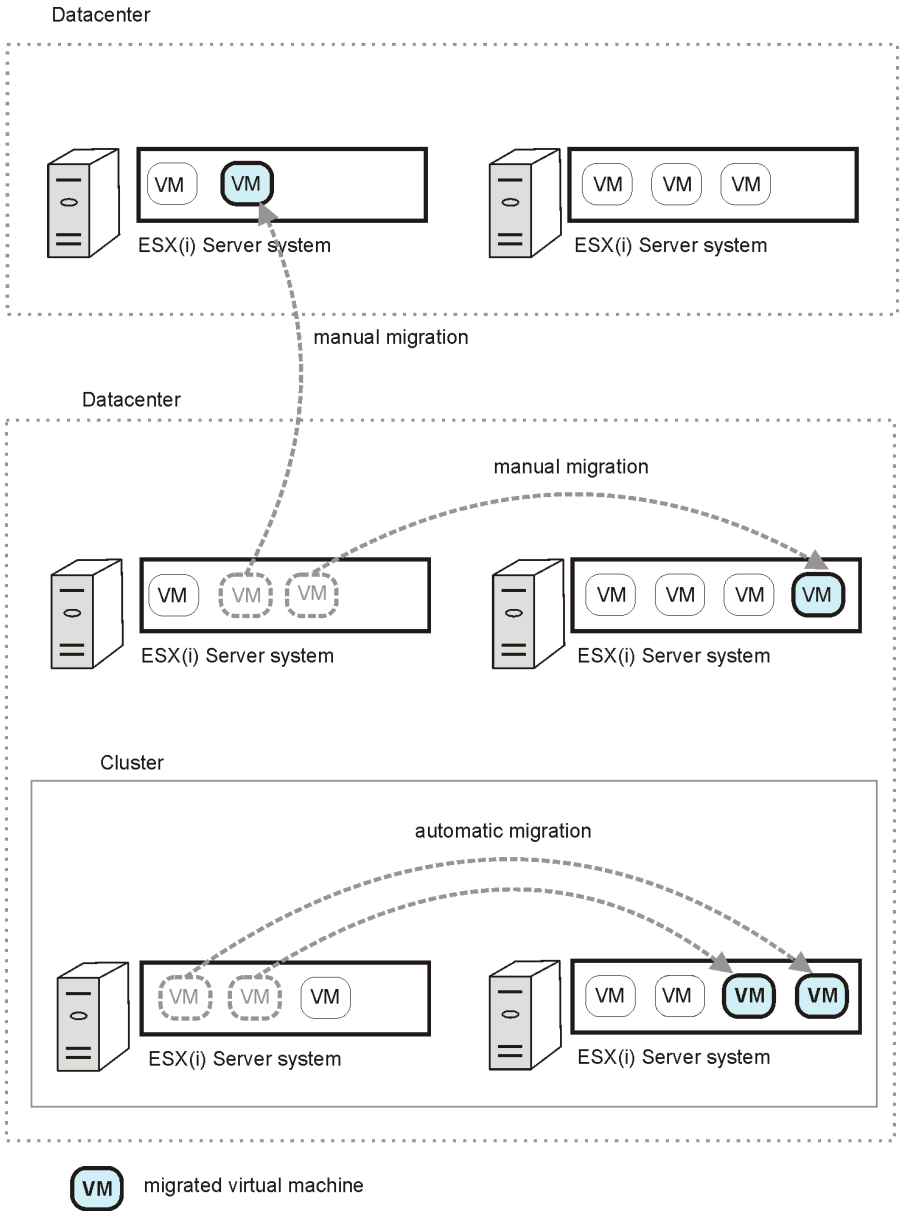

Whatever the reason for the migration, you do not need to create a new backup specification afterwards. Data Protector will automatically find the migrated virtual machines and back them up.

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

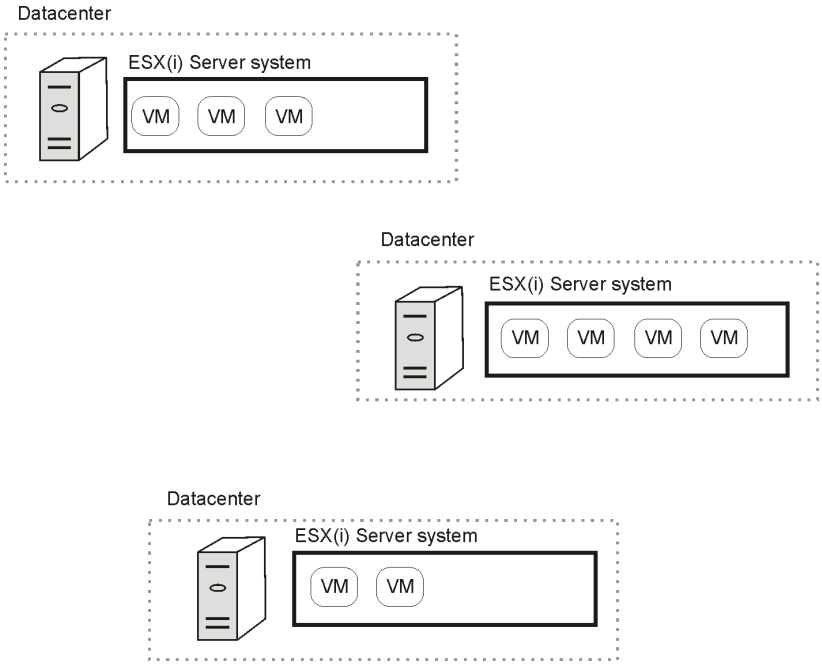

vCloud Director environment

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, ensure that you import all vCenter Server systems used by VMware vCloud Director in the Data Protector Cell as **VMware vCenter** clients.

Backup

To back up virtual machines from a virtual datacenter, first discover in which VMware vSphere datacenter they are located. After that, the procedure is the same as for other virtual machines in VMware vSphere environments.

Figure 5 vCenter Server systems used by VMware vCloud Director

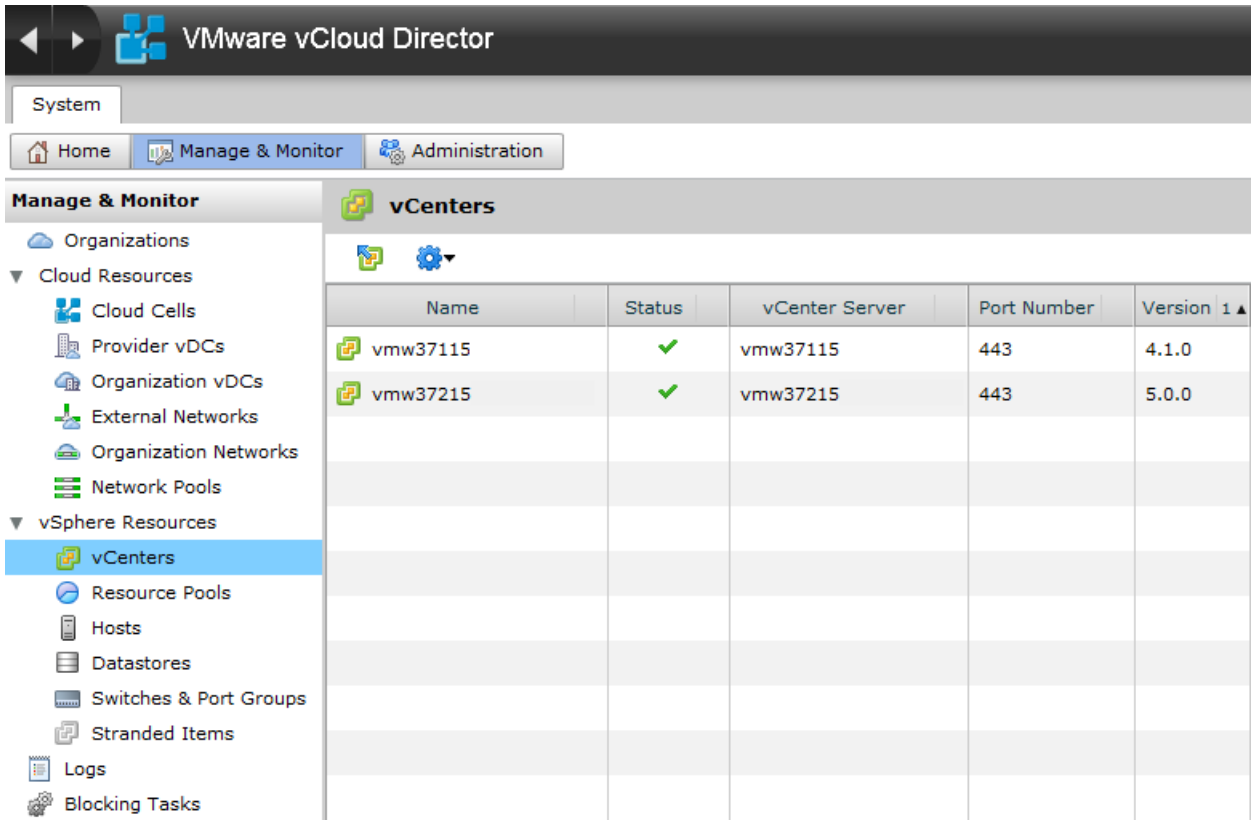

Example

Suppose you created the following organization in a VMware vCloud Director environment:

Table 4 An organization in a VMware vCloud Director environment

Organization	`GlenMorangie`
Virtual datacenter (that belongs to this organization)	`16-Men-Of-Tain`
vApp virtual machines	• `jeos_test1` • `ubuntu_01`

Figure 6 An organization in a VMware vCloud Director environment

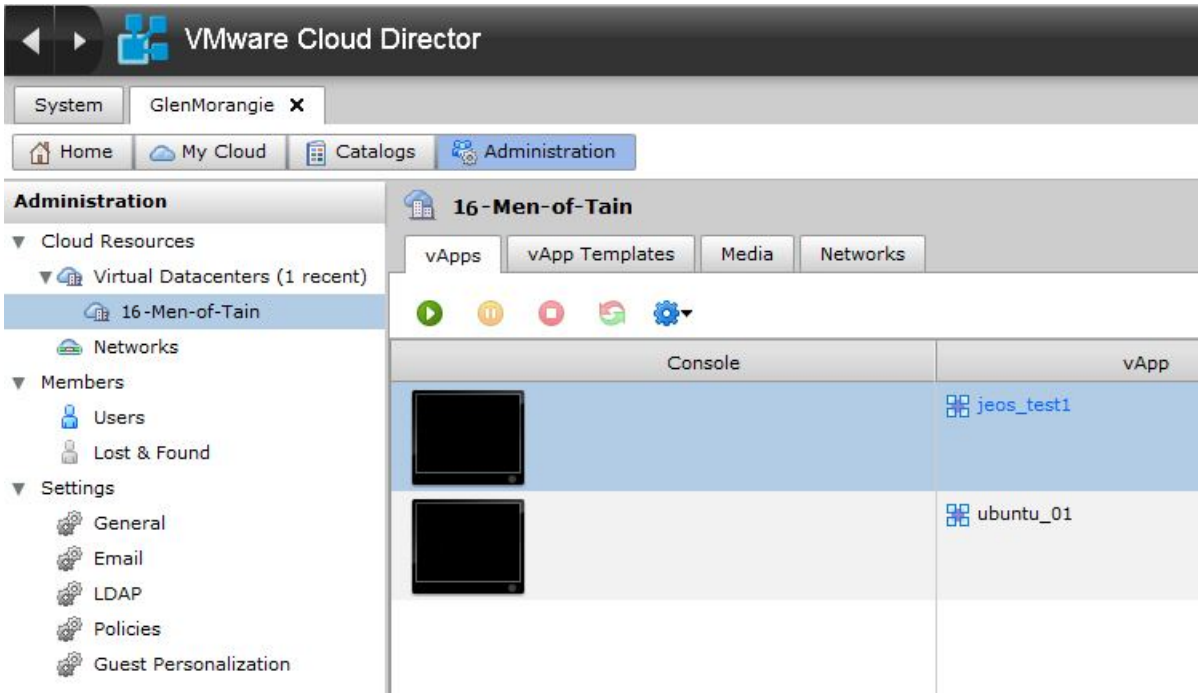

In our example, virtual machines from the virtual datacenter `16-Men-Of-Tain` are located in the VMware vSphere datacenter `vCD` that is managed by the vCenter Server system `vmw37115` See Figure 7 (page 26).

Figure 7 VMware vSphere infrastructure

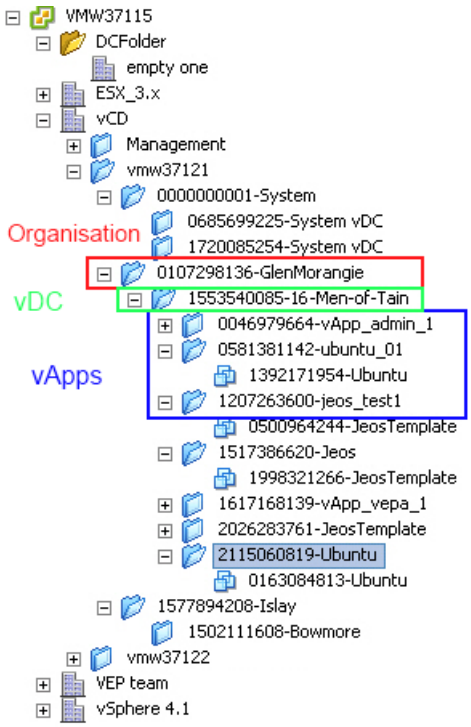

TIP: If you do not know exactly in which VMware vSphere datacenter the virtual machines are located, browse all available datacenters in the VMware vSphere infrastructure.

Restore

VMware vCloud Director uses its own directory structure in VMware vSphere. See Figure 7 (page 26). Because of this, it is recommended that you first restore virtual machines to a different location and then import them in the VMware vCloud Director environment using the native Import functionality. See Figure 8 (page 27).

Figure 8 Importing virtual machines in a VMware vCloud Director environment

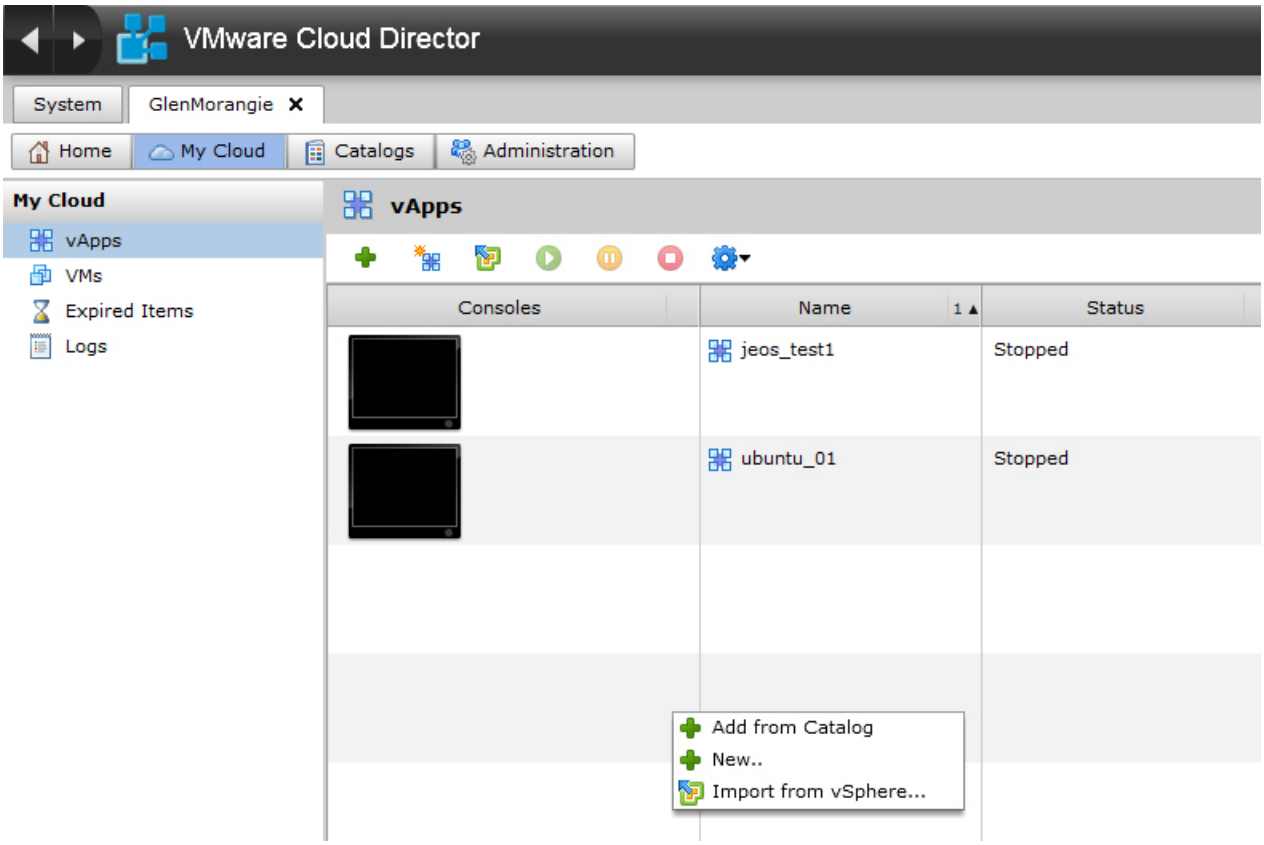

Data Protector components

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_hyperv.dll`, a dynamic link library for Microsoft Hyper-V specific backup, restore, query, and browsing tasks.[1]

NOTE: VEPA stands for Virtual Environment Protection Agent.

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 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:

- Virtual machines
- Virtual machine disks
- Virtual machine templates

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`

Virtual machines

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

- .vmx
- .vmdk
- .vmsn

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.

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 (so that, it applies to all virtual machines selected in the backup specification).

Virtual machine disks

Data Protector supports backup of individual virtual machine disks. 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 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 via a vCenter Server system (in the case of a vCenter environment), or direct (in the case of a stand-alone 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 variable to specify the time interval between opening two different disks. By default, a new disk is opened every 2 seconds.

Figure 9 vStorage Image method

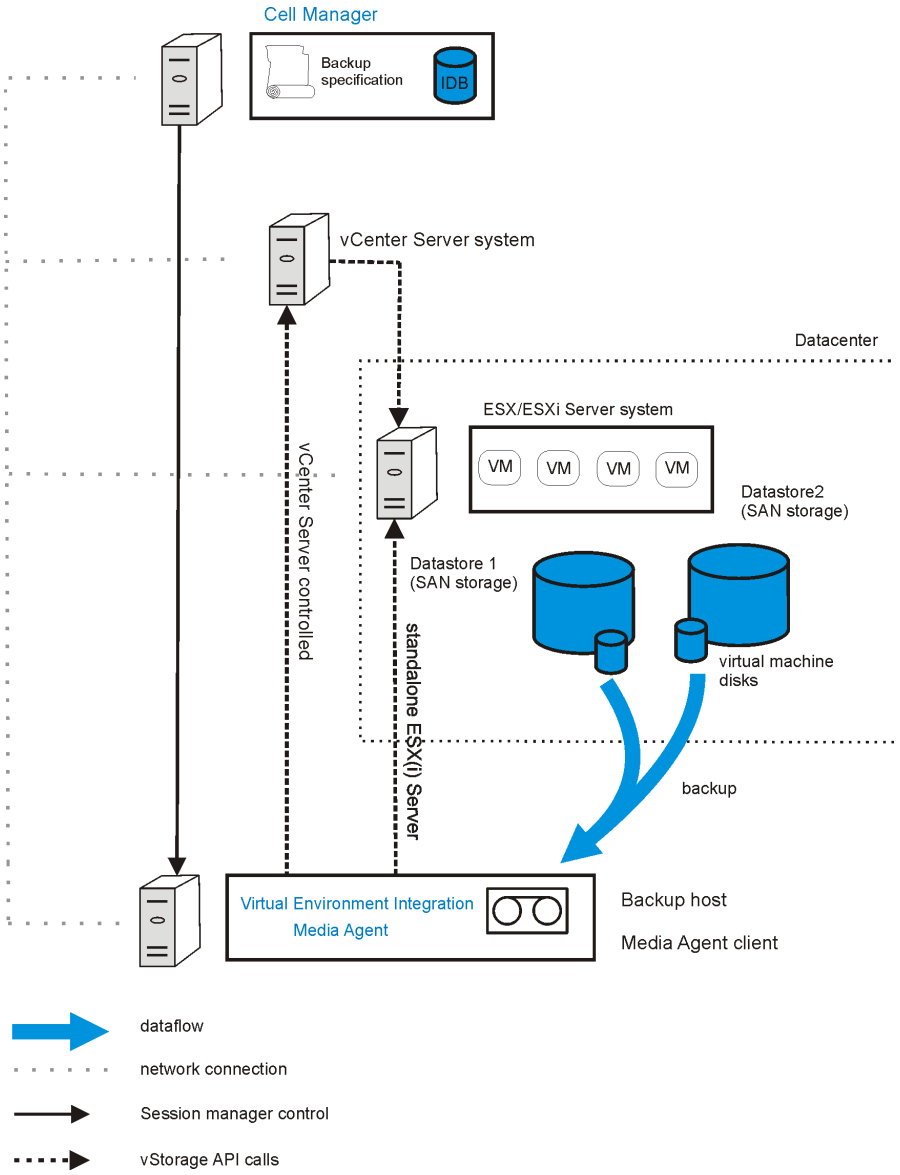

In "vStorage Image method" (page 30), the backup host is also a media agent client (it has the `Media Agent` component installed and a device connected to it).

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 then copies data from the consistent state to Data Protector media. Snapshots created by Data Protector (**DP snapshots**) are distinguished from other snapshots by the label `_DP_VEPA_SNAP_` and a description containing the product name and a timestamp.

NOTE: It does not matter if a user-created snapshot exists on 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 10 Virtual machine snapshot tree

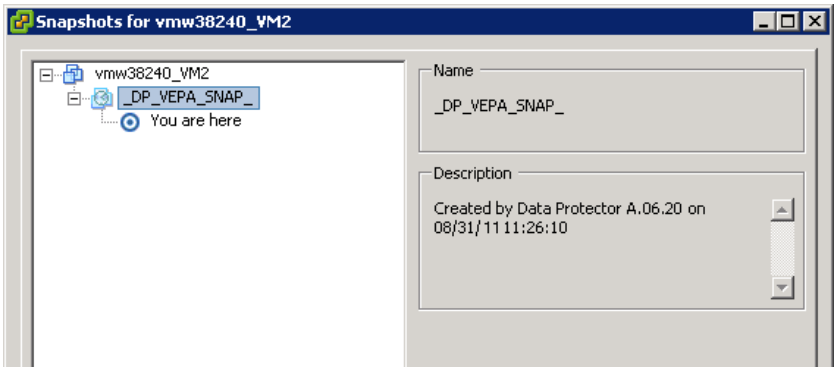

The presence of virtual-machine snapshots degrades a virtual machine's performance, so DP snapshots are deleted when they are no longer needed.

The number of DP snapshots that remain on 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 on a snapshot tree play a great role for Incremental and Differential virtual machine backups. For details, see "Backup types" (page 31).

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

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 backup method, you can perform the following backup types:

Table 5 Backup types

Full	Backs up the complete virtual machine (disk), including the virtual machine memory file (if specified).
Incr	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 whether Data Protector should identify changes:

- At *file level* or
- 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 32).

To identify changes at disk block level, Data Protector uses the VMware Changed Block Tracking functionality. For details, see "Changed Block Tracking" (page 38).

The possible combinations of snapshot handling mode and changed block tracking for VMware backups are summarized in Table 6 (page 32).

Table 6 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 allows you to control the number of Data Protector snapshots left on 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. However, the types of backup that can be performed depend on the mode setting as follows:

- **Disabled**: This mode 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 33).
- **Single**: This mode supports Full, Differential, and Incremental backups. The following backup chains are supported:

  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 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 34).

- **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 36).

IMPORTANT: A backup chain is broken if you do any of the following:

- 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
- Enable changed block tracking

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 11 Full backup (disabled mode)

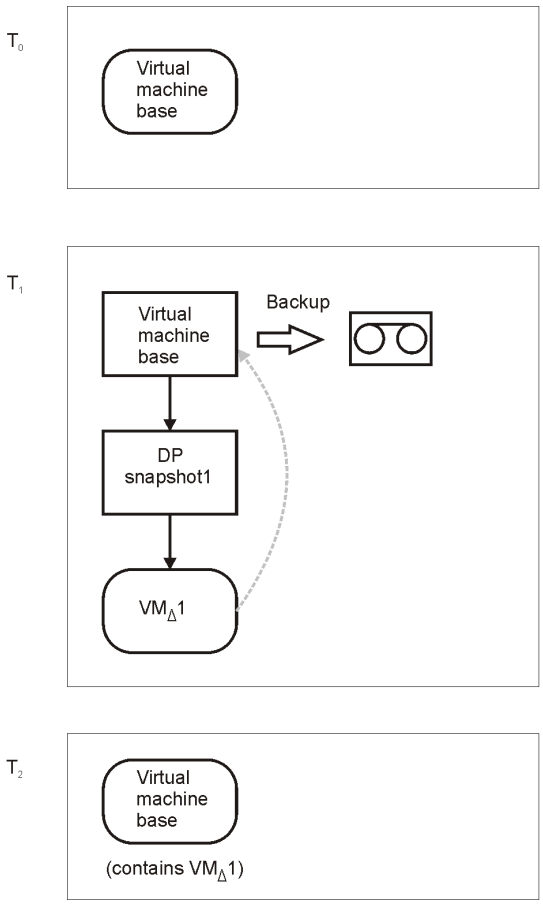

Table 7 Legend

Tn	Boxes denoted by Ti 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.
VMΔ	A rectangle denoted by VMΔ 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 `_DP_VEPA_SNAP_`.

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 12 Differential backup (single mode)

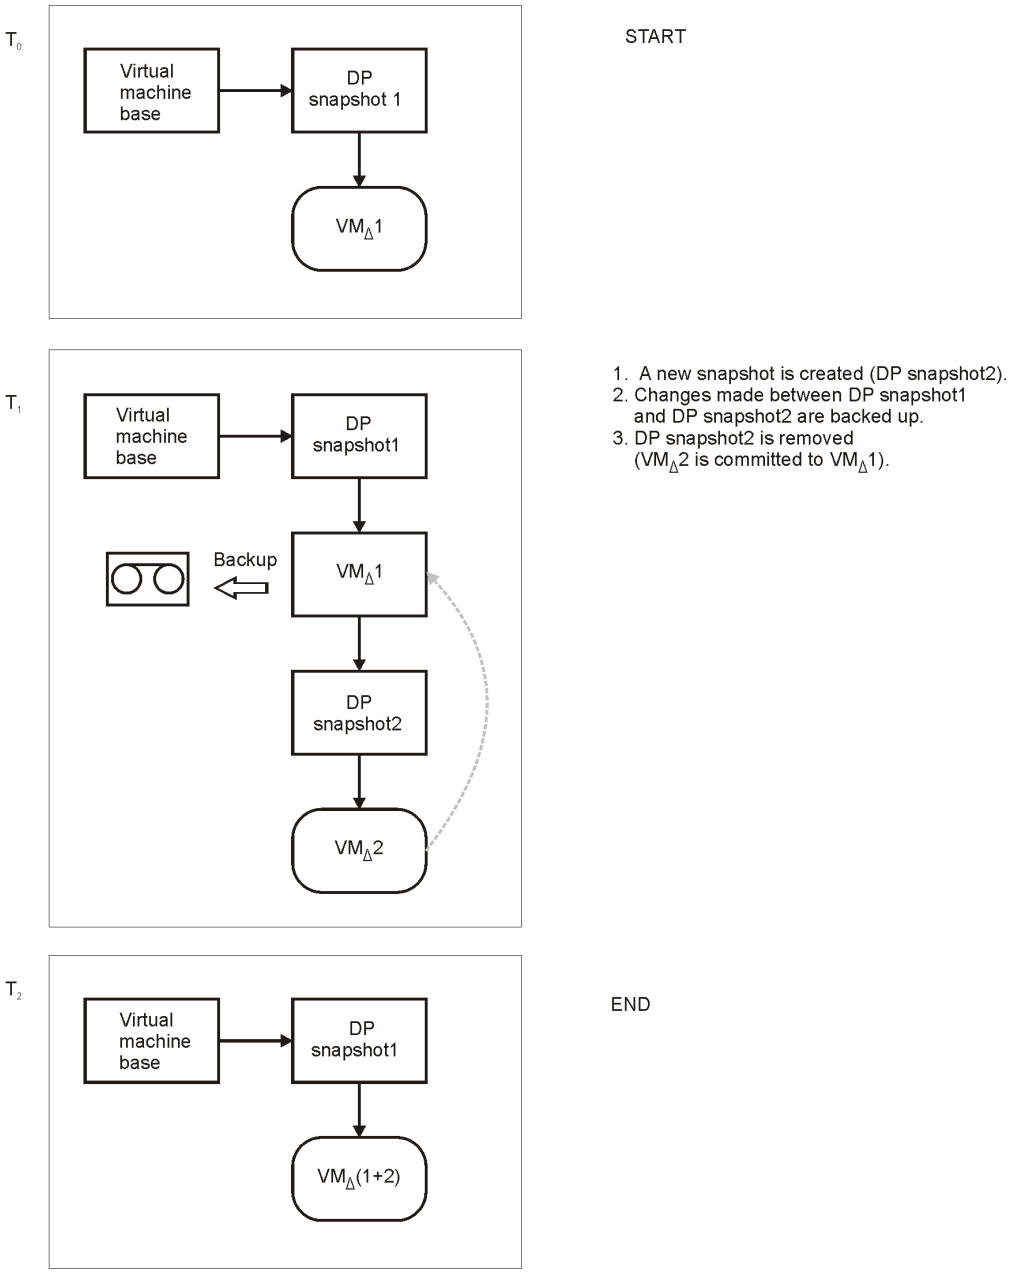

DP snapshot1 remains on the snapshot tree to track changes made since the last Full backup.

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 "Incremental backup (single mode)" (page 36)).

Figure 13 Incremental backup (single mode)

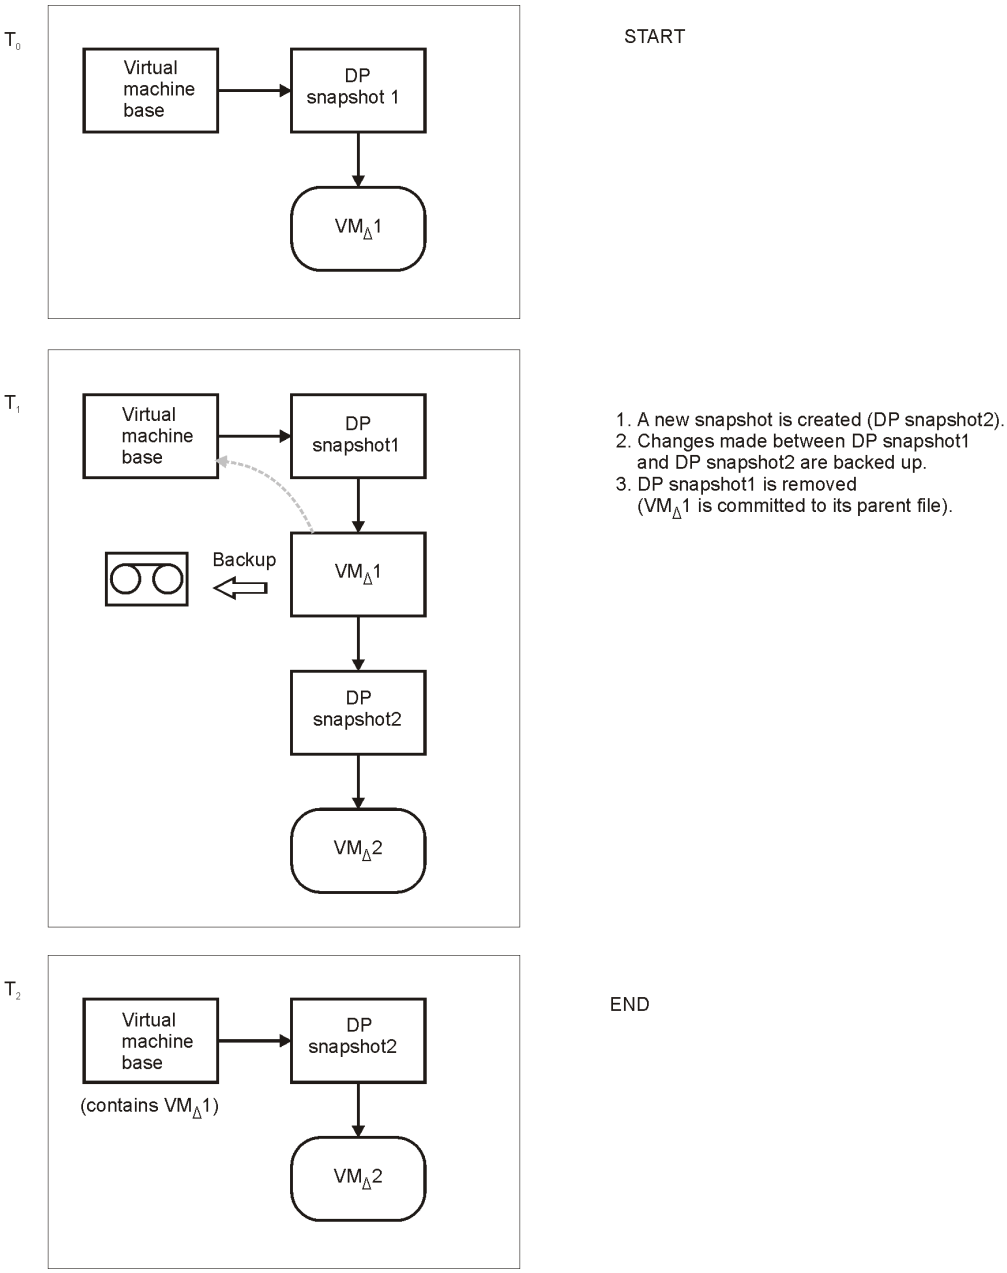

DP snapshot2 remains on the snapshot tree to track changes made since the last Incremental backup.

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_0

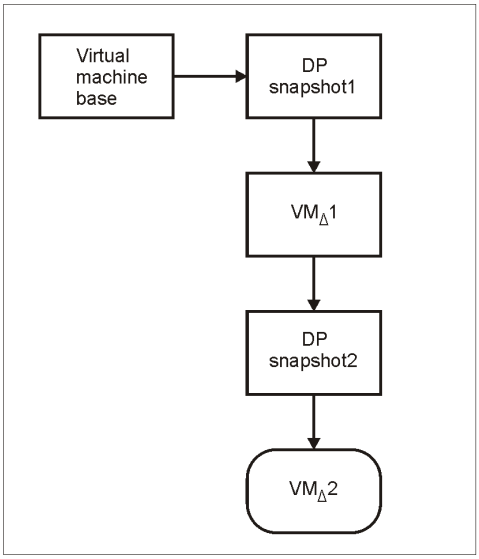

START

T_1

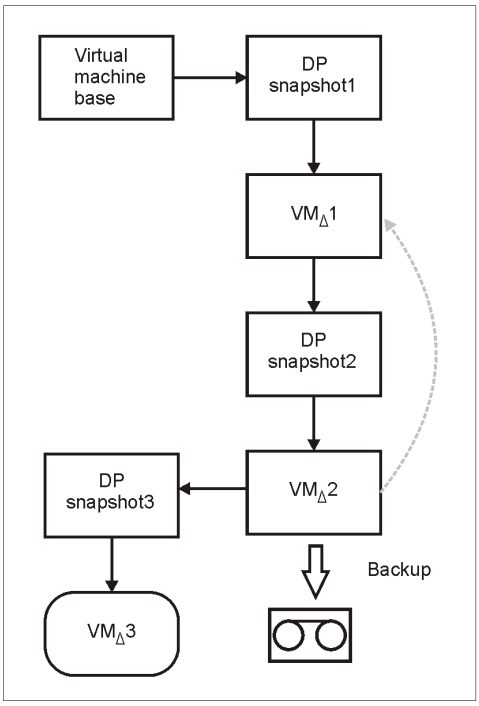

1. A new snapshot is created (DP snapshot3).
2. Changes made between DP snapshot2 and DP snapshot3 are backed up.
3. DP snapshot2 is removed ($VM_{\Delta}2$ is committed to the parent file $VM_{\Delta}1$).

Figure 14 Incremental backup (mixed mode)

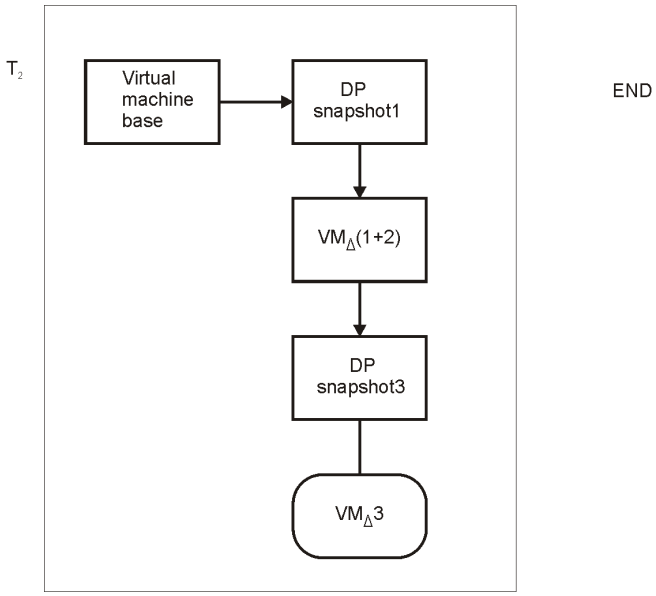

DP snapshot1 and DP snapshot3 remain on the snapshot tree to track changes made since the last Full and the last backup respectively.

A subsequent Differential backup progresses in the same way as an Incremental session described in "Incremental backup (mixed mode)" (page 38), 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 newer 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 8 Example 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 in 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 enough disk space on the datastore on which the virtual machine resides. You can use the Data Protector **Free space required %** option to ensure that a virtual machine is not backed up (the backup is skipped) unless the specified condition is fulfilled.

Before a virtual machine is backed up, Data Protector checks the datastore on which it is located. It sums the sizes of all virtual machines residing on this datastore, applies the specified percentage to the sum, and then checks whether the resulting size is still free on the datastore.

The check is performed before each virtual machine is backed up. If you specify 0%, the check is not performed.

Example

Suppose you have two virtual machines that reside on the same datastore. One has 20 GB of data and the other 30 GB. Suppose you specified 50% in the **Free space required %** option. Before the first virtual machine is backed up, Data Protector checks if there is at least 25 GB (50% of 20+30) of free space available. Suppose the first virtual machine size grows up to 30 GB during its backup and the second virtual machine's size remains the same. Before the second virtual machine is backed up, Data Protector checks if there is at least 30 GB (50% of 30+30) of free space available on the datastore.

Table 9 Disk space requirements

Backup method	Required disk space on datastores	Explanation
vStorage Image	The sum of the sizes of all the virtual machine disks, plus: • The size of any quiescence zip files, if quiescence is specified. • The size of the virtual machine memory file, if specified. • The size of the virtual machine snapshots if they are present (this applies to Single and Mixed snapshot handling modes with CBT disabled).	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.

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 variable 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 variables, see the online Help index: “omnirc options”.

Backup considerations

- **Concurrent backup sessions:** Backup sessions that use the same devices cannot run in parallel.
- **Transportation modes:** You can use various transportation modes for backup. For details, see “Advanced virtual machine settings” (page 51).

- **Large disks:** Backup of virtual machines that have disks (`*.vmdk`) larger than 1 TB minus 1 MB is not supported.

  Workaround: Install the Data Protector `Disk Agent` component inside such a virtual machine and perform a filesystem backup.

- **Thick and thin disks:** Data Protector is not able to 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.

Restore concepts

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

- 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, but you can change the destination datacenter if you want.

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 67).

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 in the datastore that need to be cleaned 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 72).
- **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.
- **Transportation modes:**
  - *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 74).
    - 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 LAN 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 not available for restore.

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

- Ensure that you have a correctly installed and configured VMware vSphere environment.

  For supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/manuals.

- Ensure 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*.

  Ensure 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.

  NOTE: The client that you intend to use as a 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 yourself).

- For 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 VMware vSphere 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 host) in your environment.

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), but to make them Data Protector clients, they must be imported into the Data Protector cell.

NOTE: If you want to install a Data Protector component on a VMware client (for example, you want to be able to perform a filesystem backup on the client), install the Data Protector component first and then import the client as a VMware vCenter or VMware ESX(i) client. Otherwise, you will have to re-import the client after the component has been installed.

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 appropriate client type (**VMware ESX(i)** or **VMware vCenter**) from the **Type** drop-down list and click **Next**.

Figure 15 Importing a VMware client (Name and Type)

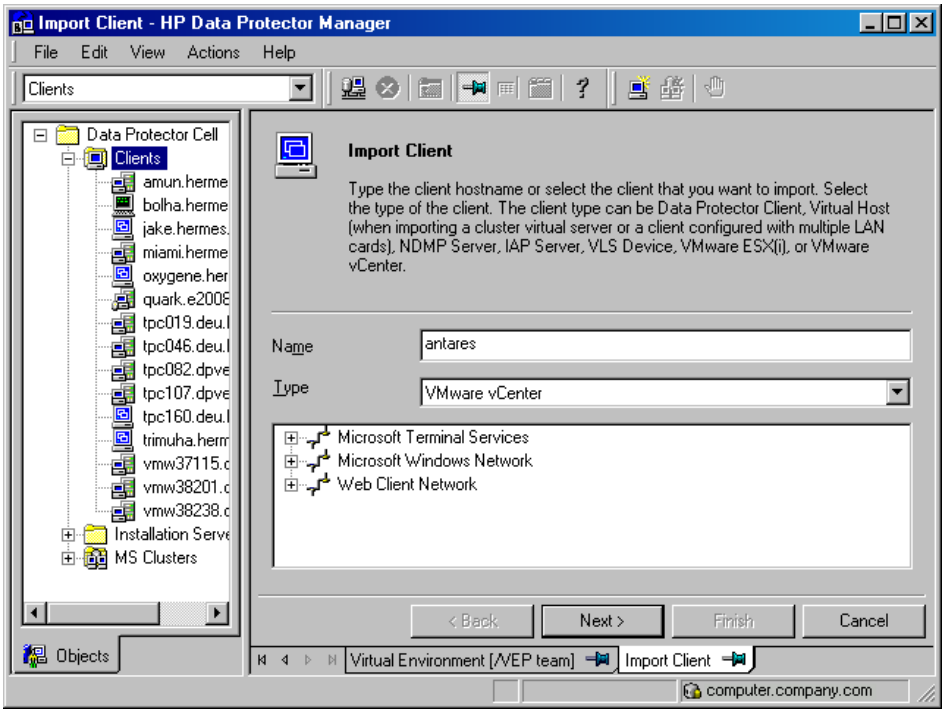

4. If you select **Standard security**, you need to manually specify which login credentials Data Protector should use to connect to the VMware client:

   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 online 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).

   Click **Finish**.

Figure 16 Importing a VMware client (login credentials)

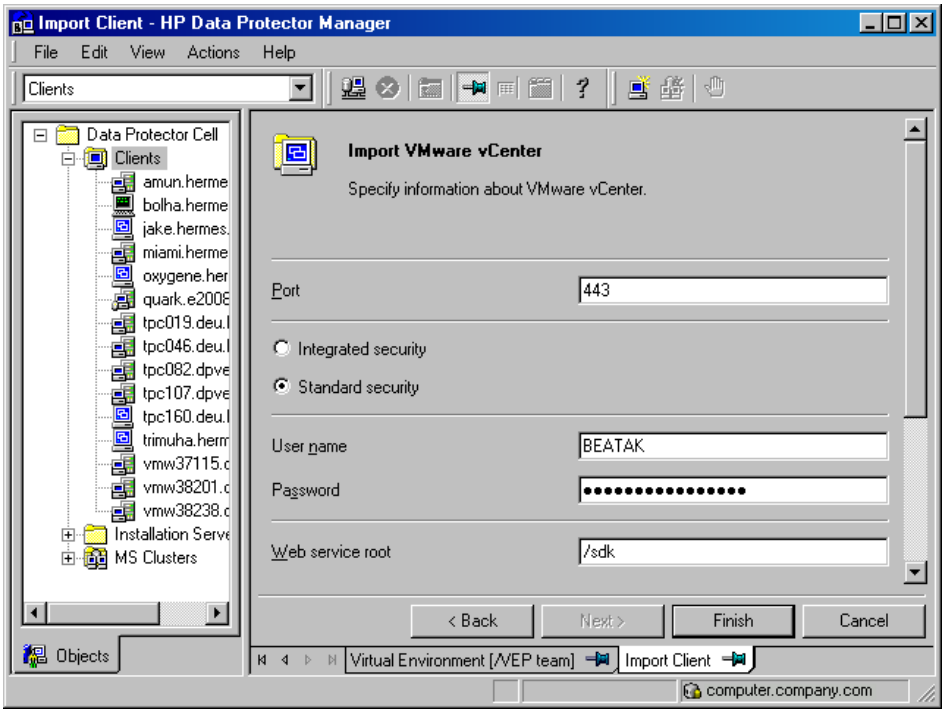

NOTE: For details on how to change or check the parameters later, see "Changing the configuration of VMware clients" (page 46) and "Checking the configuration of VMware clients" (page 48).

Changing the configuration of VMware clients

When you update the credentials for connecting to a VMware client (vCenter Server or ESX(i) Server 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 OLH 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 click **Configure**.

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

Figure 17 Changing the configuration of a VMware client

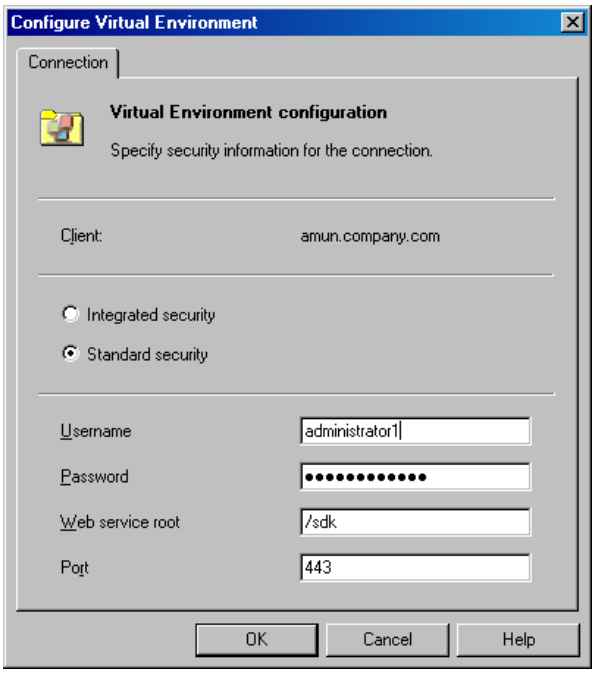

Using the Data Protector CLI

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.

   For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

2. Run:

   For **Integrated security**:

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

   For **Standard security**:

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

   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 click **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.

   For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

2. Run:

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

   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 click **Configure Virtual Machines**.

Figure 18 Configure virtual machines (Settings)

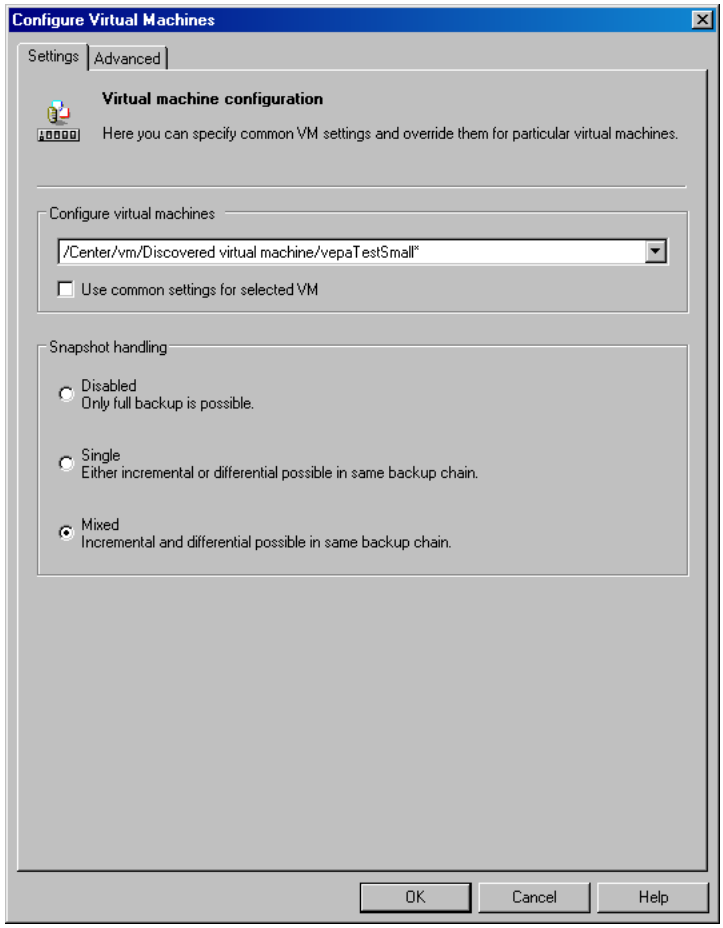

In the Configure Virtual Machines dialog box **Settings** page, specify the following settings:

Table 10 Virtual machine settings

Configure virtual machine	Select whether you want to specify common virtual machine settings (**Common VM Settings**) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.	
	Use common settings for selected VM	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
	Use default settings	Available only if **Common VM Settings** is selected. Select this option to set default values for the common virtual machine settings. Default: selected
Snapshot handling	**Disabled** (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 "Snapshot mode: disabled" (page 33).
	Single	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 "Snapshot mode: single" (page 34).
	Mixed	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 "Snapshot mode: mixed" (page 36).

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

Figure 19 Configure virtual machines (Advanced)

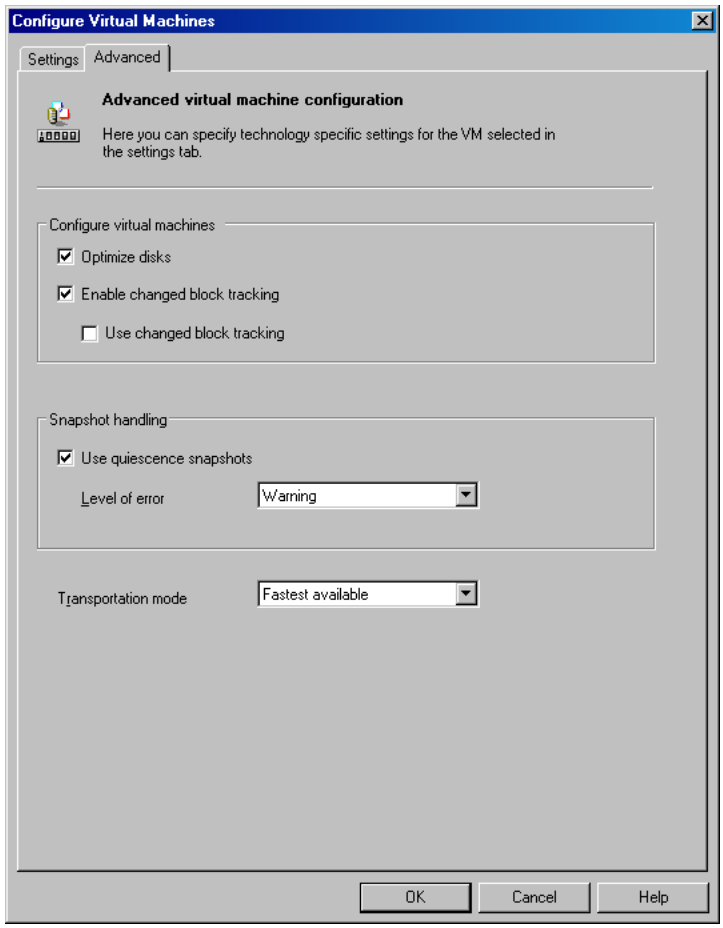

Table 11 Advanced virtual machine settings

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.	
	Optimize disks	This option is applicable only for thin virtual machine disks. 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. 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. Default: not selected.
	Enable changed block tracking	Select this option to enable the VMware Changed Block Tracking functionality for the selected virtual machines. If the functionality has already been enabled using the VMware vSphere client, you cannot disable it by clearing this option. Default: not selected.
	Use changed block tracking	Available only if **Enable changed block tracking** is selected. Select this option to use the VMware Changed Block Tracking functionality for Incremental and Differential backup sessions. For details, see "Changed Block Tracking" (page 38). 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.

Table 11 Advanced virtual machine settings *(continued)*

		Default: not selected.
Snapshot handling	**Use quiescence snapshots**	Applicable for Windows virtual machines. 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. Default: not selected. For details, see "Quiescence" (page 39).
	Level of error	Available only if **Use quiescence snapshots** is selected. Specify the level of error to be reported if a quiescence snapshot fails.
Transportation mode		Select the transportation mode to be used when backing up a virtual machine. • **LAN**: 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. • **LAN (SSL)**: Same as LAN except that the communication over the network is encrypted using the Secure Socket Layer (SSL) cryptographic protocol. • **Hotadd**: 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. • **SAN**: 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. For details, see the *VMware SAN configuration guide*. This transportation mode requires that the storage volumes on which the virtual machines are located are presented to the client with the `Virtual Environment Integration` component installed (**backup host**). **Caution:** Do not reformat these storage volumes. Otherwise, you will delete all your virtual machines. For details on these VMware transportation modes, see the *VMware Virtual Machine Backup Guide*. If you are not concerned which mode is used, select **Fastest available**. Default: **Fastest available**

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.

   For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

2. Run:

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

VM_OPTIONS
  --snapshots { 0 | 1 | 2}
  --transportation–mode {san | lan | lanssl | 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 `util_vmware.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, run:

```
vepa_util.exe
  command
  --configvm
  --virtual-environment vmware
  --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.

run:

```
util_vmware.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` variables

The `omnirc` variables 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 variables, see the online 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 28).

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 click **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 or VMware ESX(i) clients. They have a corresponding label appended at the end of their names, such as ***(VMware vCenter)***, or ***(VMware ESX(i))***.

     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 ***(LegacyVMware)*** label appended at the end. If you want 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 77).

     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 43).
   - 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**, select a datacenter 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.

- In **Backup method**, the **vStorage Image** backup method is displayed.
- In **Free space required [%]**, specify the percentage of disk space that should be free in a datastore before a virtual machine is backed up. For example, if a virtual machine's size is 50 GB and you select 50%, the virtual machine is backed up only if you have at least 25 GB disk space available in the datastore. Otherwise, backup is skipped.

  The check is performed separately for each virtual machine.

Figure 20 Selecting a VMware client, backup host, and datacenter

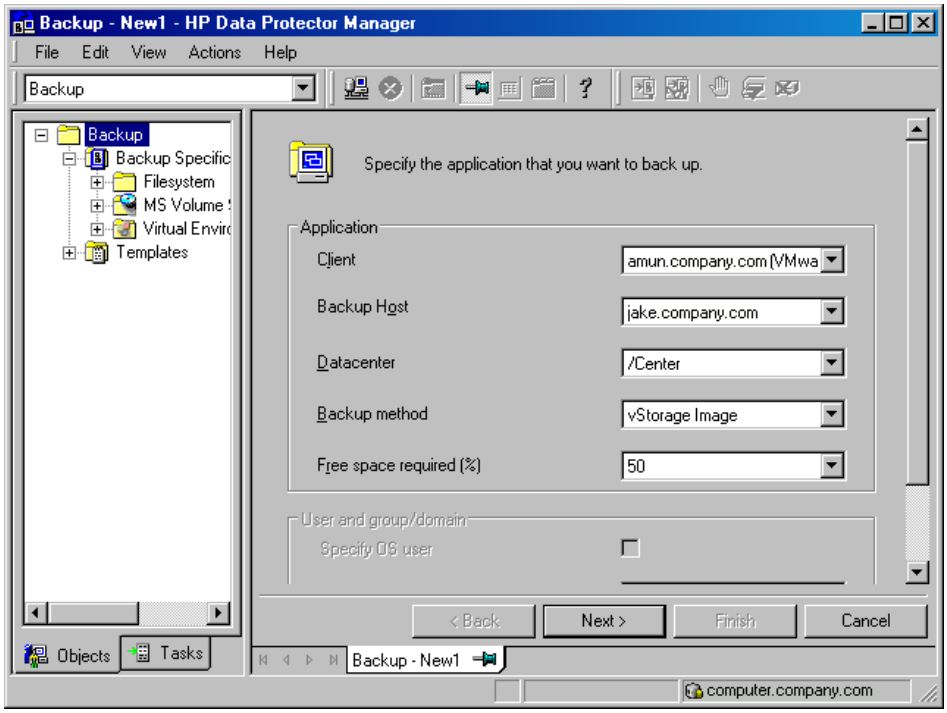

Click **Next**.

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 virtual machines, individual virtual machine disks, or virtual machine templates that you want to back up.

   You can make your selections at the following levels:

   - ESX/ESXi Servers systems
   - Pools
   - vApps
   - VM folders
   - Individual VMs
   - VM disks
   - VM templates

NOTE: If you select any level above individual VMs (for example, a vApp), 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 21 Selecting VMware objects

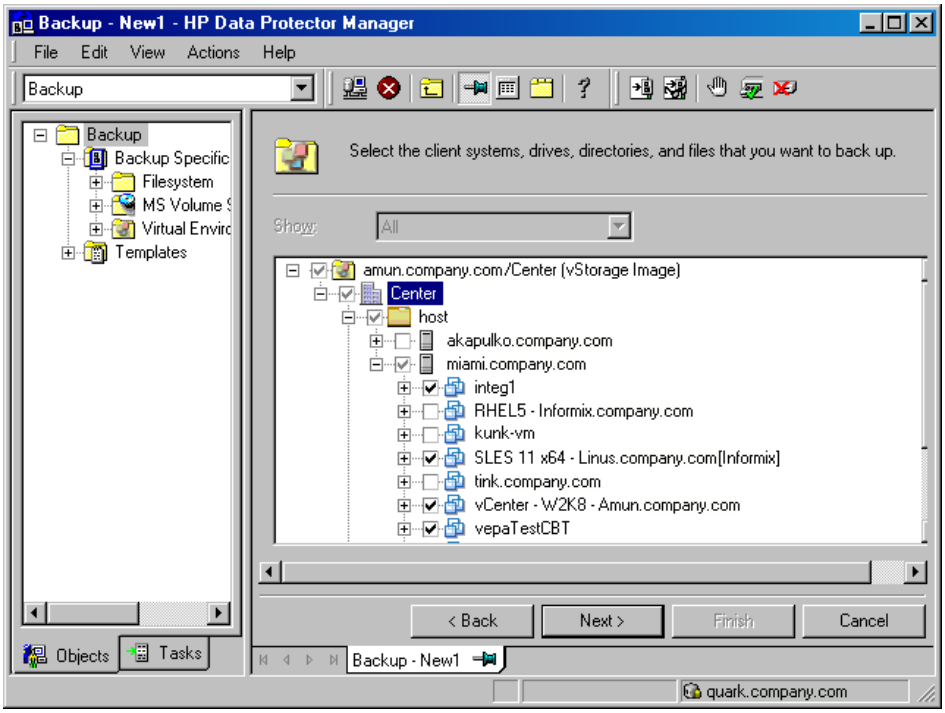

If your virtual machines are not configured yet, right-click the client system at the top or any of the virtual machines listed below, and click **Configure Virtual Machines**. For details, see "Configuring virtual machines" (page 48).

Click **Next**.

6. Select which 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 which media pool to use.

   Click **Next**.

7. Set backup options.

   For information on application-specific backup options, see "VMware backup options" (page 57).

Figure 22 Application-specific options

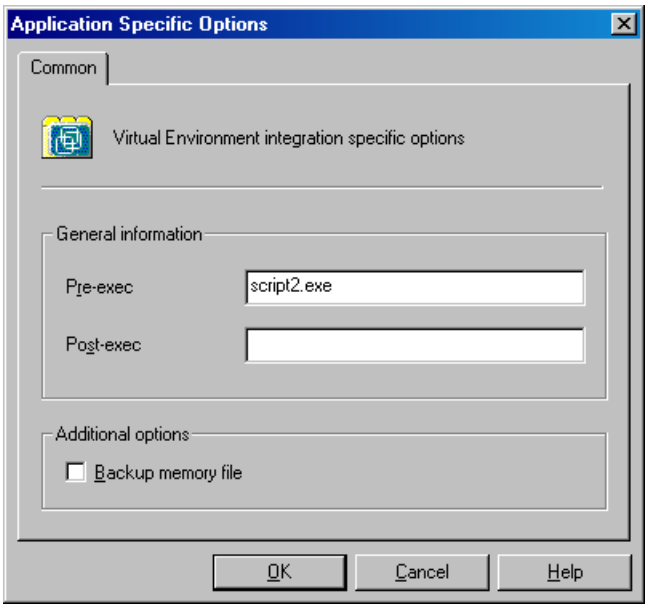

Click **Next**.

8. Optionally, schedule the backup. See "Scheduling backup specifications" (page 58).
   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 59).

Table 12 VMware backup options

Options	Description
Pre-exec, **Post-exec**	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: ***Windows Vista, Windows Server 2008:*** *`Data_Protector_program_data`*`\bin` ***Other Windows systems:*** *`Data_Protector_home`*`\bin` ***Linux:*** `/opt/omni/lbin`
Backup memory file	If this option is selected, the memory of a running virtual machine is saved into a file and backed up.

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.

Figure 23 Modifying a backup specification

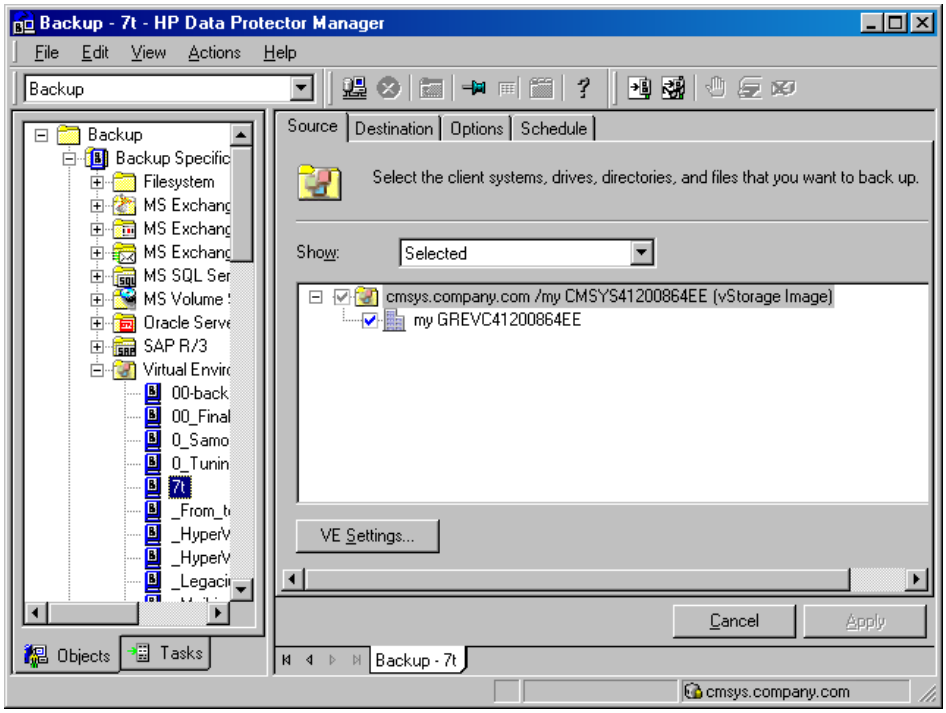

To display Virtual Environment settings, click the **VE Settings** button. Some of the settings can be modified.

Figure 24 Virtual Environment Settings

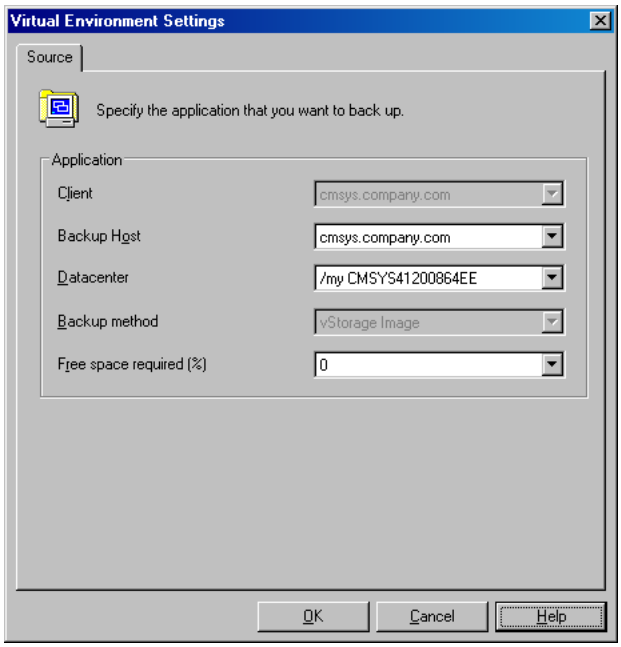

NOTE: To see all available backup objects in the source page, not just those you selected, select **All** in the **Show** option.

Scheduling backup specifications

You can run unattended backups at specific times or periodically. For details on scheduling, see the online 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 specification" (page 59). 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 25 Scheduling a backup specification

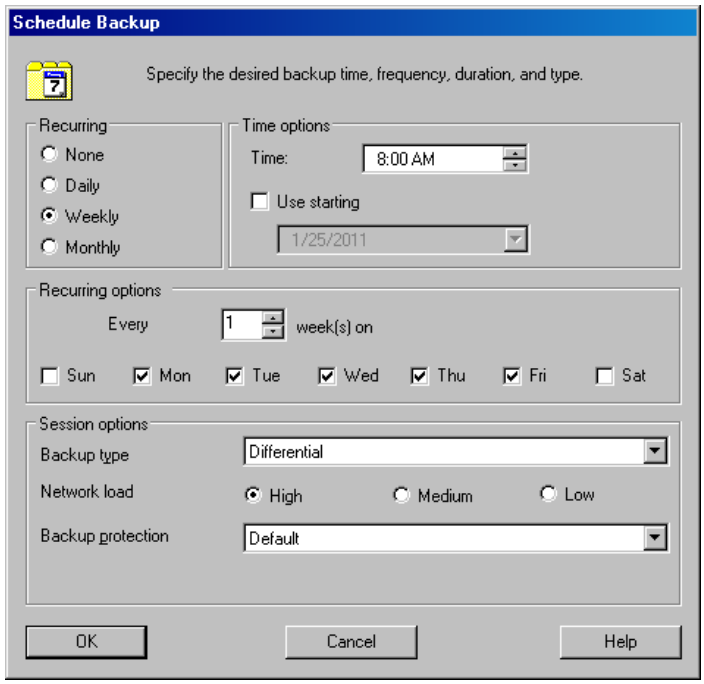

Previewing backup sessions

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 click **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.

   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:

   ```
   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 click **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.

   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:

   ```
   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`, run:

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

To start a Differential backup using the same backup specification, run:

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

Preparing for disaster recovery

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

Table 13 What must be backed up

VMware objects	How to back up
ESX/ESXi Server console	*ESX Server systems:* 1. Ensure that the Data Protector `Disk Agent` component is installed on all the ESX Server systems. 2. In the Backup context of the Data Protector GUI, right-click **Filesystem** and click **Add backup** 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 http://kb.vmware.com/selfservice/microsites/microsite.do. 3. Start a backup using the newly created backup specification. *ESXi Server systems:* With ESXi Server systems, it is not possible to install a Data Protector `Disk Agent`, so you will need to use VMware utilities to back up the configuration. A tool, `esxcfg-cfgbackup` is available from VMware. For information, refer to the VMware web site.
vCenter configuration database *(applicable only for VirtualCenter environments)*	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: 1. Ensure that the Data Protector `Oracle Integration` component is installed on the vCenter Server system. 2. In the Backup context of the Data Protector GUI, right-click **Oracle Server** and click **Add backup** to create a backup specification of the Oracle type. In **Application database**, type the name of the vCenter configuration database. Continue with the backup specification creation as described in the *HP Data Protector Integration Guide for Oracle and SAP*. 3. Start a backup using the newly created backup specification.
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 41).

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`.

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 26 Backup object information

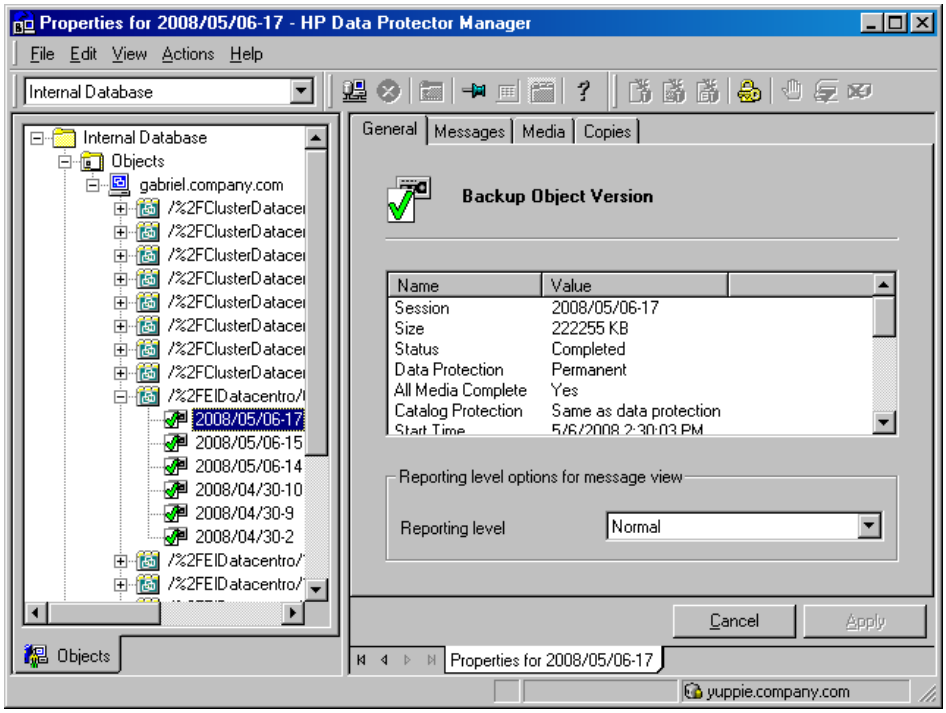

TIP: To view the messages displayed during the session, click the **Messages** tab.

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.

   For the command location, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.
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::/%2FElDatacentro/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 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 27 Selecting VMware objects for restore

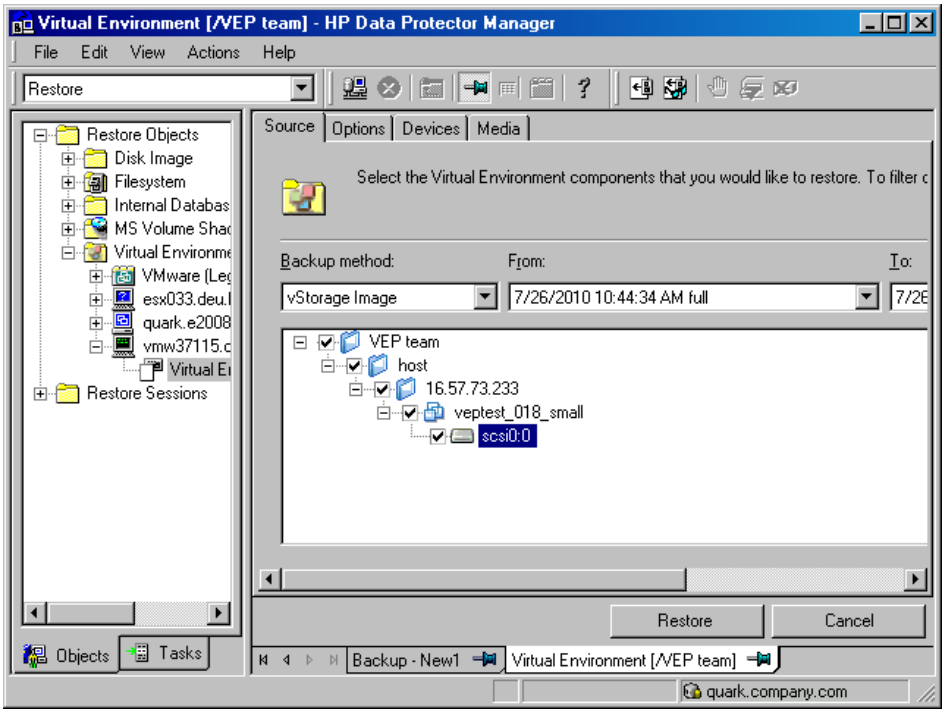

4. In the **Options** page, specify the VMware restore options. For details, see "Restore options" (page 64).

Figure 28 Restore options

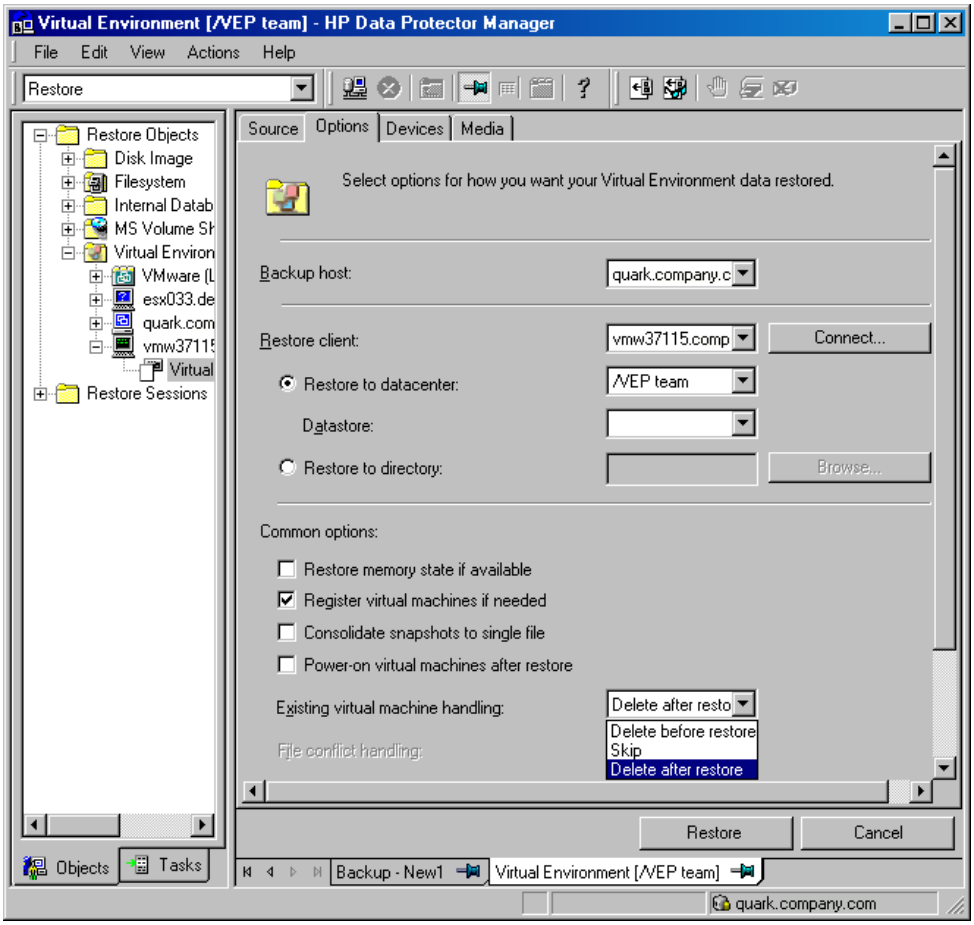

5. In the **Devices** page, select devices to use for restore.
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. Specify **Report level** and **Network load**.

   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 72).

Table 14 Restore options

GUI/CLI option	Description
Backup host/ `-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.
Restore client/ `-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 **Connect** button.
Datastore/ `-store`	Specifies the datastore to which the virtual machines should be restored. You can choose among all datastores that are accessible by the selected restore client. If you leave this option empty, the virtual machines are restored to the original datastore.

Table 14 Restore options *(continued)*

GUI/CLI option		Description
Restore to datacenter/ `-instance, -newinstance`		Select this option to restore a virtual machine to a datacenter. By default, virtual machines are restored to the original datacenter.
Restore to directory/ `-directory`		Select this option to restore virtual-machine files to a directory (outside of the datacenter) on the backup host. You can use the **Browse** button to find the target directory. 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 "Recovering virtual machines after restore to a directory" (page 67).
Restore memory state if available/ `-memory`		Available if **Restore to datacenter** is selected. Select this option to restore the memory file if it is included in the backup.
Register virtual machines if needed/ `-register`		Available if **Restore to datacenter** 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 "Recovering virtual machines manually" (page 67). Default: selected.
Consolidate snapshots to single file/ `-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 **Restore to datacenter** is selected.
Power-on virtual machines after restore/ `-poweron`		Select this option to power the virtual machines on once they are restored. Available if **Restore to datacenter** is selected.
Existing virtual machine handling	Specifies Data Protector's behavior when restoring existing virtual machines.	
	Delete before restore/ `-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.
	Skip restore/ `-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.
	Delete after restore	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. Default: selected.
File conflict handling	Specifies Data Protector's behavior when restoring existing files.	
	Overwrite/ `-overwrite`	Select this option to overwrite existing files with those from the backup. Default: selected.
	Keep latest/ `-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.
	Skip/ `-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.

   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 vmware
 -barhost BackupHost
 -apphost OriginalVMwareClient
 –instance OriginalDatacenter
 –method vstorageimage
 [-session BackupID]
 VirtualMachine [VirtualMachine...]
 [VMwareClient | Directory]

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

 VMwareClient
 [–newinstance TargetDatacenter]
 [-store TargetDatastore]
 [–destination TargetVMwareClient]
 [-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*.

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.
- If the restore fails, see "Cleaning up a datastore after a failed restore" (page 72).

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, run:

```
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`, run:

```
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 67).
- If you have restored the virtual machines to a datacenter (**Restore to datacenter**), but without selecting the restore option **Register virtual machines if needed**.

  For details, see “Recovering virtual machines after restore to a datacenter” (page 71).

Recovering virtual machines after restore to a directory

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 the **Convert Machine** button 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 29 Conversion (Source System)

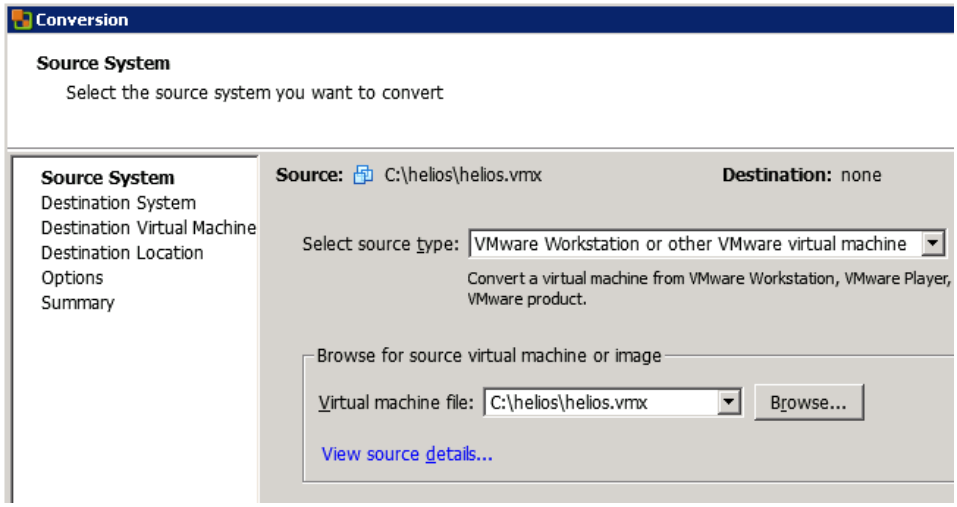

NOTE: In our example, the VMware Converter is installed on the backup host.

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 30 Conversion (Destination System)

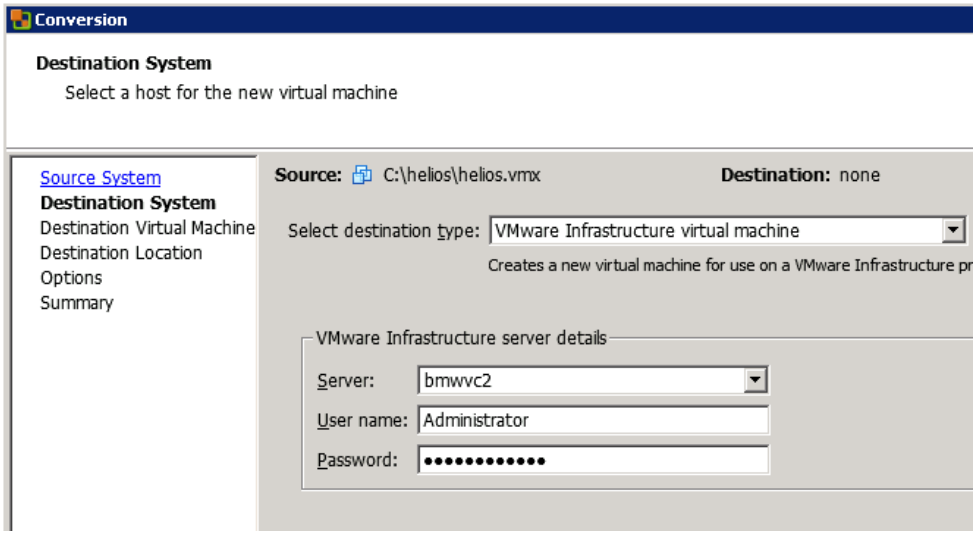

7. In the **Destination Virtual Machine** page, specify the name under which the virtual machine should be recovered.

   Click **Next**.

Figure 31 Conversion (Destination Virtual Machine)

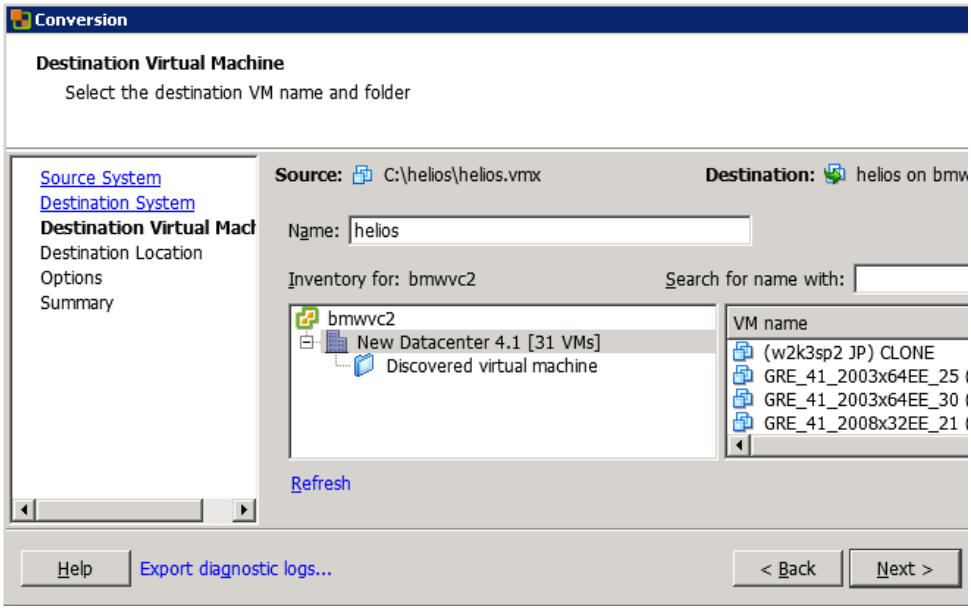

8. In the **Destination Location** page, select the destination ESX(i) Server system and datastore.

Figure 32 Conversion (Destination Location)

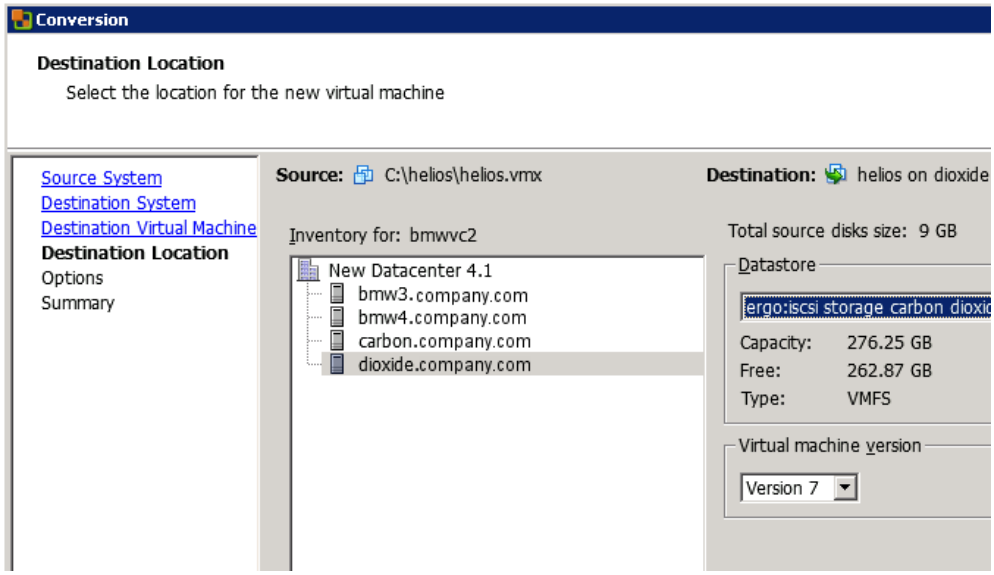

9. In the **Options** page, edit the options and click **Next**.

Figure 33 Conversion (Options)

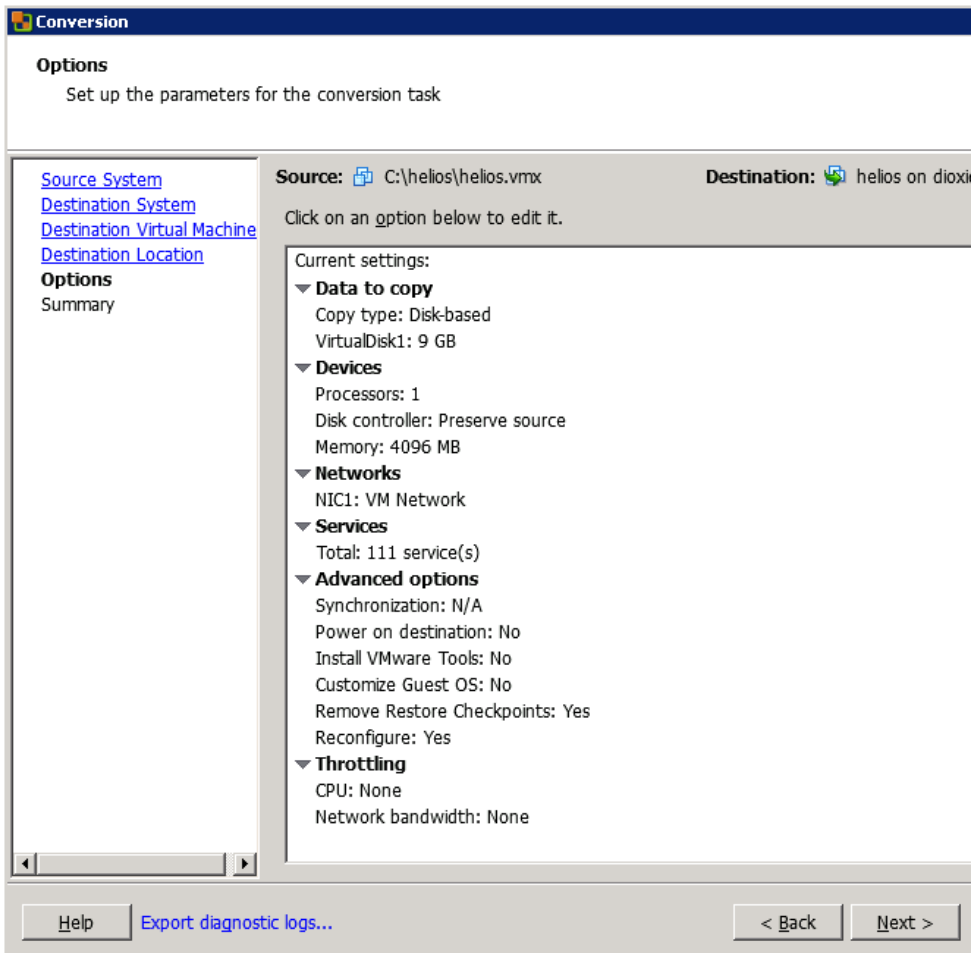

10. In the **Summary** page, review your selection and click **Finish**.

Figure 34 Conversion (Summary)

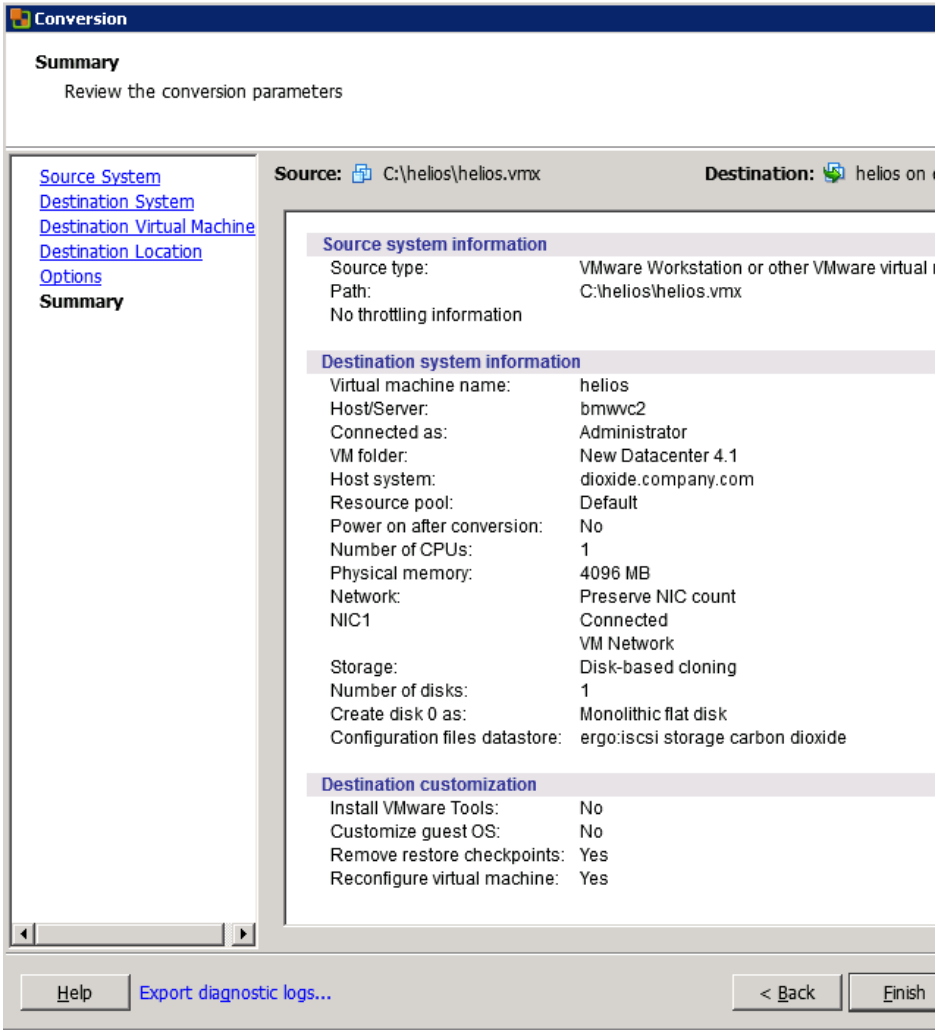

11. Open the **Datastore Browser** and upload the files created in the Incremental and Differential backup sessions to the virtual machine directory:

```
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.

NOTE: The procedure described in this section cannot be used if the virtual machine `*.vmx` file has not been restored (and an `*.xml` file has been restored instead). In such cases, to recover the virtual machine, import the virtual machine disks `*.vmdk` into an existing virtual machine.

Recovering virtual machines after restore to a datacenter

If you have restored a virtual machine to a datacenter, but 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 `*.vmx` file, and click **Add to Inventory**.
3. Follow the wizard and click **Finish**.

Restoring using another device

You can restore using a device other than that used for backup. For details, see the online 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 in 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 click **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 click **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 online 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 online 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 online 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://www.hp.com/support/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` located in `Data_Protector_home\log` on the backup host.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the online 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

Backup fails

A backup session fails with an error similar to the following:

```
[Major] From: OB2BAR_VEPA_BAR@vmw37239.company.com "/ESX5.0_Testing"  Time: 8/5/2011 10:03:15 AM

        Virtual Machine 'vmw35086_VM_Judith': Could not backup disk scsi0:0 ...

[Major] From: OB2BAR_VEPA_BAR@vmw37239.company.com "/ESX5.0_Testing"  Time: 8/5/2011 10:03:15 AM

        Virtual Machine 'vmw35086_VM_Judith': No disk backed up ...
```

This may happen if the datastore names contain non-alphanumeric characters (for example, @). The Virtual Disk Development Kit does not support such characters. For details, see:

http://www.vmware.com/support/developer/vddk/VDDK-1.2.1-Relnotes.html

Action

1. Ensure that the datastore names contain only letters, numbers, single quotes, spaces, and hyphens.
2. Restart the session.

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.

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

Reconfigure the virtual machine:

1. Open the backup specification.
2. In the Source page, right-click the VMware client and click **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
\a1f9f4e3-482d-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. Run `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 onlined 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.

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[2] systems (**ESX(i) Server systems**), and VirtualCenter Server systems, to back up and restore the following **VMware objects**:

- Virtual machines
- Filesystems of virtual machines[3]

The following backup methods are available:

- Snapshot
- Suspend
- VCBimage[4]
- VCBfile [4]

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 83).

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

- Full
- Incremental
- Differential

For details on the backup types see, "Backup types" (page 93).

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 online 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 80).
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://www.hp.com/support/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 80).

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 35 VirtualCenter environment

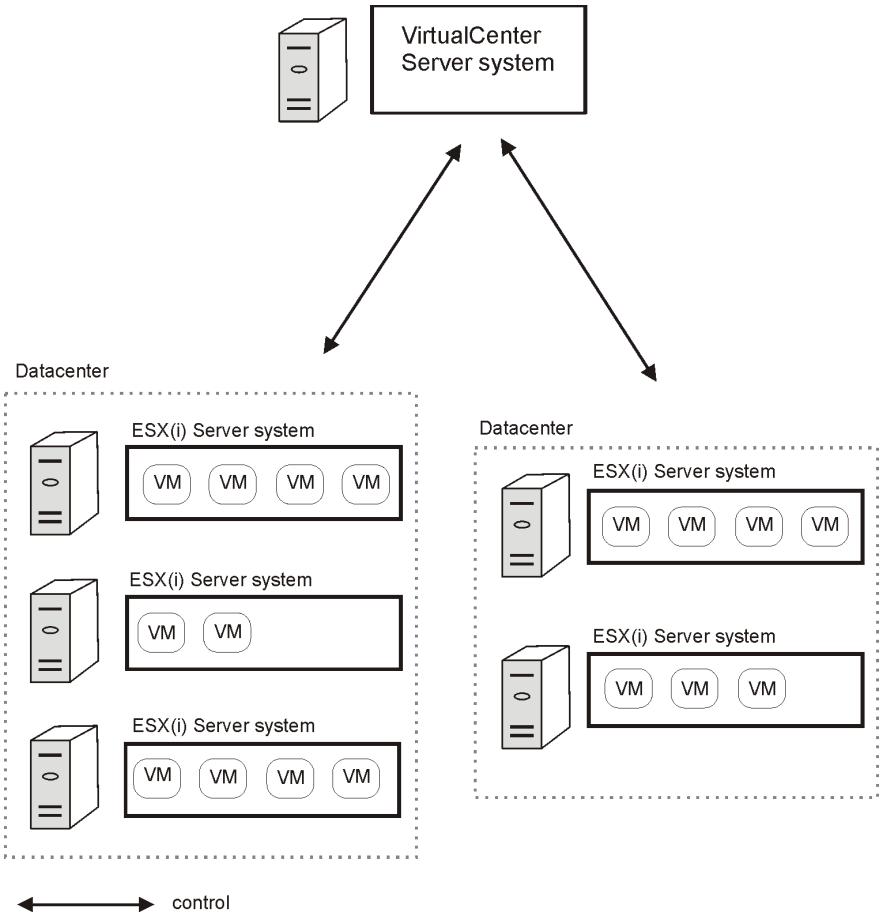

ESX(i) Server system	VMware platform capable of hosting multiple virtual machines.
VM	Virtual machine. Virtualized x86 or x64 PC environment, in which a guest operating system and associated application software can run.
Datacenter	An organizational unit that consists of one or more ESX(i) Server systems 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 36 Migration of virtual machines

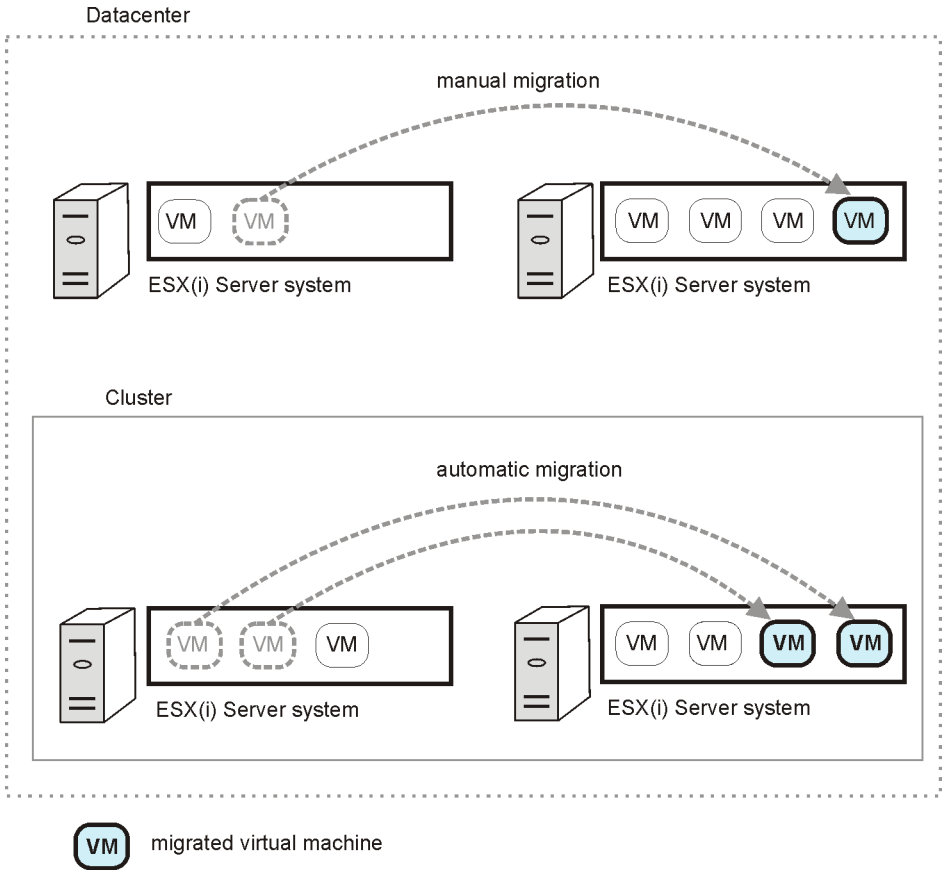

Whatever the reason for the migration, you do not need to create a new backup specification afterwards. Data Protector will automatically find the migrated virtual machines and back them up.

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 37 Standalone ESX Server environment

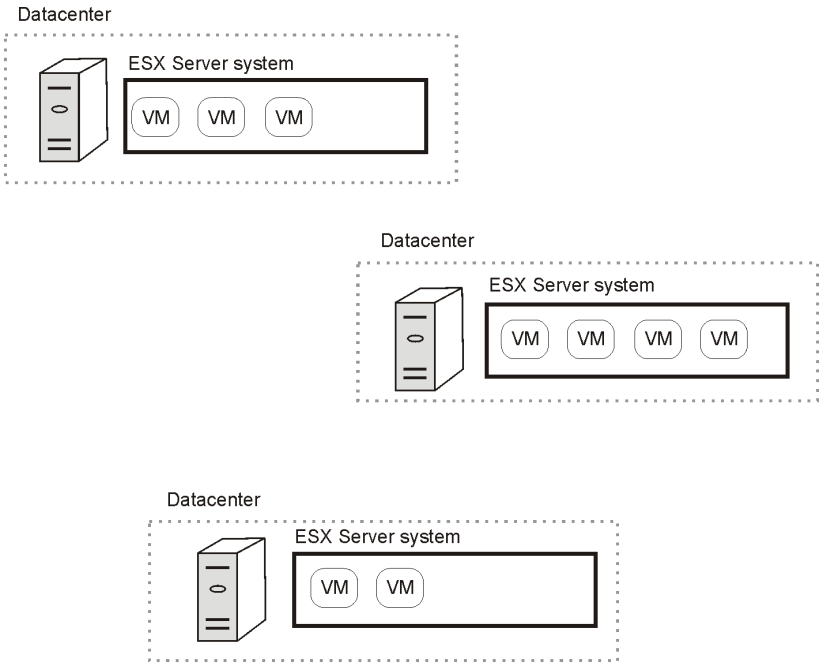

ESXi Server systems

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 78).

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 113) and "Restore options (VCBfile)" (page 114).

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.

`vmdk`
: 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.

`vmsn/vmss`
: Snapshot or suspend memory file (not applicable for the **VCBimage** backup method).

`vmsd`
: 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 `vmdk` 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 118).

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 83))
- **Suspend** (see "Suspend method" (page 89))
- **VCBimage** (see "VCBimage method" (page 90))
- **VCBfile** (see "VCBfile method" (page 91))

Table 15 Backup method overview

Backup method	How is backup consistency achieved?	Needs a backup proxy system	Supported backup types			Backs up all snapshot branches (including 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 93).

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 38 Snapshot tree

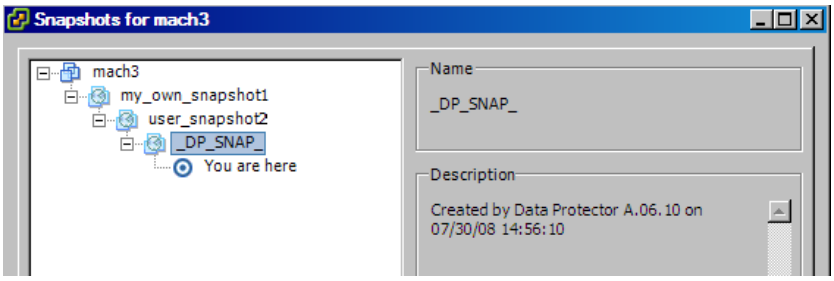

Existing virtual machine snapshots reduce the overall performance of a virtual machine. For this reason, Data Protector removes DP snapshots once they are no longer needed. The number of DP

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 84).
- **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 85).
- **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 87).

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 39 Full backup (disabled mode)

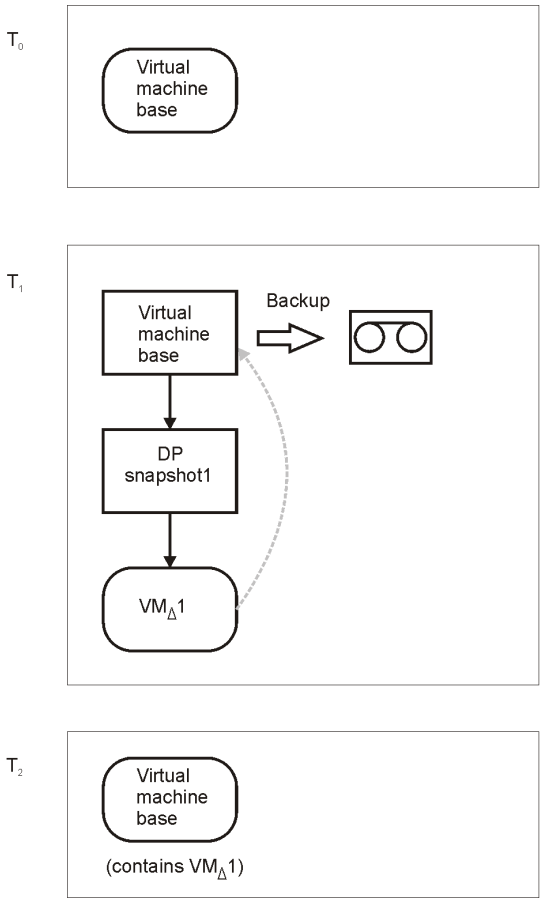

Table 16 Legend

Tn	Boxes denoted by Ti 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.
VMΔ	A rectangle denoted by VMΔ 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 40 Differential backup (single mode)

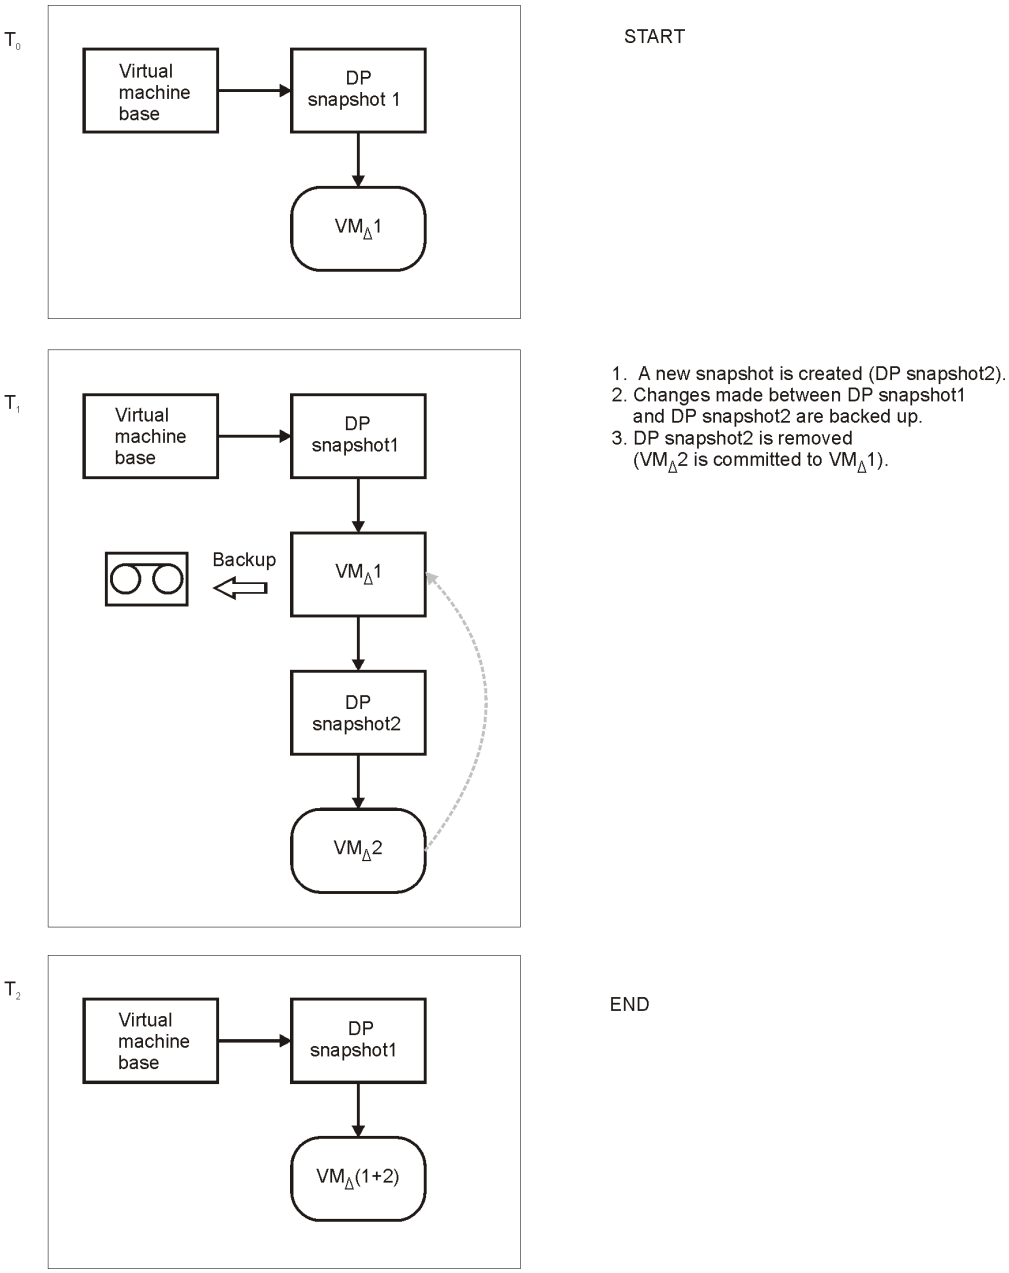

DP snapshot1 remains on the snapshot tree to track changes made since the last full backup.

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 41 (page 87)).

Figure 41 Incremental backup (single mode)

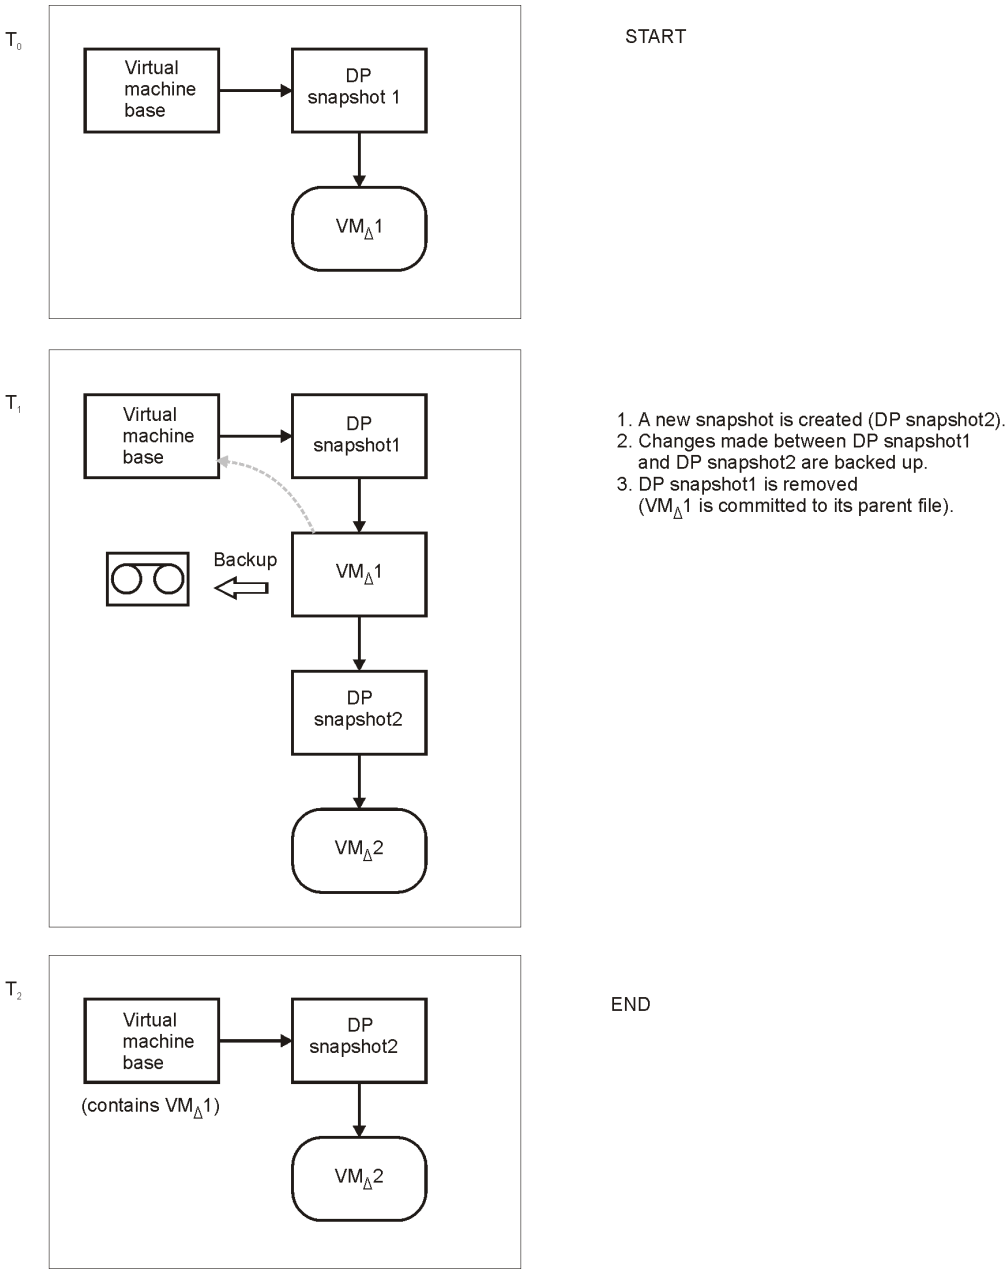

DP snapshot2 remains on the snapshot tree to track changes made since the last incremental backup.

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_0

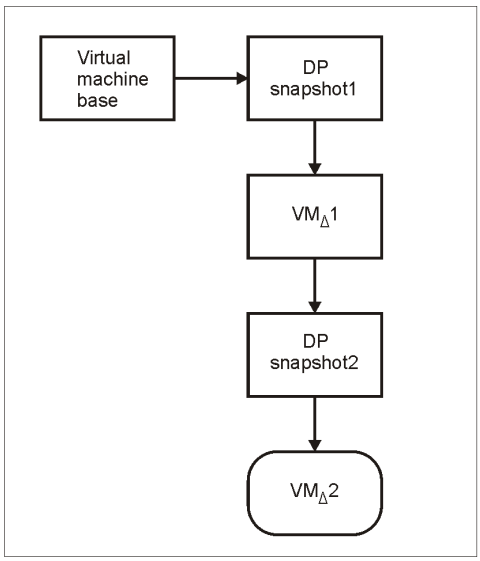

START

T_1

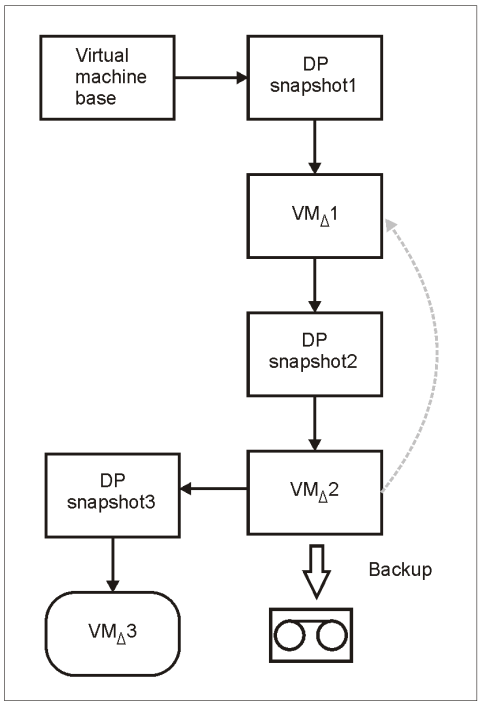

1. A new snapshot is created (DP snapshot3).
2. Changes made between DP snapshot2 and DP snapshot3 are backed up.
3. DP snapshot2 is removed (VM$_\Delta$2 is committed to the parent file VM$_\Delta$1).

Figure 42 Incremental backup (mixed mode)

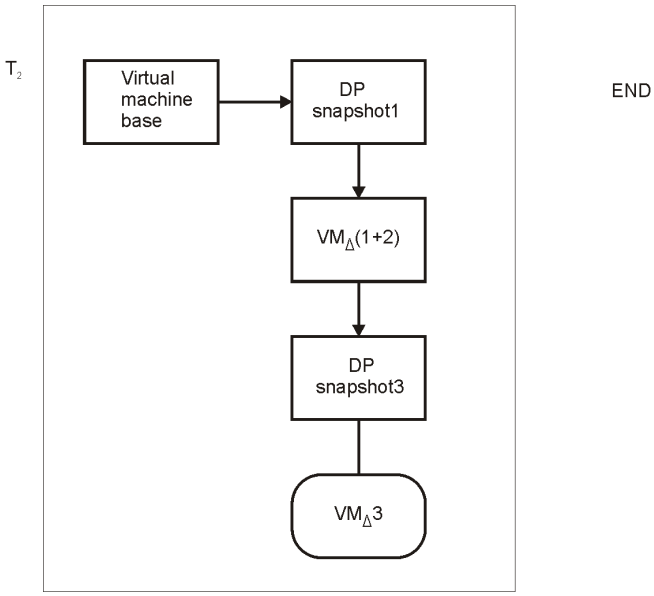

DP snapshot1 and DP snapshot3 remain on the snapshot tree to track changes made since the last full and the last backup respectively.

A subsequent differential backup progresses in the same way as an incremental session described in "Incremental backup (mixed mode)" (page 89) 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 43 Snapshot and Suspend backup methods

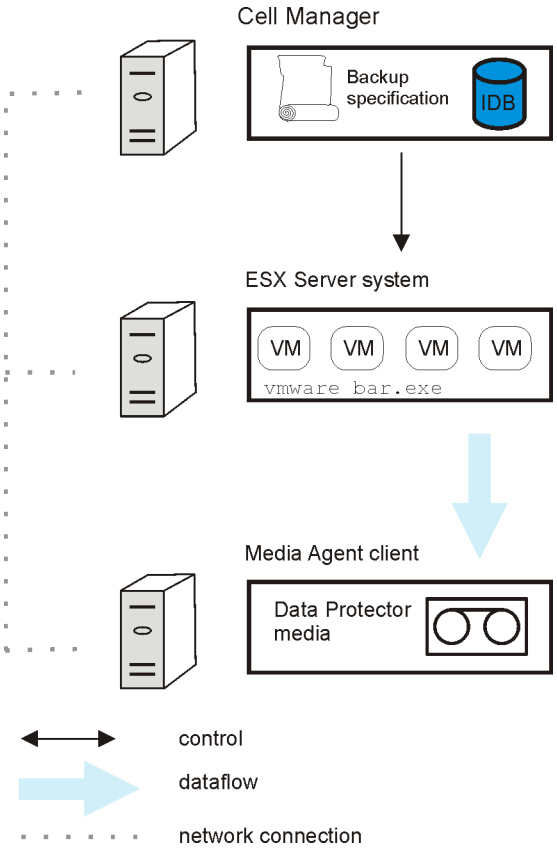

VCBimage method

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 options variable. For details of how to set the variable, see the online 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 unmounted. 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 44 VCBimage backup method

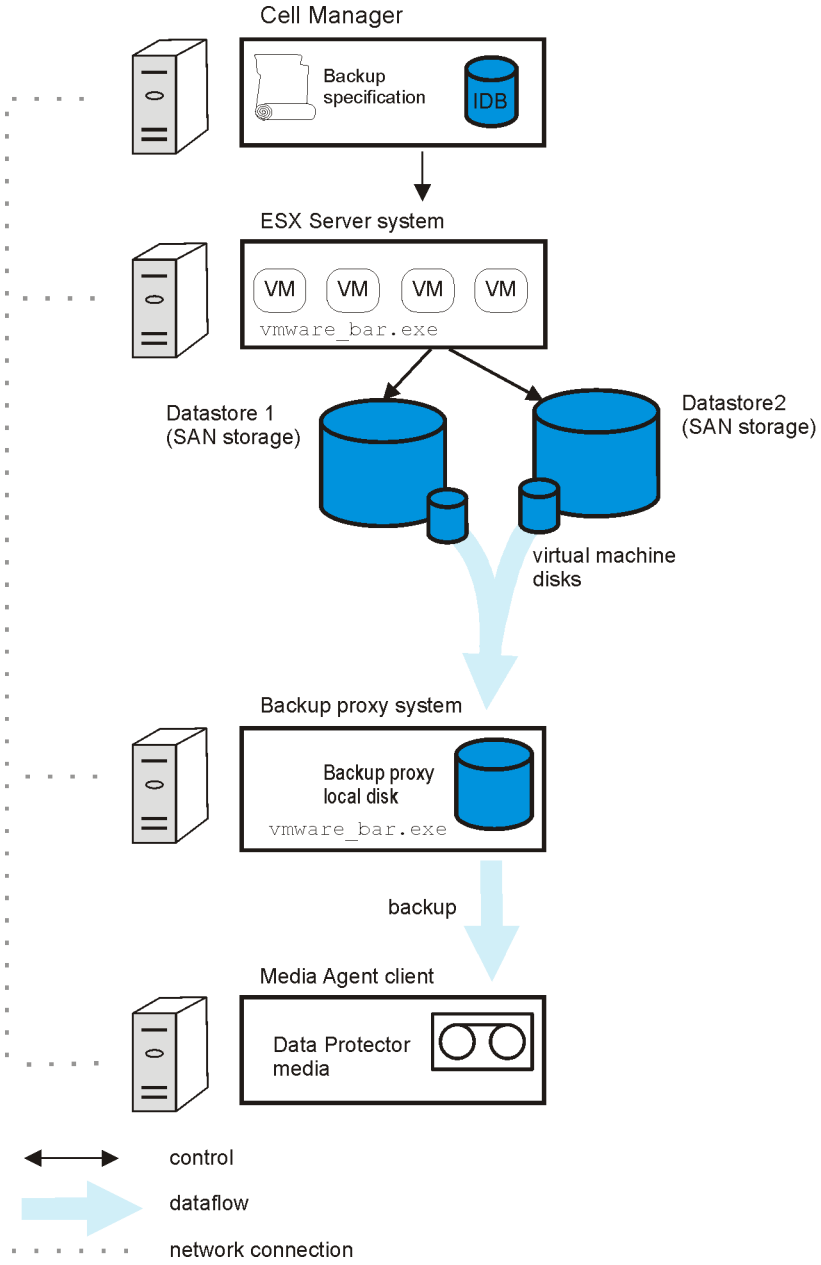

VCBfile method

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 unmounted 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 45 VCBfile method

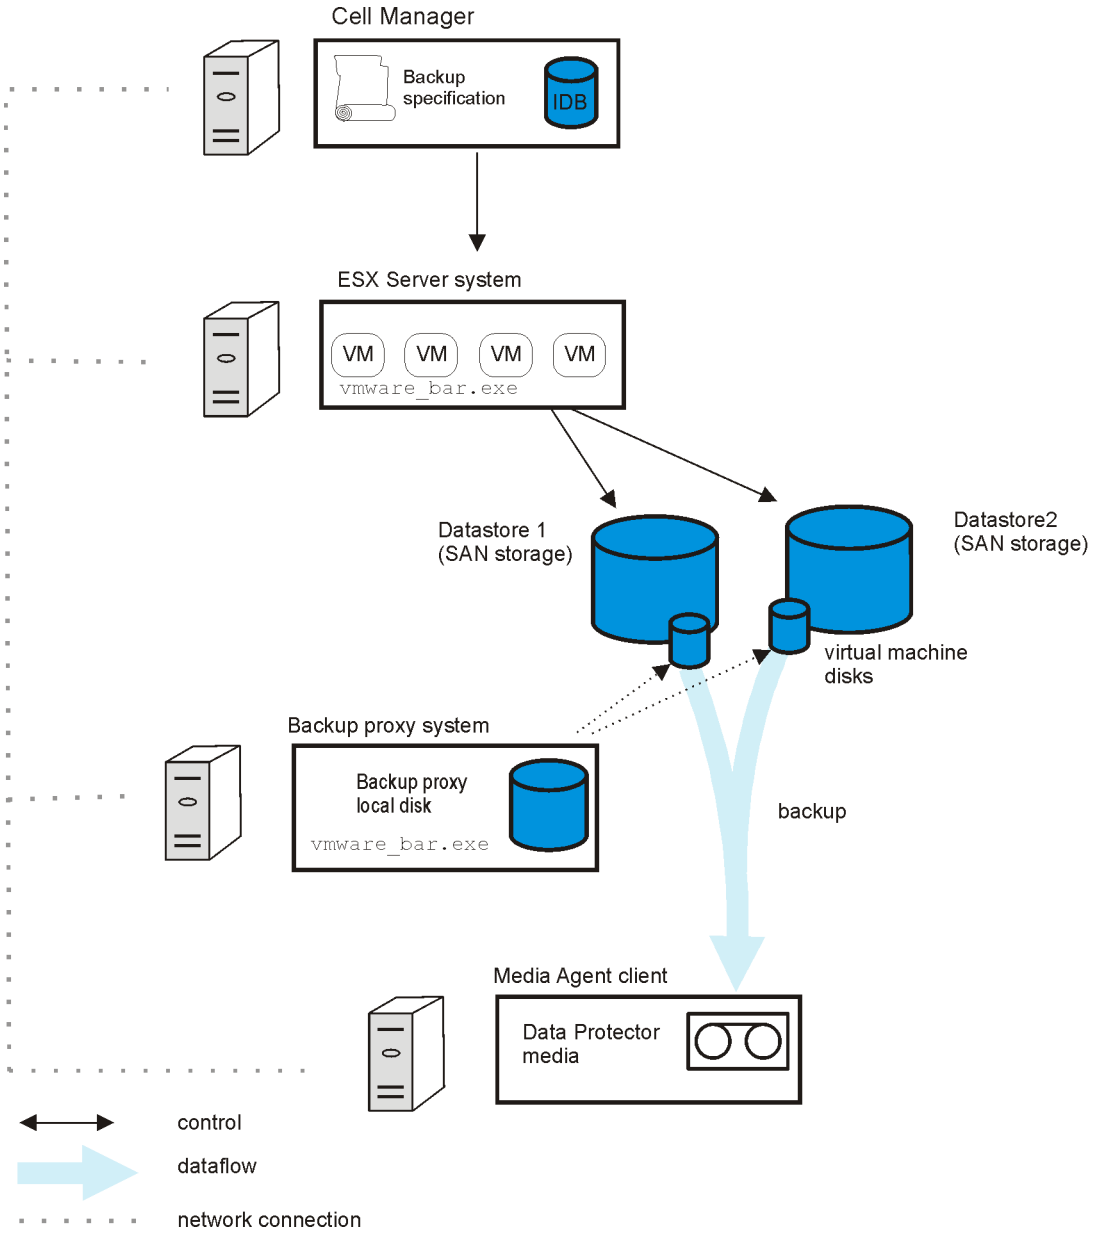

Backup types

The integration provides the following backup types:

Table 17 Backup types

Full	**Snapshot and Suspend:** Backs up the complete virtual machine, including the virtual machine snapshot tree and the virtual machine memory file (if specified). **VCBimage:** Backs up the complete virtual machine. The virtual machine snapshot tree and the virtual machine memory file are not included. **VCBfile:** Backs up all selected files and folders.
Incr	**Snapshot and Suspend:** Backs up changes made to the virtual machines since the last backup of any type. **VCBimage:** Not available. **VCBfile:** Backs up only those files and folders that have changed since the last backup of any type.
Differential	**Snapshot and Suspend:** Backs up changes made to virtual machines since the last full backup. **VCBimage:** Not available. **VCBfile:** 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 18 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` variable `OB2_VMWARE_SKIPSPACECHECK` on the backup proxy system to `1`.

For details of how to set the variable, see the online Help index: "omnirc options" and "Setting the environment variables using GUI" (page 94).

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` variable `OB2_VMWARE_VCBOPTIMISED` on the backup proxy system to `1`.

For details of how to set the variable, see the online Help index: "omnirc options" and "Setting the environment variables using GUI" (page 94).

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` variable `OB2_VMWARE_PATH` (for example, `/tmp/`). Therefore, before you start the restore session, ensure that the `omnirc` variable `OB2_VMWARE_PATH` is set. For details of how to set `omnirc` variables, see the online 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 118).

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 95)). However, if the `omnirc` variable `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` variables `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` variable to `1` and the `OB2_VMWARE_PATH` variable 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` variables, see the online 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` variables.

However, note that restoring a non-optimized backup session by setting these variables 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 97).
- Provide Data Protector with login information to VMware management clients as described in "Configuring VMware management clients" (page 97).
- If your ESX Server systems are configured in a cluster, check the cluster settings as described in "Configuring clusters" (page 97).
- For each virtual machine, specify details of how to perform various backup methods as described in "Configuring virtual machines" (page 100).

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://www.hp.com/support/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 19 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 online 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 constrains**.

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 online 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 46 Configuring a VirtualCenter Server system

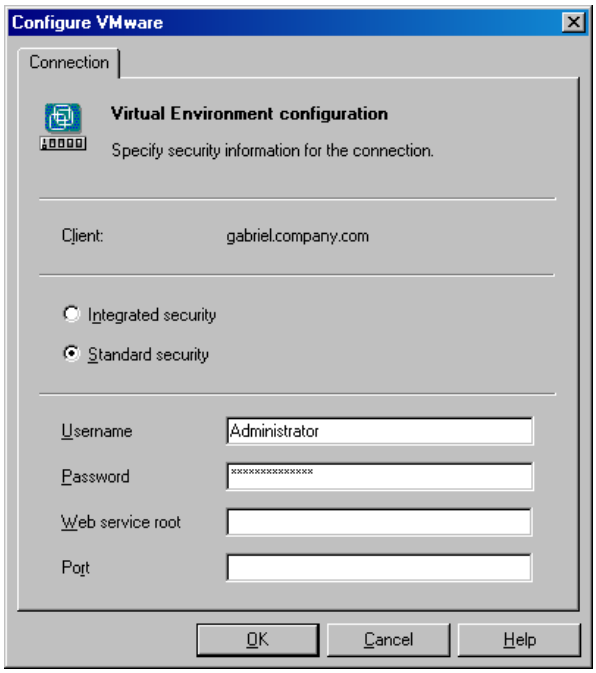

Select either **Integrated security** or **Standard security**.

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

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

  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 43)).

- **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` variable `OB2_VMWARE_HTTP` to `1`. For information of how to set the `omnirc` variable, see the online 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="vnmx46578"
```

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 97).
2. Go to the following directory:

   VirtualCenter Server system: *`Data_Protector_home`*`\bin`

Standalone ESX Server system: `/opt/omni/lbin`

3. Run:

   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 47 Configuring virtual machines

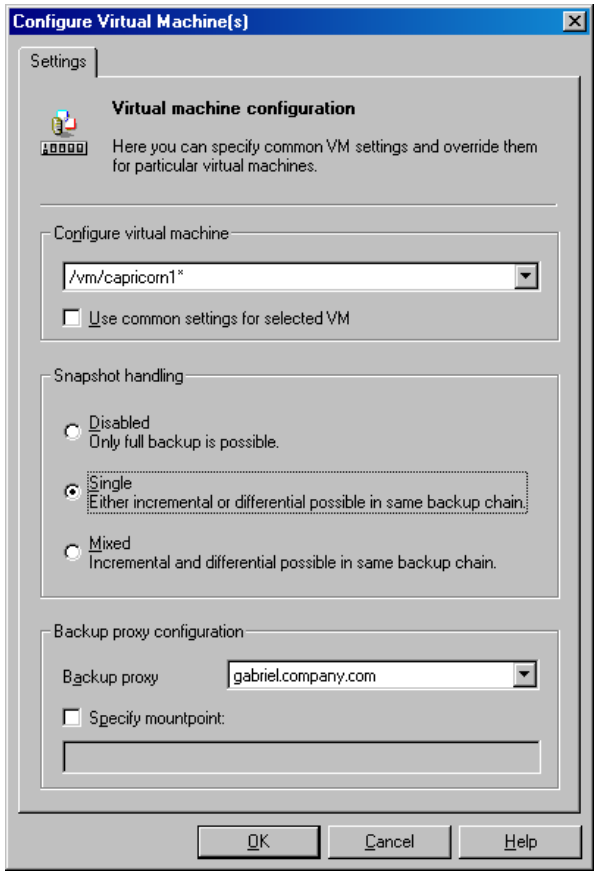

In the Configure Virtual Machines dialog box, specify the following options:

Table 20 Virtual machine options

Configure virtual machine	Select whether you want to specify common virtual machine settings (**Common VM Settings**) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.	
	Use common settings for selected VM	Available only if a virtual machine is selected. Select this option to apply the common virtual machine settings for the selected virtual machine. Default: ON
	Use default settings	Available only if **Common VM Settings** is selected. Select this option to set default values for the common virtual machine settings. Default: ON
Snapshot handling	**Disabled** (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 "Snapshot mode: disabled" (page 84).
	Single	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 "Snapshot mode: single" (page 85).
	Mixed	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 "Snapshot mode: mixed" (page 87).
Backup proxy configuration	**Backup proxy**	Select a backup proxy system to be used for the **VCBfile** and **VCBimage** 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.
	Specify mountpoint	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 **VCBimage** 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: *`Data_Protector_home`*`\tmp`

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 97).
2. Go to the following directory:

   VirtualCenter Server system: *`Data_Protector_home`*`\bin`

   Standalone ESX Server system: `/opt/omni/lbin`

3. Run:

```
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, run:

```
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`, go to the directory *`Data_Protector_home`*`\bin`, and run:

```
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 97).
2. Go to the following directory:

   VirtualCenter Server system: *`Data_Protector_home`*`\bin`

Standalone ESX Server system: `/opt/omni/lbin`

3. Run:

   ```
   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 81).

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 97).

   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 48 Selecting VMware objects (Snapshot, Suspend, VCBimage)

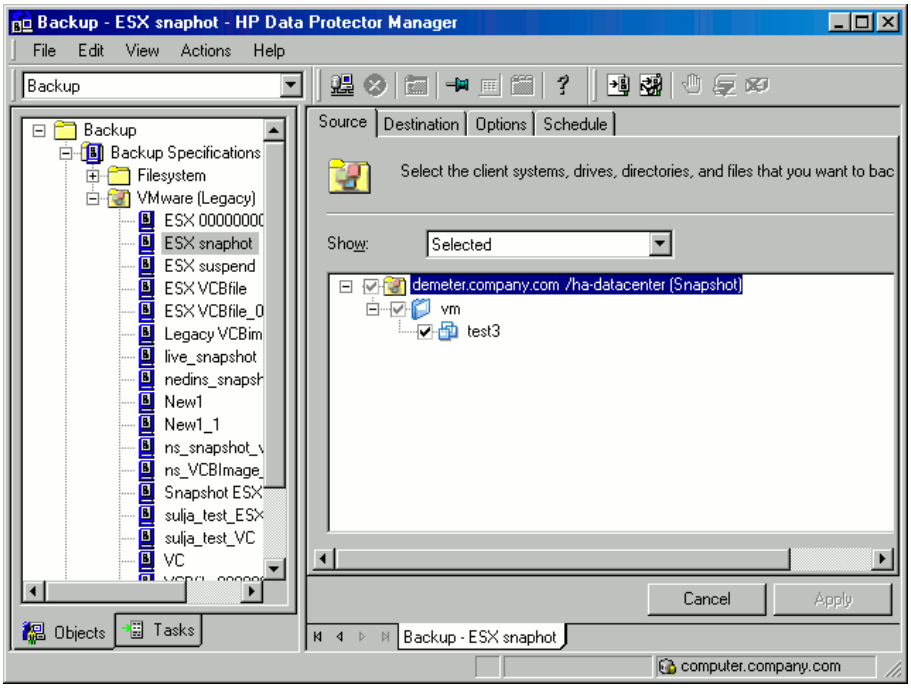

VCBimage backup method: To use the method in optimized mode, right-click the VMware datacenter, select **Set Environment Variables...**, and set the `OB2_VMWARE_VCBOPTIMISED` variable to `1`.

For details, see "Setting the environment variables using GUI" (page 94).

Figure 49 Selecting VMware objects (VCBfile)

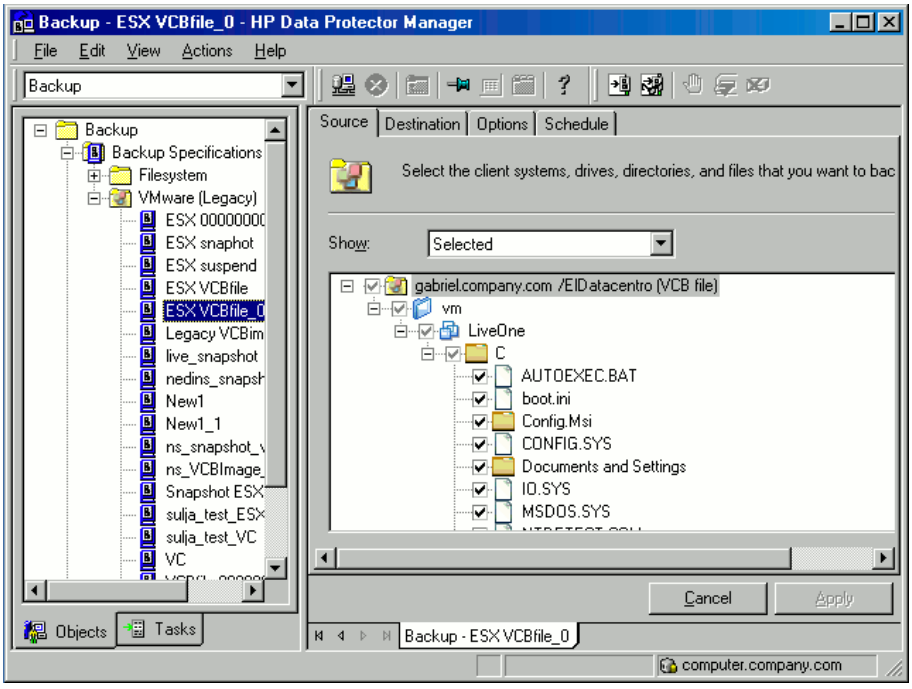

VCBfile backup method: Right-click a virtual machine and click **Mount filesystem** to mount virtual machine filesystems. This may take some time. Then, select the files and folders that you want to back up.

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 100).

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 107).

Figure 50 Application-specific options

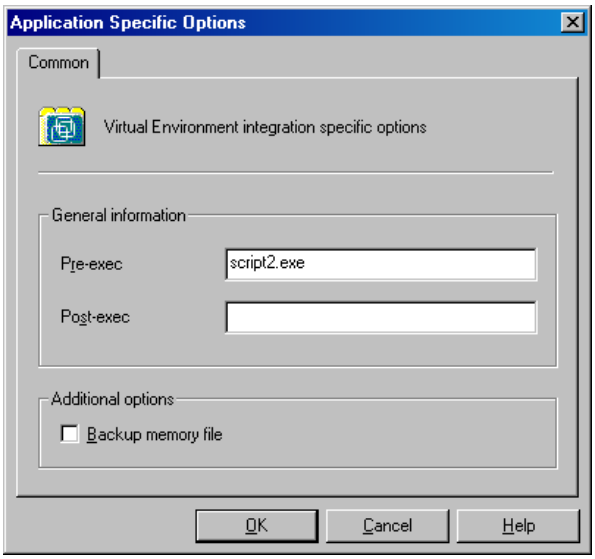

Click **Next**.

8. Optionally, schedule the backup. See "Scheduling backup specifications" (page 107).

   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 108).

Table 21 VMware backup options

Options	Description
Pre-exec, **Post-exec**	The command specified here is run by `vmware_bar.exe` on the VMware management client before the backup (`pre-exec`) or after it (`post-exec`). Do not use double quotes. Type only the name of the command and ensure that the command resides in the following directory: ***VirtualCenter Server system:*** *`Data_Protector_home`*`\bin` ***Standalone ESX Server system:*** `/opt/omni/bin`
Backup memory file	Available only for the **Snapshot** and **Suspend** backup methods. 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.

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 specifications

You can run unattended backups at specific times or periodically. For details on scheduling, see the online 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 specification” (page 108). 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 51 Scheduling a backup specification

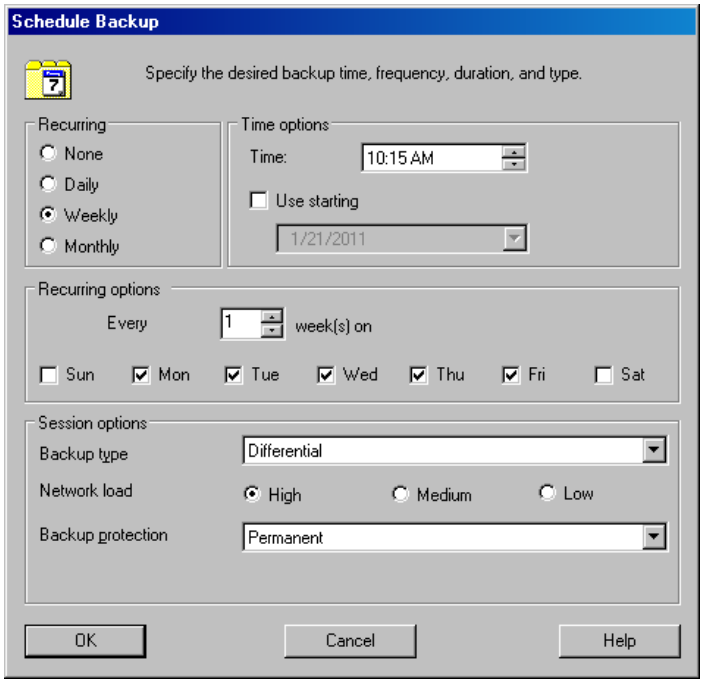

Previewing backup sessions

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 97).
2. Go to the following directory:

   VirtualCenter Server system: `Data_Protector_home\bin`

   Standalone ESX Server system: `/opt/omni/lbin`

3. Run:

   ```
   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 97).
2. Go to the following directory:

   VirtualCenter Server system: `Data_Protector_home\bin`

   Standalone ESX Server system: `/opt/omni/lbin`

3. Run:

   `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`, run:

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

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

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

Preparing for disaster recovery

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

Table 22 What must be backed up

VMware objects	How to back up
ESX Server console	1. Ensure that the Data Protector `Disk Agent` component is installed on all the ESX Server systems. 2. In the Backup context of the Data Protector GUI, right-click **Filesystem** and click **Add backup** 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 http://kb.vmware.com/selfservice/microsites/microsite.do. 3. Start a backup using the newly created backup specification.
VirtualCenter configuration database *(applicable only for VirtualCenter environments)*	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: 1. Ensure that the Data Protector `Oracle Integration` component is installed on the VirtualCenter Server system. 2. In the Backup context of the Data Protector GUI, right-click **Oracle Server** and click **Add backup** to create a backup specification of the Oracle type. In **Application database**, type the name of the VirtualCenter configuration database. Continue with the backup specification creation as described in the *HP Data Protector Integration Guide for Oracle and SAP*. 3. Start a backup using the newly created backup specification.
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 94).

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`.

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 52 Backup object information

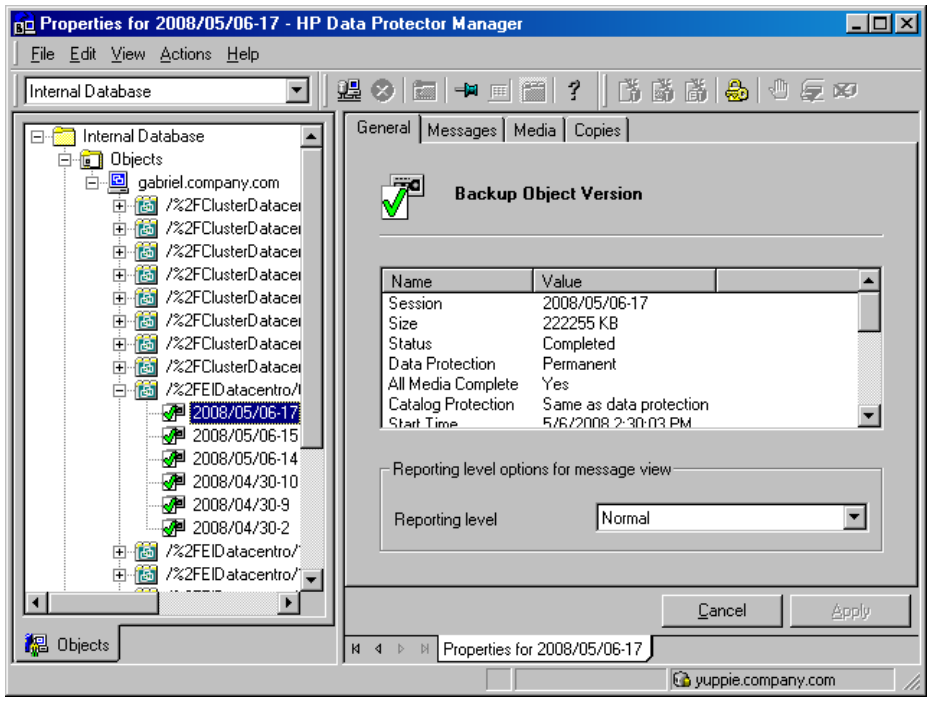

TIP: To view the messages displayed during the session, click the **Messages** tab.

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 97).
2. Go to the following directory:

   VirtualCenter Server system: `Data_Protector_home\bin`

   Standalone ESX Server system: `/opt/omni/lbin`
3. Get a list of VMware backup objects created in a particular backup session:

   `omnidb -session SessionID`
4. 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::/%2FElDatacentro/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 53 Selecting VMware objects for restore (Snapshot, Suspend, VCBimage)

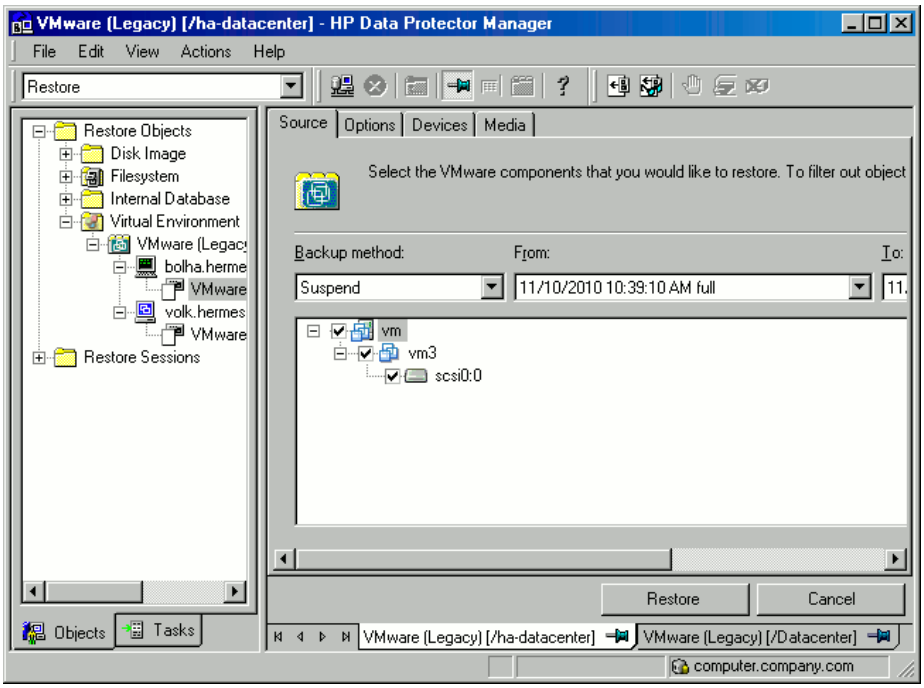

Figure 54 Selecting VMware objects for restore (VCBfile)

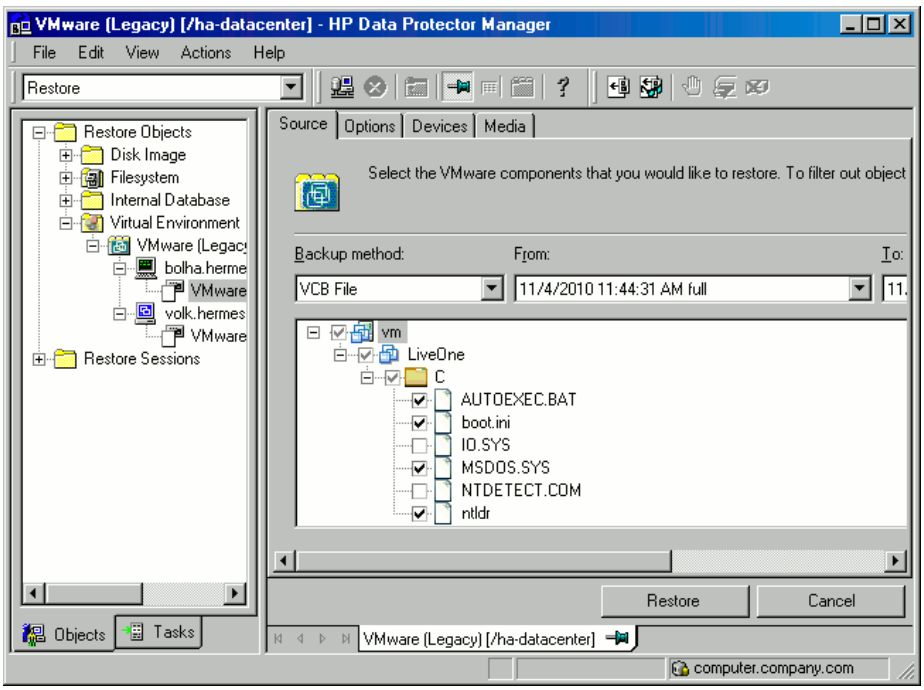

4. In the **Options** page, specify the VMware restore options. For details, see "VMware restore options" (page 115).

Figure 55 Restore options (Snapshot, Suspend, VCBimage)

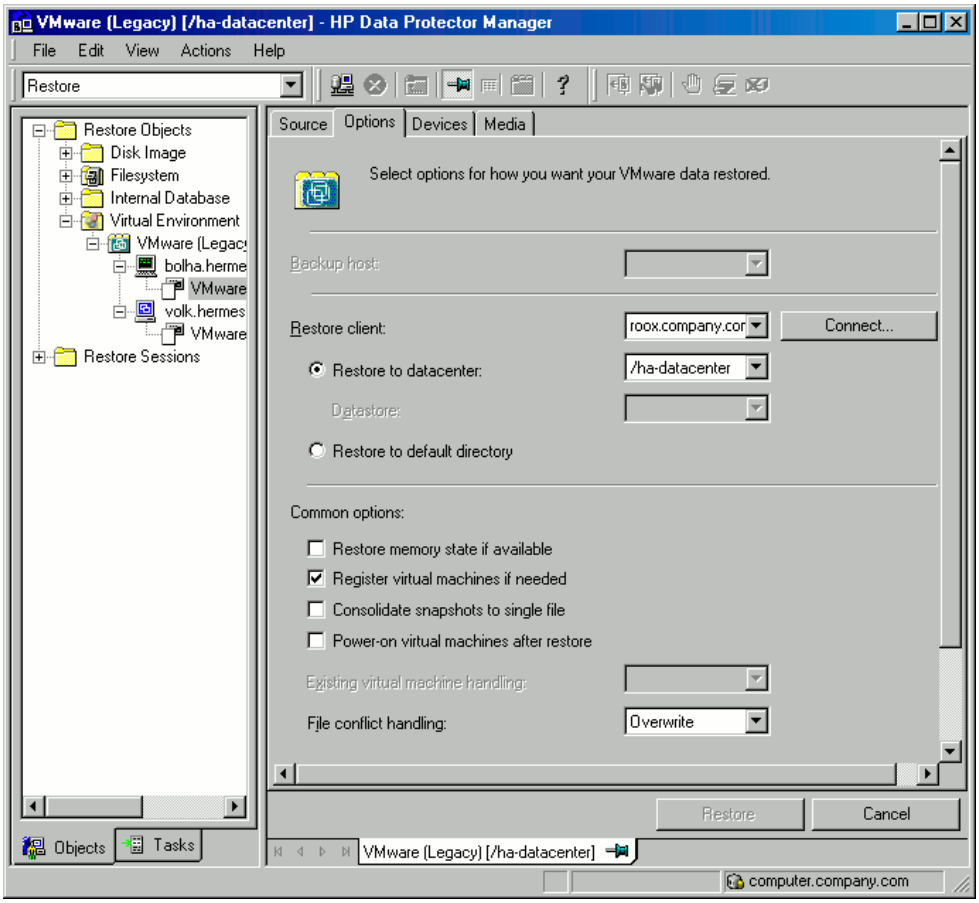

Figure 56 Restore options (VCBfile)

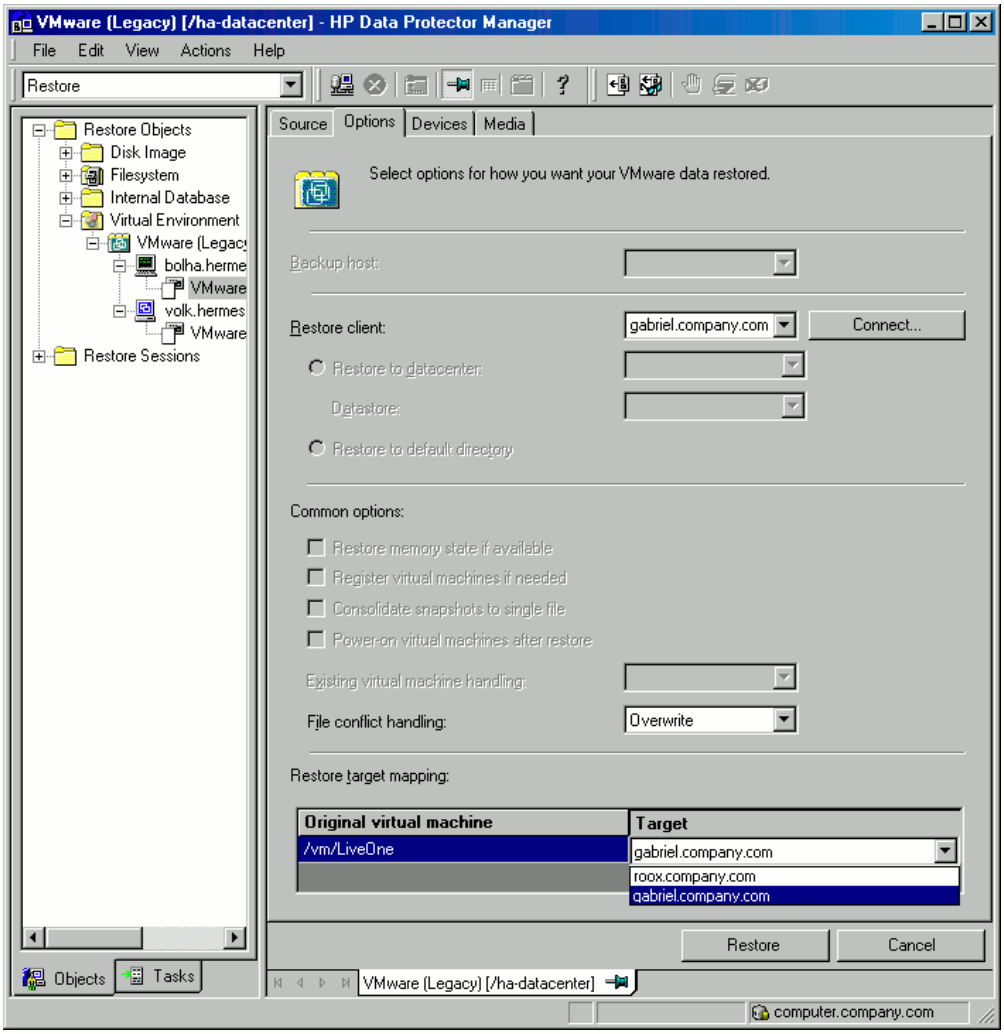

IMPORTANT: If you restore from a **Snapshot** backup in which individual virtual machine disks were backed up, clear the option **Register virtual machines if needed**. Otherwise, the restore fails.

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 123).

5. 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 online Help index: "restore, selecting devices for".
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. 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 23 VMware restore options

GUI/CLI option	Description
Backup host	Not available.
Restore client/ `-destination`	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. Change the default only in the case of a disaster recovery when you restore from a **Suspend**, **Snapshot**, or **VCBimage** 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. 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. If you select a client that has not been configured for use with the Data Protector VMware (Legacy) integration, a warning is displayed. Click **OK** and configure the client as described in "Configuring VMware management clients" (page 97). However, to restore outside a datacenter, you do not need to configure the destination client, just click **Cancel**. To change the configuration of the selected client, click the **Connection** button.
Restore to datacenter/ `-newinstance, -instance`	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. This option is not available if you restore from a **VCBfile** backup.
Restore to default directory	Select this option to restore virtual machines outside datacenter, to the client selected in **Restore client**, and to the directory specified in the `OB2_VMWARE_PATH omnirc` variable on the selected client. This option is not available if you restore from a **VCBfile** backup.
Restore memory state if available/ `-memory`	Select this option to restore the virtual machine memory file as well, if it was backed up. This option is not available if **Restore to default directory** is selected or if you restore from a **VCBfile** backup. Default: not selected.
Register virtual machines if needed/ `-register`	Select this option to register restored virtual machines. Clear this option if you restore from a **Snapshot** backup in which individual virtual machine disks were backed up. Otherwise, the restore fails. If this option is not selected, you need to manually recover the restored virtual machines as described in "Recovering virtual machines manually" (page 118). This option is not available if **Restore to default directory** is selected or if you restore from a **VCBfile** backup. Default: not selected.
Consolidate snapshots to single file/ `-consolidate`	This option is applicable only if **Register virtual machines if needed** is selected. Select this option to commit all existing snapshots (including non-Data Protector ones) to the virtual machine base when the restore completes. This option is not available if **Restore to default directory** is selected or if you restore from a **VCBfile** backup. Default: not selected.
Power-on virtual machines after restore/ `-poweron`	This option is applicable only if **Register virtual machines if needed** is selected. Select this option to put the newly-restored virtual machines online when the session completes.

Table 23 VMware restore options *(continued)*

GUI/CLI option		Description
		This option is not available if **Restore to default directory** is selected or if you restore from a **VCBfile** backup. Default: not selected.
File conflict handling/ `-overwrite [older]`	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 **Snapshot**, **Suspend**, **VCBimage**, or **VCBfile** backup. If you restore from a **Snapshot**, **Suspend**, or **VCBimage** backup and you select **Restore to default directory**, the description for restore from a**VCBfile** backup applies.	
	Overwrite (default)	**Snapshot**, **Suspend**, **VCBimage**: Data Protector unregisters such virtual machines, deletes their files, and then restores them from the backup. **VCBfile**: Data Protector overwrites such files with those from the backup.
	Keep latest	**Snapshot**, **Suspend**, **VCBimage**: Data Protector leaves such virtual machines intact. They are not restored from the backup. **VCBfile**: 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.
	Skip	**Snapshot**, **Suspend**, **VCBimage**: Data Protector leaves such virtual machines intact. They are not restored from the backup. **VCBfile**: Data Protector leaves such files intact. They are not restored from the backup.
Target hosts for restore/ `-target`		This option is available only if you restore from a **VCBfile** backup. Specifies where to restore filesystems of virtual machines. The target system can be any Windows system (physical or virtual).

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 97).
2. Go to the following directory:

   VirtualCenter Server system: *`Data_Protector_home`*`\bin`

   Standalone ESX Server system: `/opt/omni/lbin`
3. Run:

   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.

Run:

```
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` variable `OB2_VMWARE_PATH` to `C:\tmp`.
3. Run:

```
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.

Run:

```
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 118).
- 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 123)).

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.vmxf"
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 57 Conversion (Specify Source)

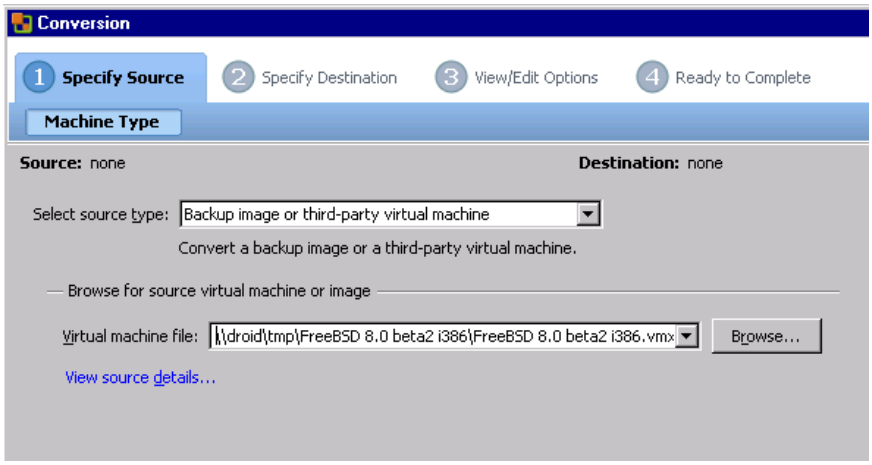

Click **Next**.

8. In the **Specify Destination (Destination Type)** page, provide the login information for the VirtualCenter Server system.

Figure 58 Conversion (Specify Destination 1)

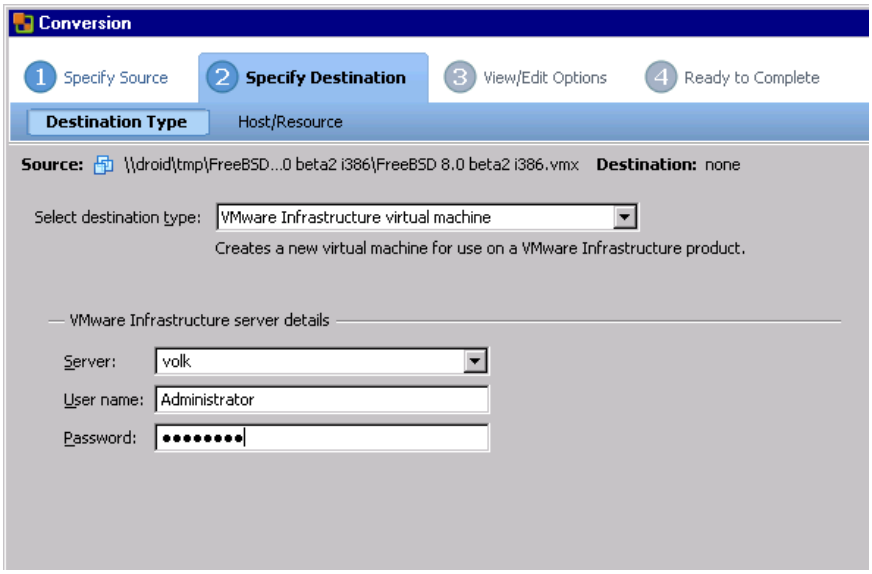

Click **Next**.

9. In the **Specify Destination (Host/Resource)** page, specify the destination ESXi Server system `srsen.company.com`.

Figure 59 Conversion (Specify Destination 2)

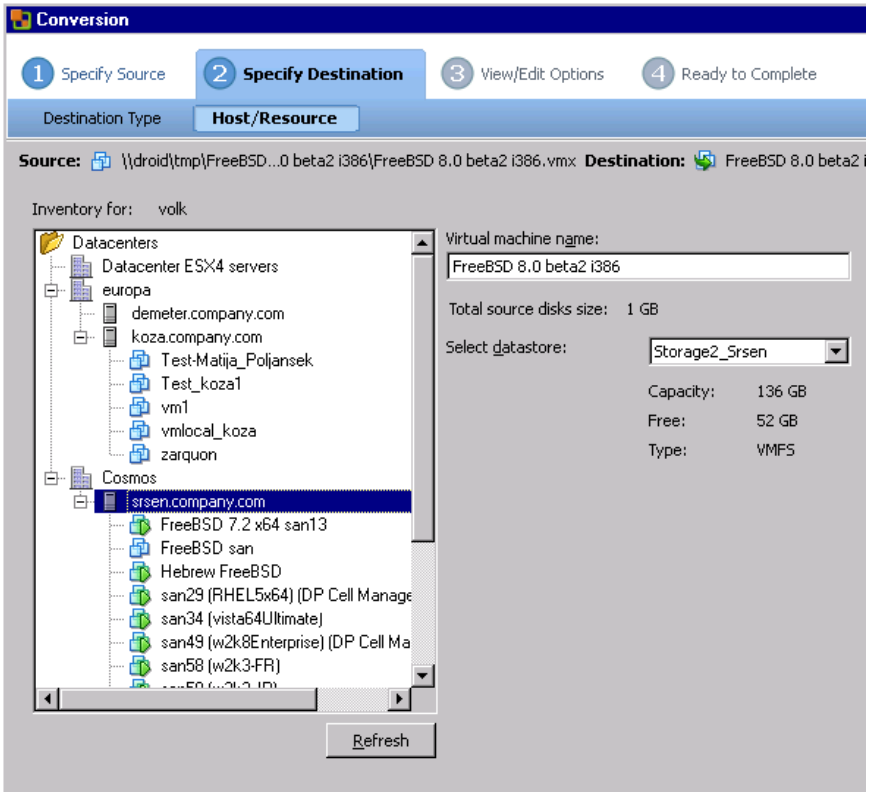

Click **Next**.

10. In the **View/Edit Options** page, modify the options (optional).

Figure 60 Conversion (View/Edit Options)

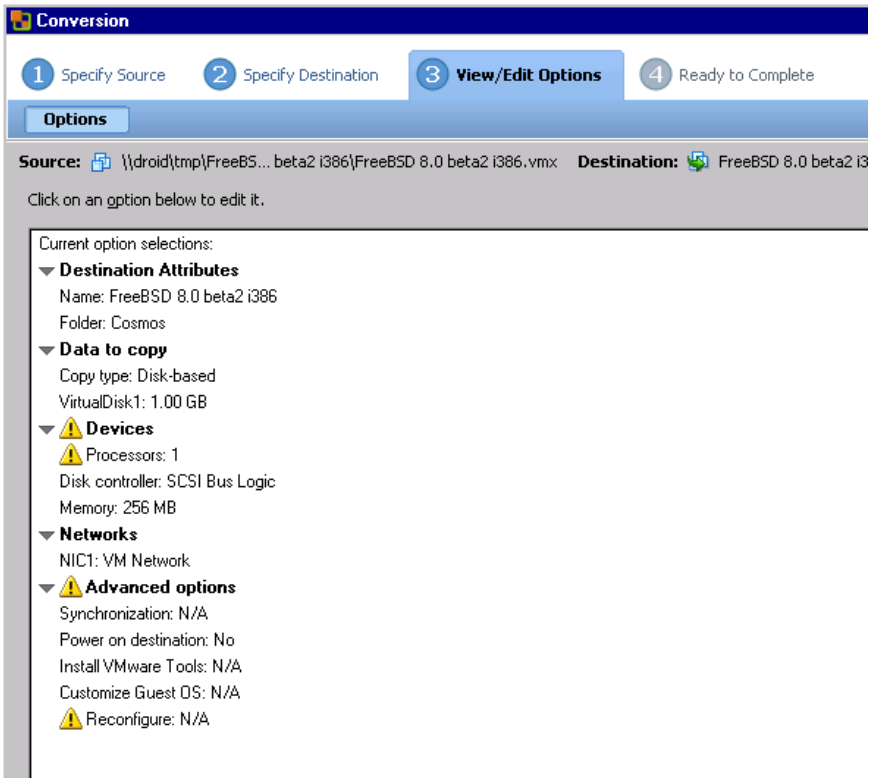

Click **Next**.

11. In the **Ready to Complete** page, review your selection and click **Finish**.

Figure 61 Conversion (Ready to Complete)

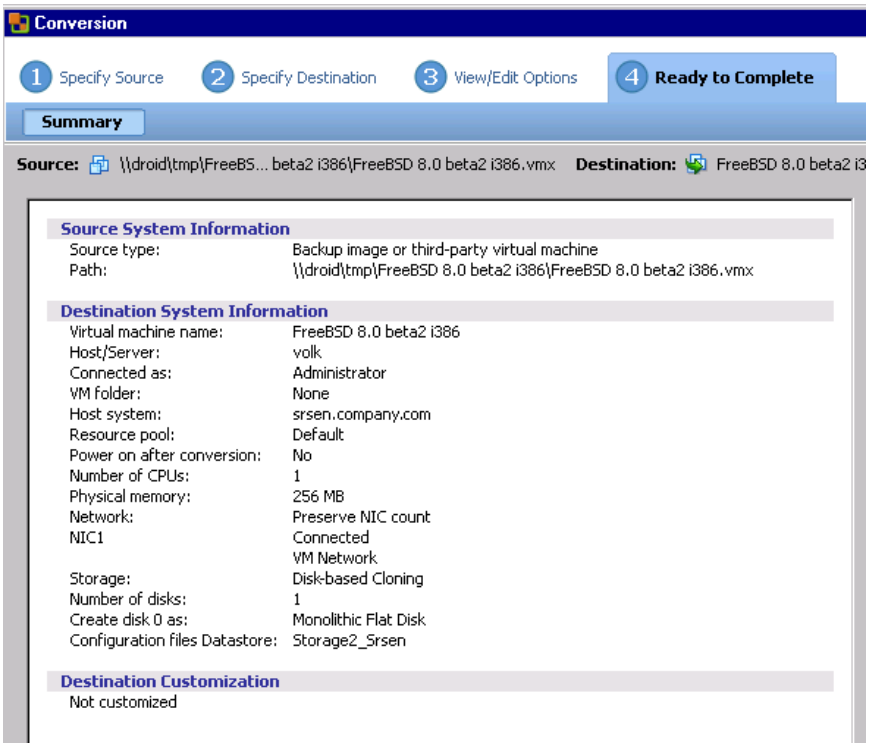

12. After a while, the VMware Converter informs you about the successful recovery.

Figure 62 Recovery completed

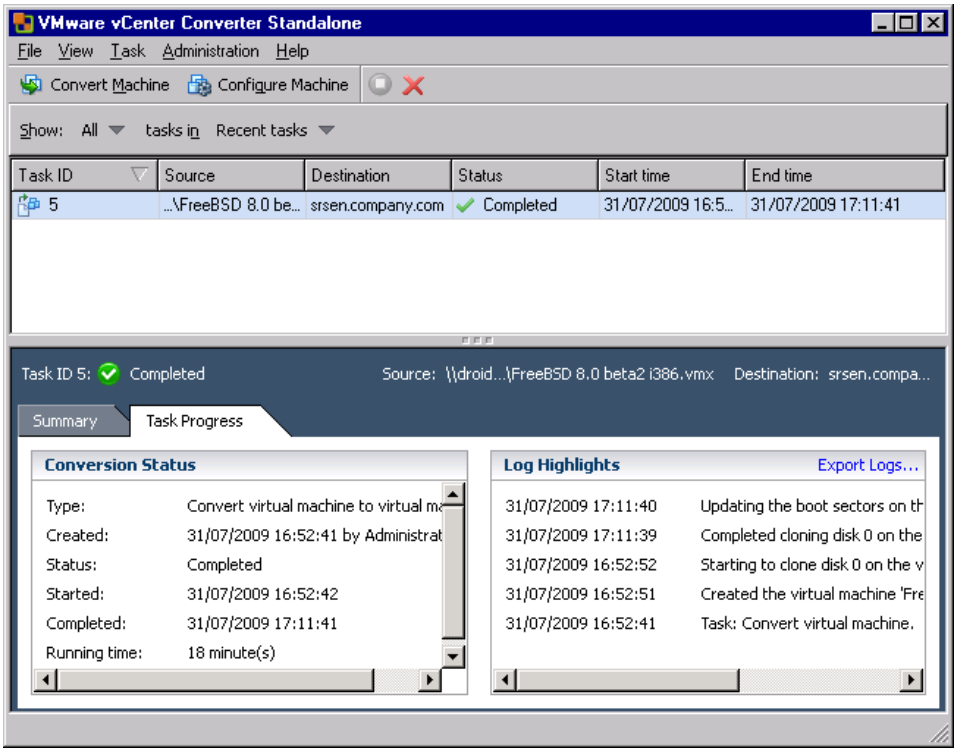

Recovery after restoring to a datacenter

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 24 Virtual machine information

Virtual machine	`helios`
Virtual machine disks	• `scsi0:0` (1.2 GB) base metadata file: `helios.vmdk` base extent file: `helios-flat.vmdk` • `scsi1:2` (1.0 GB) base metadata file: `helios_1.vmdk` base extent file: `helios_1-flat.vmdk`
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.vmxf.2010_09_10_0001
helios.vmxf.2010_09_10_0003
helios.vmxf.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. Run:

   `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 63 Virtual Infrastructure Client

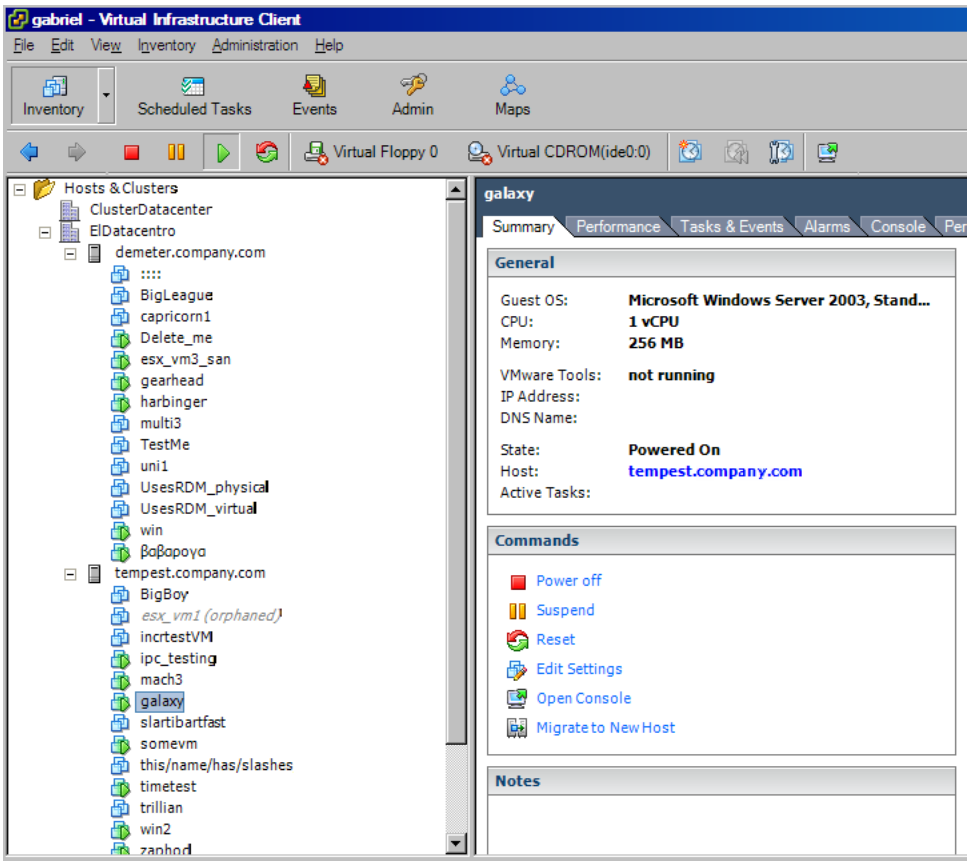

Right-click the virtual machine `galaxy` and click **Edit Settings**.

2. In the Virtual Machine Properties dialog box, select **Hard Disk 1** and click **Add**.

Figure 64 Virtual machine properties

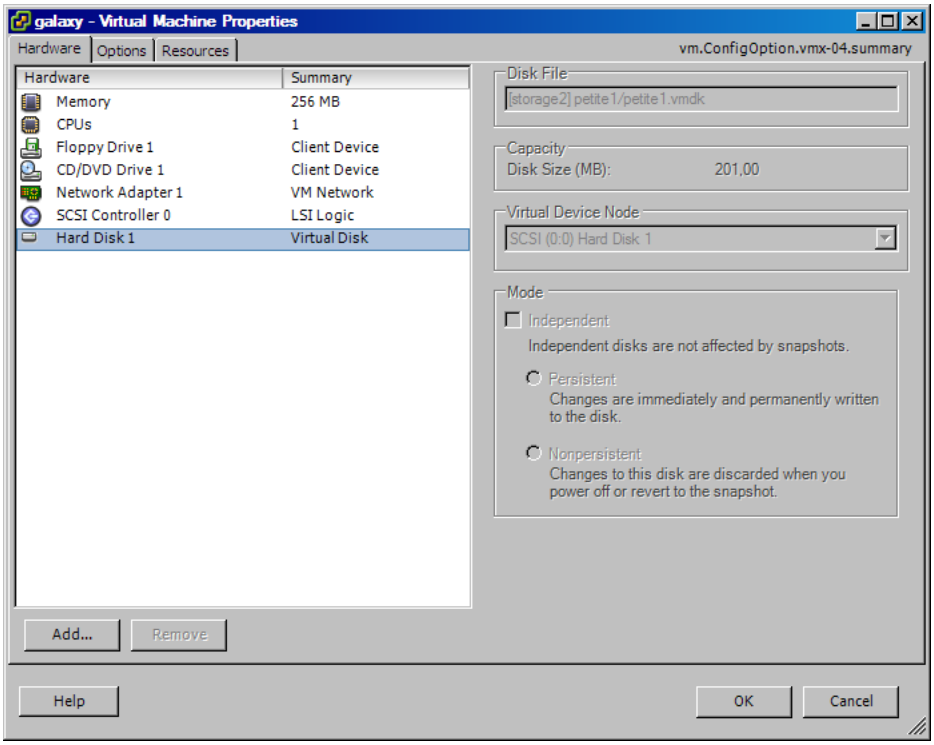

3. In the Select Device Type page of the Add Hardware wizard, select **Hard disk** and click **Next**.
4. In the Select a Disk page of the Add Hardware wizard, select **Use an existing virtual disk** and click **Next**.
5. In the Select Existing Disk page of the Add Hardware wizard, click **Browse**.
6. In the Browse Datastores dialog box, browse to `/vmfs/volumes/storage2/helios`, select `helios_1_recovered.vmdk` and click **Open**.

Figure 65 Browsing datastores

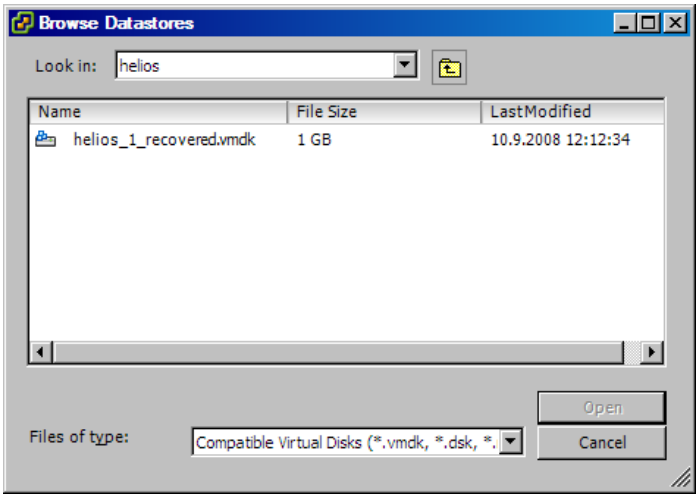

7. In the Advanced Options page of the Add Hardware wizard, select a Virtual Device Node for the new disk and click **Next**.
8. In the Ready to Complete page of the Add Hardware wizard, review your selection and click **Finish**.

Figure 66 Add hardware summary

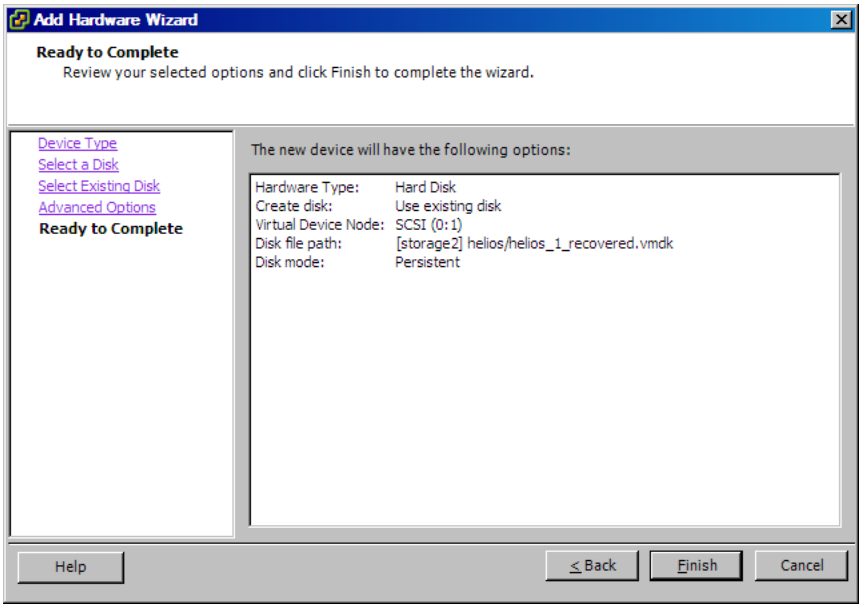

Restoring using another device

You can perform a restore using a device other than that used for the backup. For details, see the online 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 online 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 110).

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 online 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 online 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://www.hp.com/support/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` located in:

  VirtualCenter Server system: *`Data_Protector_home`*`\log`.

  Standalone ESX Server system: `/opt/omni/log`.

- Check if you can do a filesystem backup and restore on the problematic client. For information, see the online 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 97).

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 online 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 online 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 constrains**.
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.

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 unmounting 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.

Action

Restart the backup. The previous failed backup resulted in unmounted VMs, so the issue will not appear.

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.

  For details, see "Data Protector Virtual Environment integration" (page 132).

- ***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.

  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 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 shut down the virtual machines before starting a backup session.

Table 25 Data Protector backup solutions for Microsoft Hyper-V

Feature		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?		✓	✓	
Is virtual machine migration within a Microsoft Hyper-V cluster supported?				✓
Extra licenses needed		Optionally, ZDB and IR	• Online Extension • Optionally, ZDB and IR	• Online Extension • Optionally, ZDB

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 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 138).

As Microsoft Hyper-V and VSS are involved, you must specify the following backup types for each backup session:

- **Microsoft Hyper-V backup type**

  The only supported type is a virtual machine Full backup.

- **VSS backup type**

  The supported types are:

  - Local or network backup
  - VSS transportable

For details on the backup types, see "Backup types" (page 139).

Restore

You can restore Microsoft Hyper-V virtual machines:

- **To the default location**

  Using this option, you can restore virtual machines to the default 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.

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 online Help.

Integration concepts

Supported environments

Data Protector supports environments where Microsoft Hyper-V systems are standalone (**standalone environments**) as well as environments where Microsoft Hyper-V systems are configured in a cluster (**clustered 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.

In one session, you can back up virtual machines from only one Microsoft Hyper-V system.

Figure 67 Standalone environments

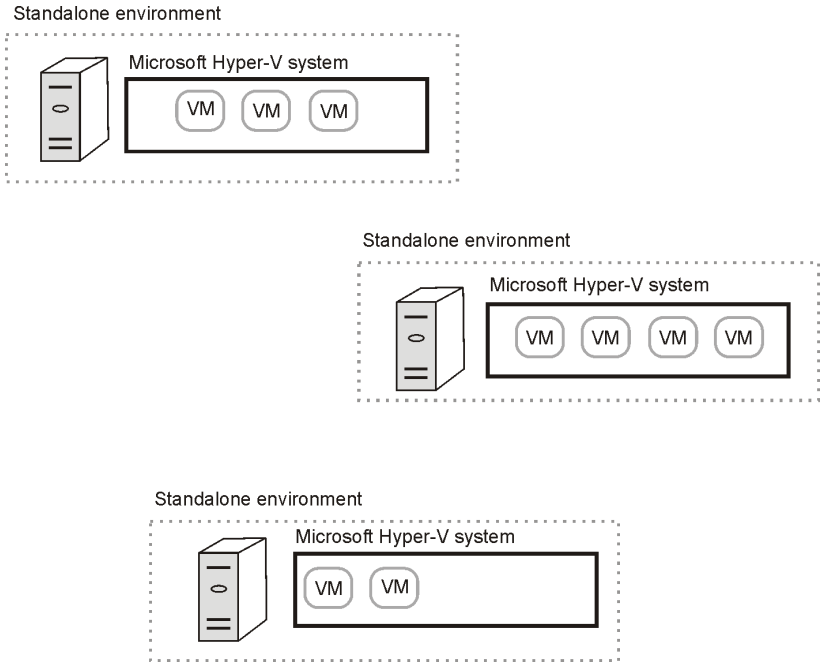

For standalone Microsoft Hyper-V systems, it is only possible to migrate a virtual machine by manually exporting it from one Microsoft Hyper-V system and then importing it into another.

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 virtual system (host). In reality, the request reaches the Microsoft Hyper-V system on which a virtual machine to be backed up is running.

Figure 68 Clustered environments

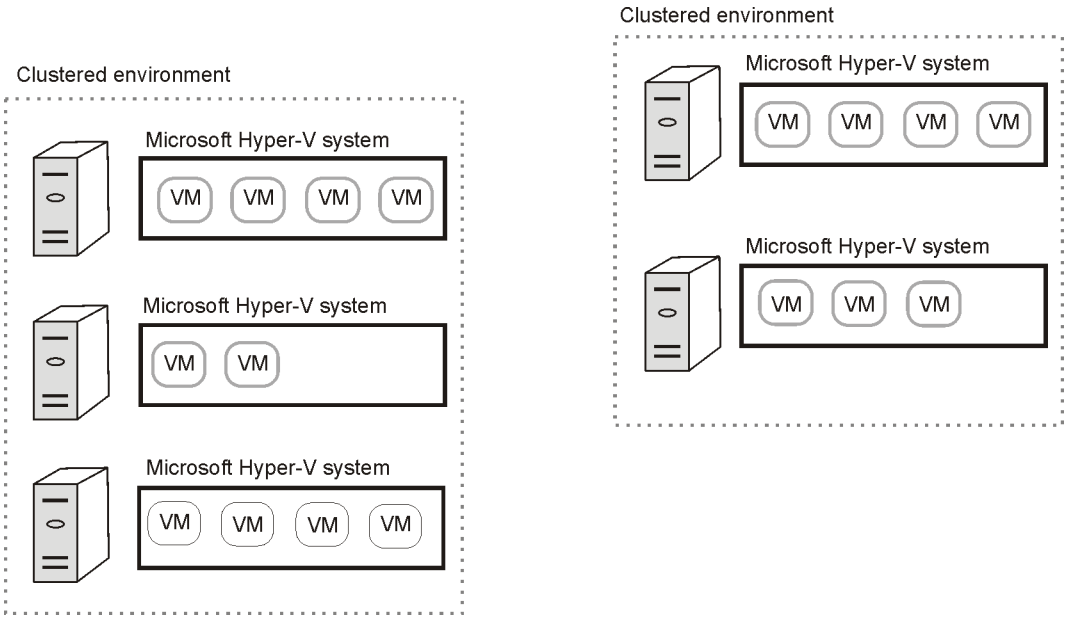

Example

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 69 Windows Server Manager — Cluster VMs

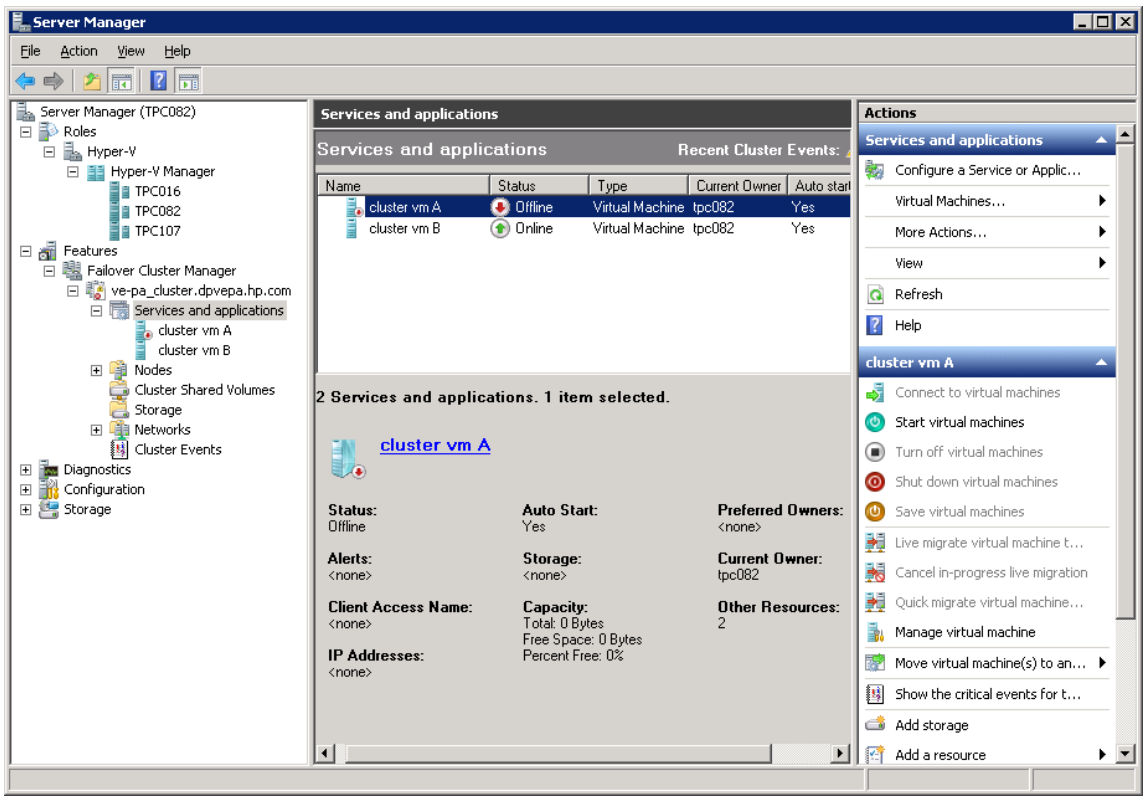

Figure 70 Windows Server Manager — VMs running on TPC082

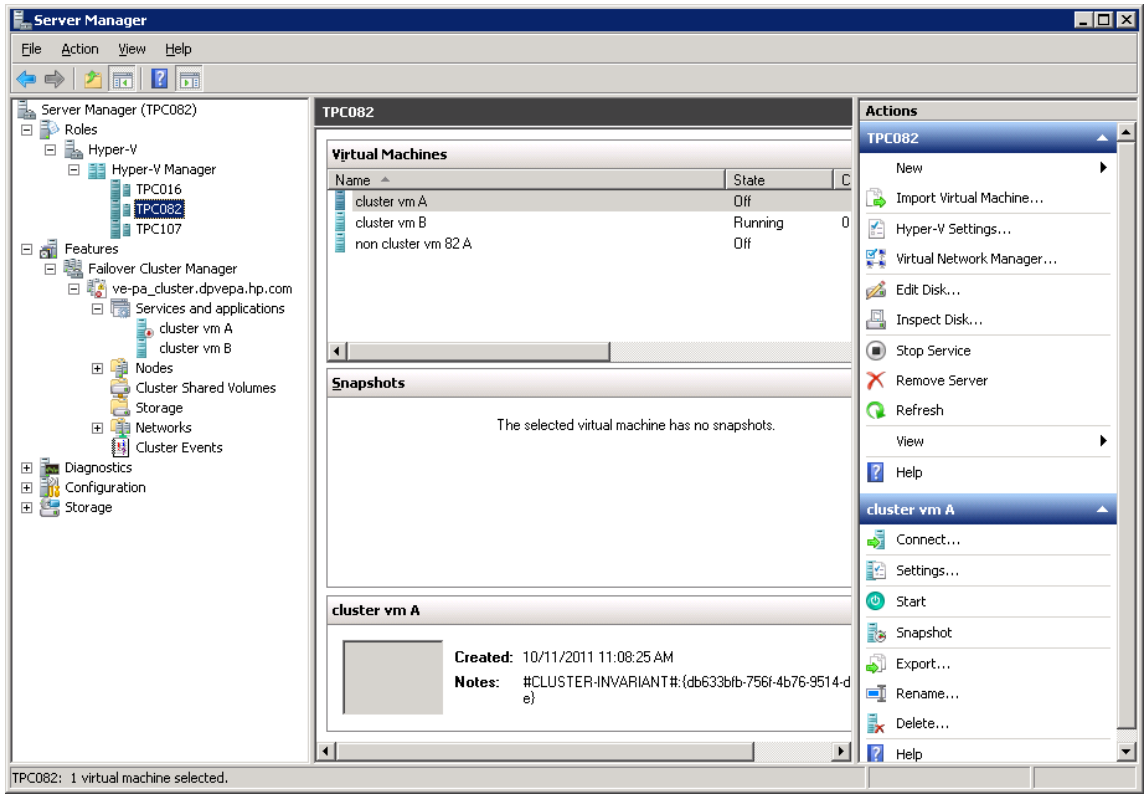

Figure 71 Windows Server Manager — VMs running on TPC107

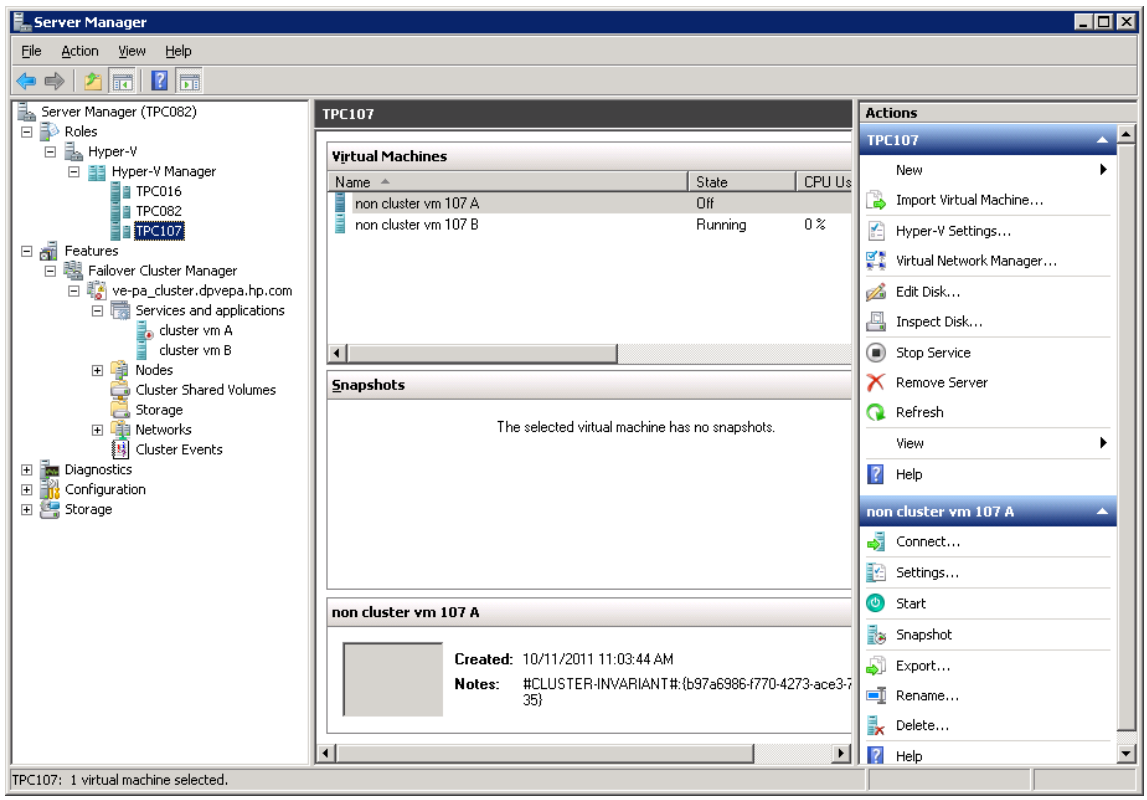

In the Data Protector backup specification, both non-cluster VMs and cluster VMs are displayed for all nodes of the cluster; non-cluster VMs are listed under their physical nodes and cluster VMs under their virtual cluster system. See Figure 72 (page 136).

Figure 72 Data Protector backup specification

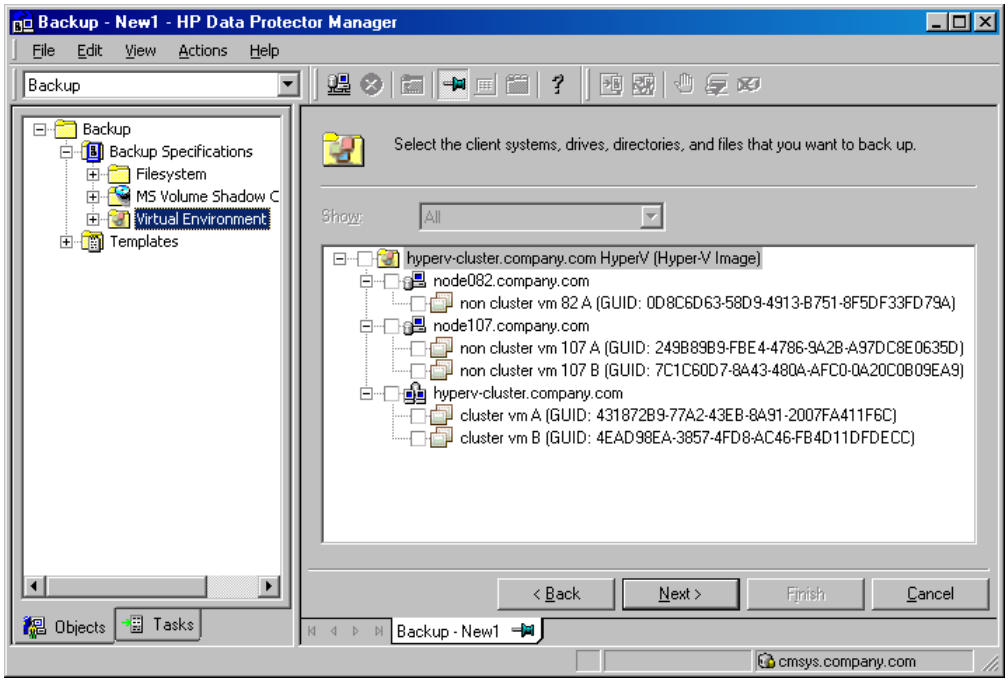

Migration of virtual machines

The Data Protector Virtual Environment integration supports migration of virtual machines between Microsoft Hyper-V systems within a clustered environment. It means that 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 WMI functions, 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 tries to back it up, the backup fails.

Figure 73 Migration of virtual machines

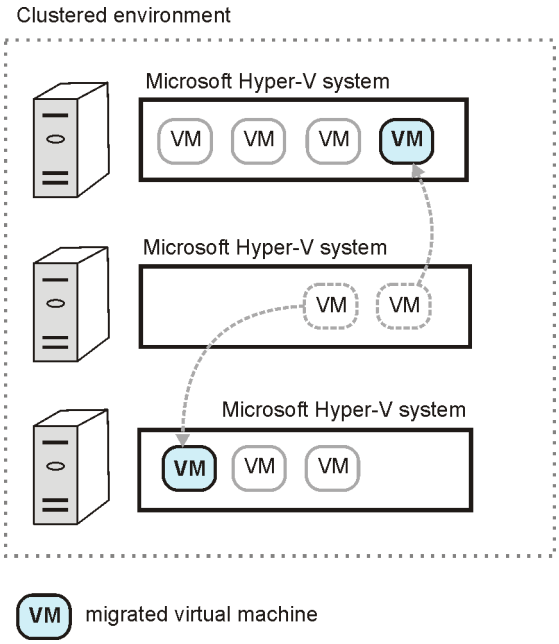

Cluster Shared Volumes

Data Protector supports Microsoft Hyper-V environments in combination with Cluster Shared Volumes (CSV). With CSV, multiple virtual machines can be configured in a cluster (that is, they use the same LUN (disk)), while still being able to move from one Microsoft Hyper-V system to another, independently of each other.

If a virtual machine from the cluster migrates to a different Microsoft Hyper-V system, Data Protector will find it and back it up. You do not need to create a new backup specification after the migration.

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 must be installed on at least one client in a Data Protector Cell (**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_vmware`[5], a dynamic link library providing VMware specific functionality for backup, restore, query and browsing tasks.
- `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 directory button (this button is used when you restore to a 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 from or restore 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: A backup system and a backup host are not one and the same thing. 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 tape.

Data Protector Media Agents

Data Protector Media Agents must be installed on the 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.

5. The Data Protector Virtual Environment integration can also be used to back up virtual machines from VMware vSphere virtualization environments. For details, see the VMware part.

Backup concepts

The Data Protector Virtual Environment integration for Hyper-V is based on the Data Protector Microsoft Volume Shadow Copy Service Integration. With Data Protector Virtual Environment integration, Data Protector allows users to back up and restore virtual machines, without having to get heavily involved in VSS operations.

Hyper-V Image backup method

The Hyper-V Image backup method enables you back up complete Microsoft Hyper-V virtual machines. 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 some DLLs (specifically, `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 tape.
5. `vss_bar.exe` informs `vepa_bar.exe` that the backup has completed.

NOTE: You can use the omnirc variable `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 variable is set to `INFINITE`).

Figure 74 Hyper-V Image method (VSS transportable)

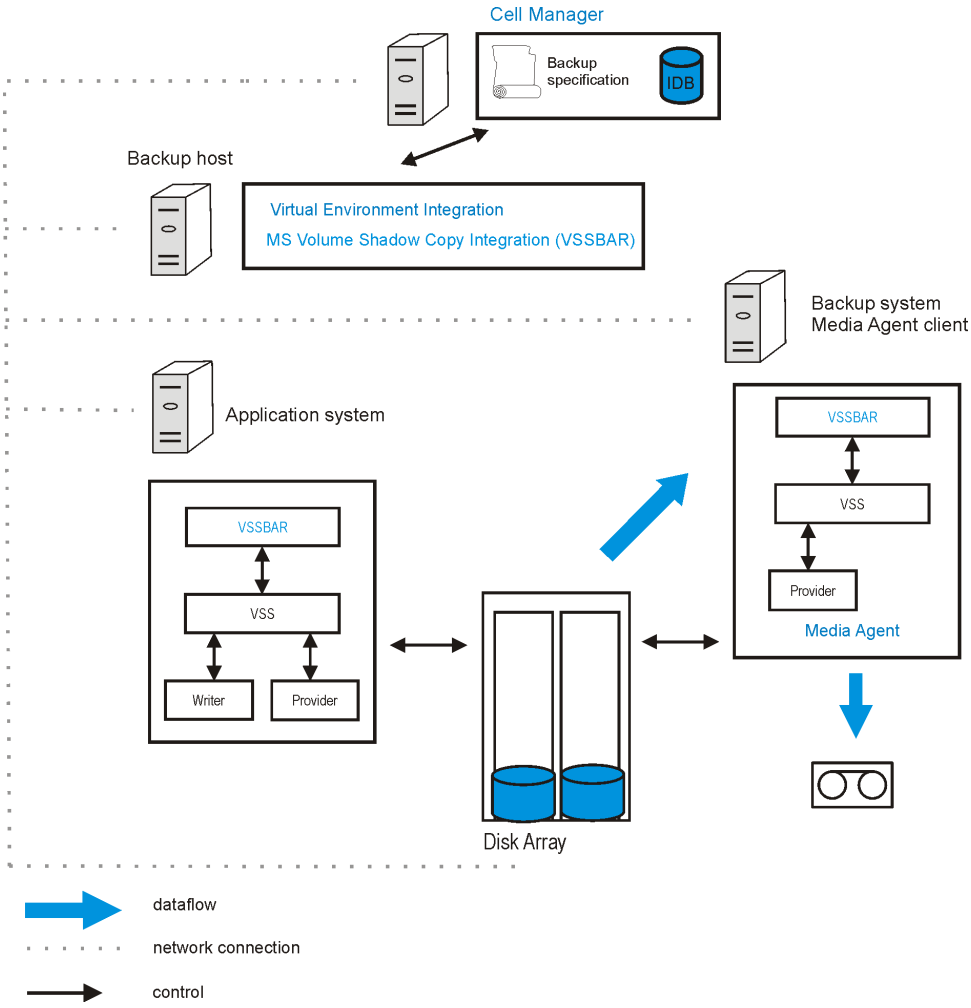

In "Hyper-V Image method (VSS transportable)" (page 139), the backup system is also a media agent client (it has the `Media Agent` component installed and a device connected to it).

For more information, see the "Concepts" section in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*. Skip the instant recovery-related sections.

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 backup specification level.

Microsoft Hyper-V backup types

You can select from among the following Microsoft Hyper-V backup types:

- **Full**

  All virtual machine files are backed up.

VSS backup types

You can select from among the following VSS backup types:

- **Local or network backup**

  This type is used for a single-host VSS configuration. It means that 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 hardware provider) or standard Data Protector backup (using the software provider).

  NOTE: In a Microsoft Hyper-V cluster environment, a local backup using the hardware provider is not supported. This is a Microsoft limitation.

- **VSS transportable**

  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 tape can be performed. A 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 limitation.

For details, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Quiescence

For Windows virtual machines that are online 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 the virtual machine 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 Hyper-V systems.

If the guest operating system is anything other than 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). This means that in the case of a ZDB, 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 so some data on the virtual machine may have changed in the meantime.

Backup considerations

- **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.
- **Storage:** Due to a Microsoft Hyper-V 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 Internet SCSI (iSCSI) initiators.

  For details, see:

  Microsoft TechNet Library

  Search for "Hyper-V, Planning for Backup".
- **ZDB only:**
  - 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 tape.
  - *HP P4000 SAN Solutions:* Due to a hardware provider limitation, virtual machines whose files reside on an HP P4000 SAN Solutions disk array cannot be backed up online.

    There are two workarounds:

    1. Use the software provider to back up the virtual machines online (this is a non-ZDB backup).
    2. Back up the virtual machines offline (ZDB backup).
- Since the Data Protector Virtual Environment integration is based on the Data Protector Microsoft Volume Shadow Copy Service integration, also consider the Microsoft Hyper-V writer specifics described in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.
- Using the Data Protector Virtual Environment integration, you cannot back up Microsoft Hyper-V system configuration data.

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 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 will be deleted before the files from the backup are restored.

If, for security reasons, you want to keep 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 will be 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 will cause the import to fail. Also after importing, you may have network problems unless you change the IP address.

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 answer to the question where a restored VM will be running depends on the restore client selection and the state of the environment at restore time (and not on the state of the environment at 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 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 a backup host. 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.

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 considerations

- **Concurrent sessions:** Restore sessions that use the same devices or restore to the same Microsoft Hyper-V system (or cluster) cannot run concurrently.
- **Different backup methods:** Hyper-V objects that were created using the Data Protector Microsoft Volume Shadow Copy Service integration cannot be restored using this integration.
- Since the Data Protector Virtual Environment integration is based on the Data Protector Microsoft Volume Shadow Copy Service integration, also consider the Microsoft Hyper-V writer specifics described in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Configuring the integration

To configure the integration:

- Enable automatic mounting as described in "Enabling automatic mounting on Microsoft Hyper-V systems" (page 144).
- Import and configure Microsoft Hyper-V systems as described in "Importing and configuring Microsoft Hyper-V systems" (page 144).
- *Microsoft Hyper-V cluster only:* Resolve every cluster node as described in "Configuring Microsoft Hyper-V clusters" (page 144).

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 in various architectures, 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 only:

    – If the virtual machine files are located on a disk array, ensure that you have configured and installed the corresponding hardware provider.

    – For VSS transportable backups, the `MS Volume Shadow Copy Integration` must also be installed on the backup system.

    For details, see the "Configuration" section in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

    NOTE: As instant recovery is not supported with the Data Protector Virtual Environment integration, the corresponding Data Protector disk array integration component does not need to be installed on Microsoft Hyper-V systems and backup systems.
- For supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/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 Microsoft Hyper-V systems and 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 on Microsoft Hyper-V systems

To be able to perform online backups, enable automatic mounting of new volumes on all Microsoft Hyper-V nodes by running the `MOUNTVOL /E` command on each node.

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:

```
omnidbvss -resolve -apphost HyperVNode
```

The resolution needs to be performed sequentially and, normally, only needs to be performed once. For special cases when the cluster nodes also need to be resolved, 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.

NOTE: **Clustered environment:** If your Microsoft Hyper-V systems are configured in a cluster, you need to import all physical cluster nodes and the virtual cluster as **Hyper-V** clients.

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 75 Importing a Microsoft Hyper-V client (Name and Type)

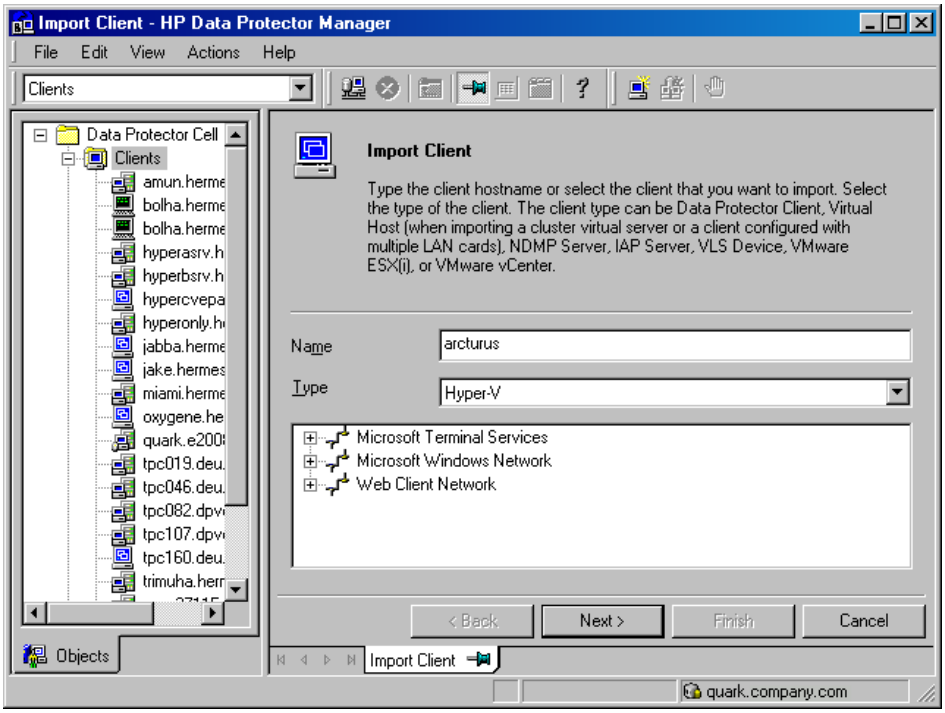

4. Specify login credentials:

   Username and **Password**: Specify an operating system user account that has appropriate permissions to access Microsoft Hyper-V WMI services. You can specify the user account in the following format:

   `DOMAIN\User`

   If the domain is not specified, it is detected automatically.

   In case of a Microsoft Hyper-V cluster, both the domain and the user name must be specified.

   NOTE: The User Account Control on the Microsoft Hyper-V system must be configured to elevate user rights automatically.

   Click **Finish**.

Figure 76 Importing a Microsoft Hyper-V client (Login credentials)

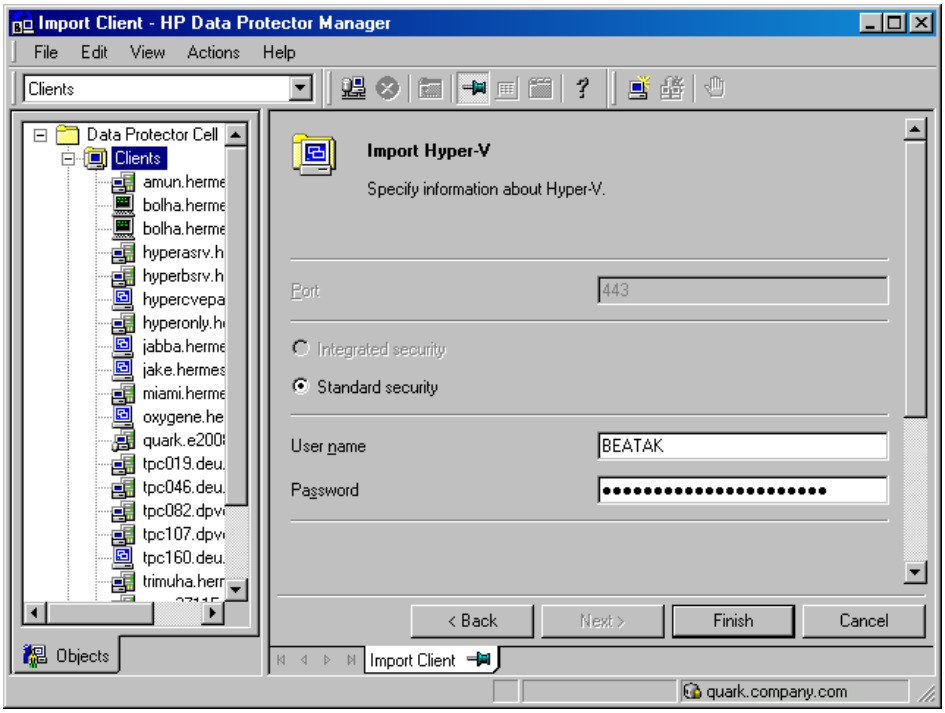

NOTE: For details of how to change the parameters later, see "Changing the configuration of Microsoft Hyper-V systems" (page 146).

Changing the configuration of Microsoft Hyper-V systems

When you update the credentials for connecting to a Microsoft Hyper-V system, 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 online Help index: "user groups".

To update the 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 and click **Apply**.

Using the Data Protector CLI

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

   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 `vepa_util.exe` man page or the *HP Data Protector Command Line Interface Reference*.

Customizing the Data Protector behavior with `omnirc` variables

The `omnirc` variables 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 variables, see the online Help index: "omnirc options".

Backup

To back up virtual machines, create a backup specification and then start a backup session.

For details on backup concepts, see "Backup concepts" (page 138).

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. In **Client**, select the Microsoft Hyper-V system that you want to back up from. The drop-down 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 cluster system. Regardless of what part of the cluster you select, you will be able to back up all VMs residing on the cluster.

   In **Backup Host**, select a system to be used to control the backup. The drop-down list contains all clients that have the Data Protector `Virtual Environment Integration` and the Data Protector `MS Volume Shadow Copy Integration` components installed.

Figure 77 Selecting a client system and backup host

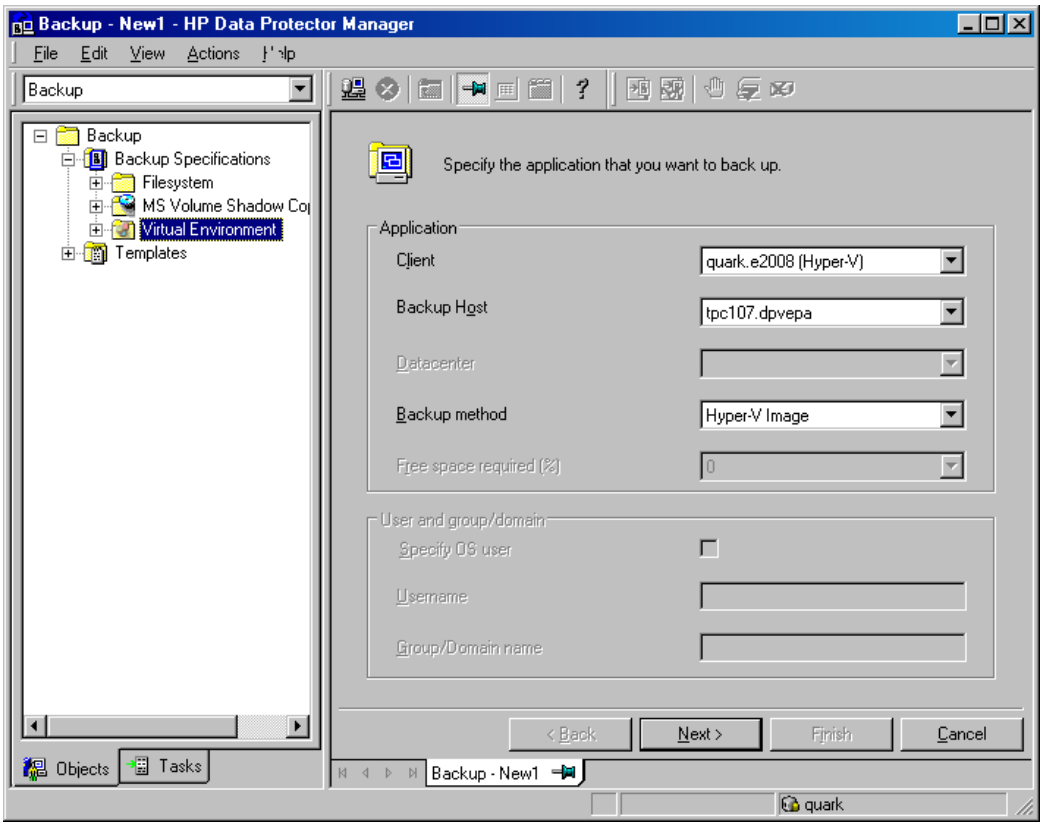

Click **Next**.

5. This page applies to a ZDB environment only. In a non-ZDB environment, click **Next**. For ZDB, 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 tape or where you want your shadow copies to be presented and mounted after the backup. The 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 78 Selecting the backup objects

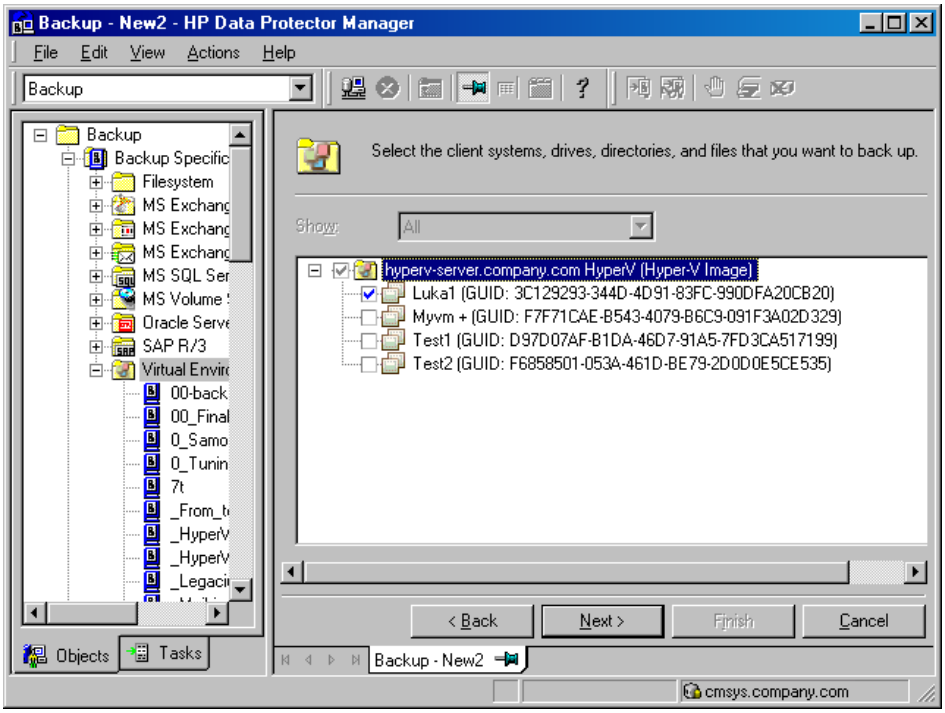

In a cluster, both non-cluster VMs and cluster VMs are displayed for all nodes of the cluster; non-cluster VMs are listed under their physical nodes and cluster VMs under their virtual cluster system.

Figure 79 Selecting the backup objects in a cluster

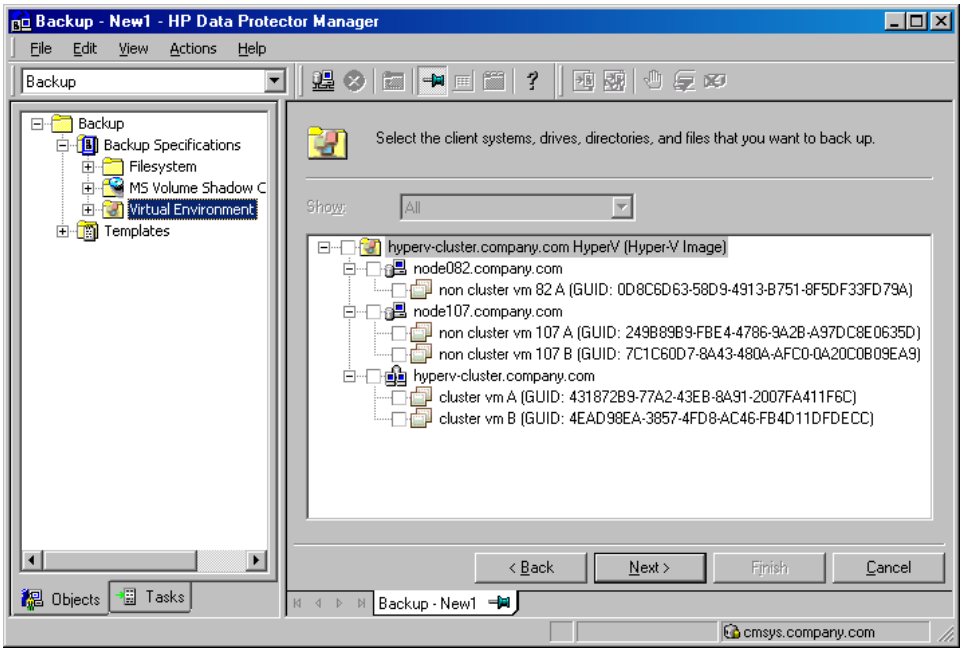

Click **Next**.

7. Select the devices you want 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 150).

Figure 80 Application-specific options

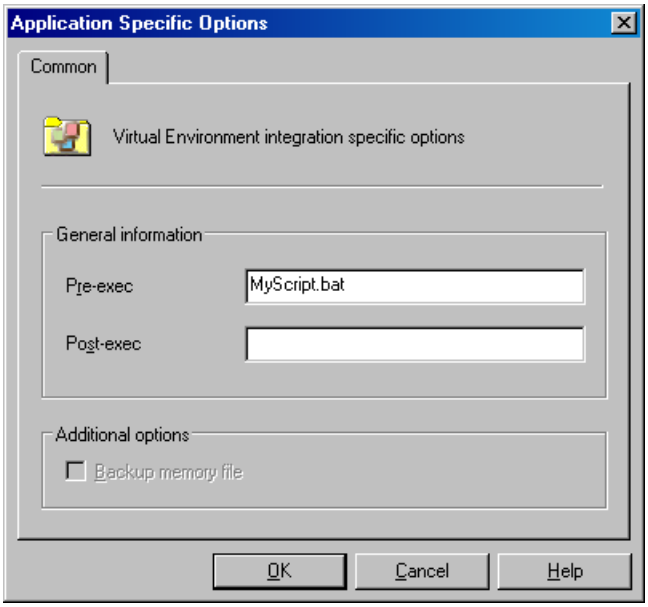

Click **Next**.

9. Optionally, schedule the backup. See "Scheduling backup specifications" (page 151).

Click **Next**.

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

Table 26 Application-specific backup options

Options	Description
Pre-exec, **Post-exec**	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: ***Windows Vista, Windows Server 2008:*** *`Data_Protector_program_data`*`\bin` ***Other Windows systems:*** *`Data_Protector_home`*`\bin`
Backup memory file	Not available.

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.

Figure 81 Modifying a backup specification

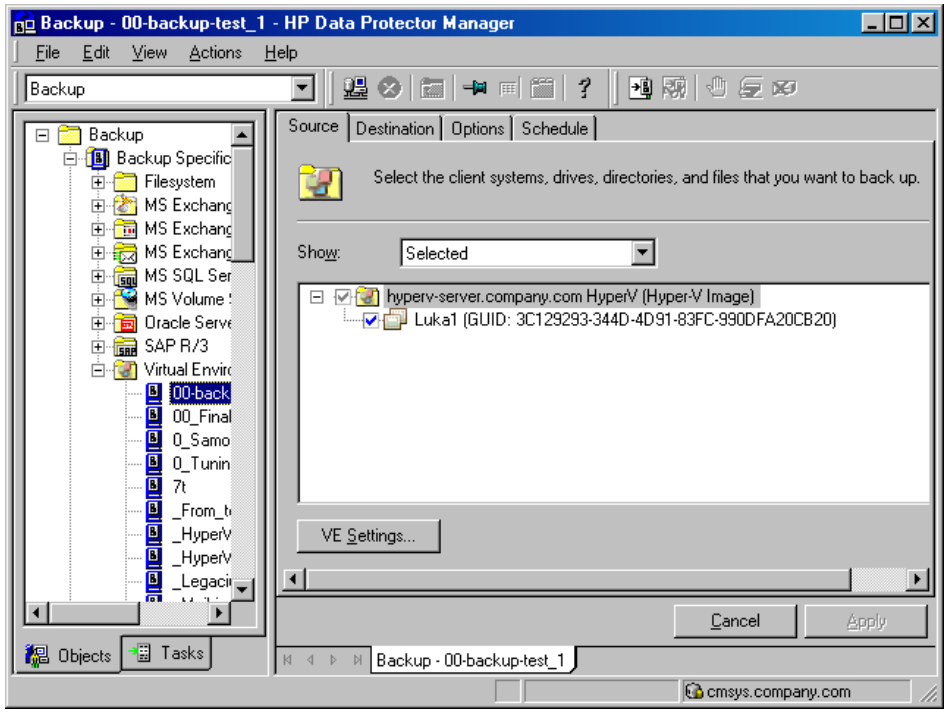

To display Virtual Environment settings, click the **VE Settings** button. Some of the settings can be modified.

Figure 82 Virtual Environment Settings

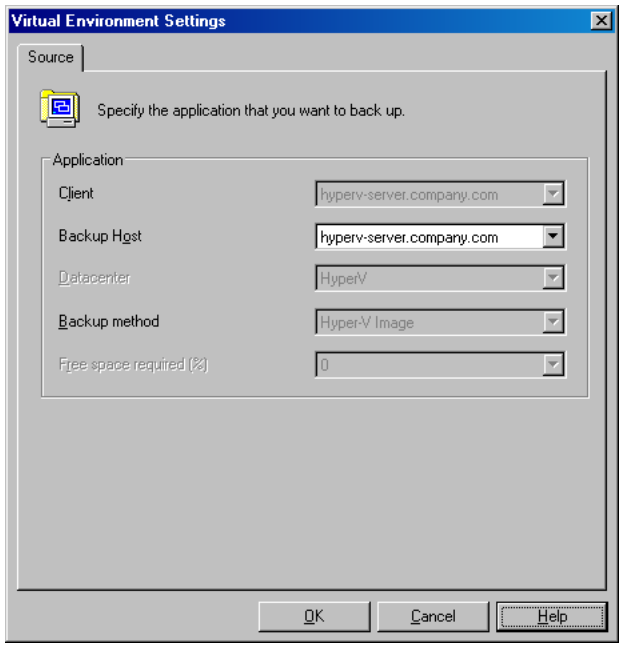

NOTE: To see all virtual machines in the source page, not just those you selected, select **All** in the **Show** option.

Scheduling backup specifications

You can schedule a backup session to start automatically at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

Scheduling example

To schedule backups at 8:00, 13:00, and 18:00 during weekdays:

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 backups" (page 152).

   For Microsoft Hyper-V backups, the backup type is automatically set to **Full**.

   Click **OK**.
3. Repeat Step 1 and Step 2 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

Figure 83 Scheduling backups

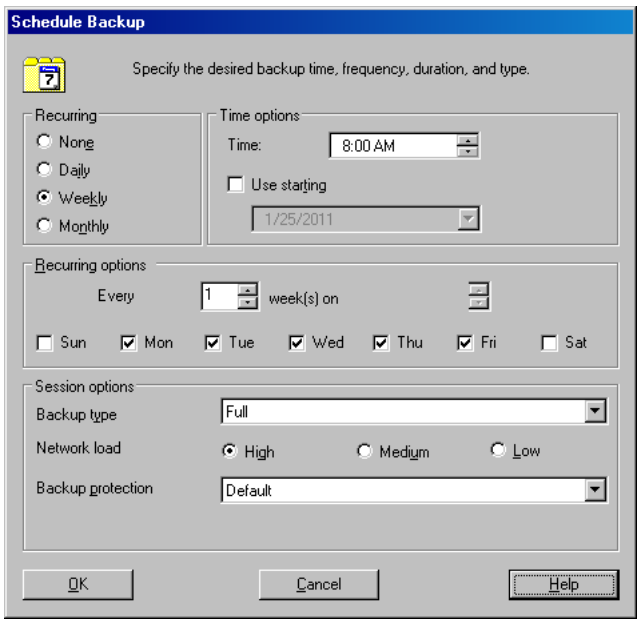

Starting backup sessions

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**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Network load** and 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.

   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:

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

   where *VirtualEnvironmentMode* can only be `full`.

   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`, run:

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

Restore

You can restore Microsoft Hyper-V virtual machines using the Data Protector GUI or CLI.

For details on restore concepts, see "Restore concepts" (page 141).

Finding information for restore

You can retrieve information about backup sessions (such as information on the media used and the messages reported during the backup) from the Data Protector IDB.

To retrieve information, 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 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 the backup object and click **Properties**.

TIP: To view the messages reported 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.

   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 `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 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 **Options** page, specify the Data Protector Microsoft Hyper-V integration restore options. For details, see " Restore options" (page 156).

Figure 84 Restore options

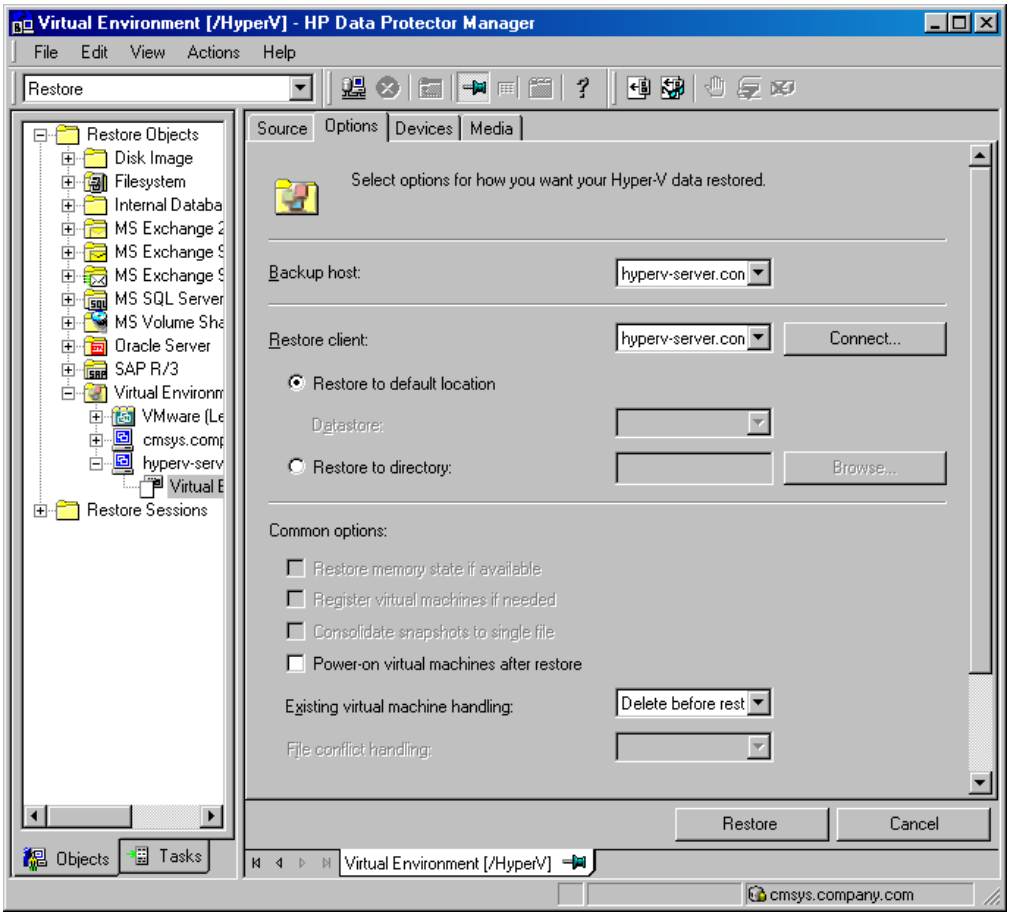

5. In the **Devices** page, specify which devices to use for restore.

   For details, see the online Help index: "restore, selecting devices for".
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.

8. Specify **Report level** and **Network load**.

   Click **Finish** to start the restore.

   If the session succeeds, the message `Session completed successfully` is displayed at the end.

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.

   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
 -instance hyperv
 [-session BackupID]
 -vm GUID [-vm GUID...]
 [HypervClient | Directory]

 HypervClient
 [-destination DifferentHyperVClient]
 [-poweron]

 Directory
 -directory RestoreDirectory
```

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

   NOTE: 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 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 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 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 default location, using the backup session `2011/01/11-1`, and to power the newly restored virtual machines on when the session completes, run:

```
omnir -veagent -virtual-environment hyperv -barhost
backuphost.company.com -apphost hyperv1.company.com -instance hyperv
```

```
-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 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 -instance hyperv
-session 2011/2/12-5 -vm 62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C -vm
54C22930-E3B9-43AA-AFCD-1E90BB99F130 -directory c:\tmp
```

Restore of cluster-aware virtual machines

After a restore of a virtual machine that was cluster-aware and was, for example, deleted, so no cluster resource for that virtual machine is available any more, the restored virtual machine is a local virtual machine running on a node. This is because a broken cluster resource configuration is not automatically restored with the virtual machine restore. The virtual machine requires manual configuration to become cluster-aware again, which is performed as follows:

1. Make sure the virtual machine is not running.
2. In the Windows Server Manager, under **Features** > **Failover Cluster Manager** > the virtual cluster, 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 would like to configure and click **Next**.
5. In the Summary page, review the summary and click **Finish**.

Restoring using another device

You can restore using a device other than that used for a backup. For details, see the online Help index: “restore, selecting devices for”.

Restore options

Table 27 Restore options

Option GUI/CLI	Description
Backup host `-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.
Restore client / `-apphost, -destination`	Select the Microsoft Hyper-V client to which the selected virtual machines should be restored. By default, the client from which the virtual machines were backed up is selected.
Restore to default location /	Select this option to restore virtual machines to the original location on the selected restore client.

Table 27 Restore options *(continued)*

Option GUI/CLI		Description
Restore to directory / `-directory`		Select this option to restore virtual-machine files to a directory on the backup host. You can use the browse button to find the target directory.
Power-on virtual machines after restore / `-poweron`		Select this option to power the virtual machines on once they are restored. This option is not available when restoring to a directory.
Existing virtual machine handling	Specifies Data Protector's behavior when restoring existing virtual machines.	
	Delete before restore/ `-deletebefore`	Select this option to delete an existing virtual machine before it is restored, and then restore it from new. Default: selected.
	Skip restore/ `-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.

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 online 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 Microsoft Volume Shadow Copy troubleshooting information, 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 online Help index: “patches”.
- 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://www.hp.com/support/manuals for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your browsing, backup, or restore failed:

- Examine system errors reported in the `debug.log` located in:

  `Data_Protector_home``\log`.

- Check if you can do a filesystem backup and restore on the problematic client. For information, see the online Help.

Problems

Problem

Backup session of a Microsoft Hyper-V virtual machine ends unexpectedly

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 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 not enough free disk space for shadow copies, use of a non-NTFS filesystem, and so on.

Action

Check if automatic mounting 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 logs 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 want 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, backup, or restore operation fails due to a wrong password

When the application client and the backup host are the same system, the configuration check of the user and password is not performed (in other words, in such a setup, the configuration check always seems to be successful, even if the credentials are wrong). In such a case, even if a wrong password is specified at configuration time, subsequent browse, backup, and restore operations will work. However, when the connection is not established locally (for example, when a different backup host is used), these operations will fail due to a wrong password.

Action

The workaround for the configuration check is to use a backup host that is not the application client.

Problem

The Data Protector Inet service configuration is missing

During a restore session, a message similar to the following is displayed:

```
[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

You can 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. You can do 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 always occurs when the Virtual Environment integration for Microsoft Hyper-V connects to the cluster namespace. The message can be ignored, as everything works normally. Details are explained in the Microsoft KB article 2590230: support.microsoft.com/kb/2590230.

Action

No action is required.

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 162).

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 appendix 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://www.hp.com/support/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. Refer to "Special considerations" (page 175).

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. Refer to "Specifying a restore" (page 173).

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 90 (page 174).

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. Refer to "Specifying a restore" (page 173).

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 85 (page 164).

Figure 85 Example backup and restore

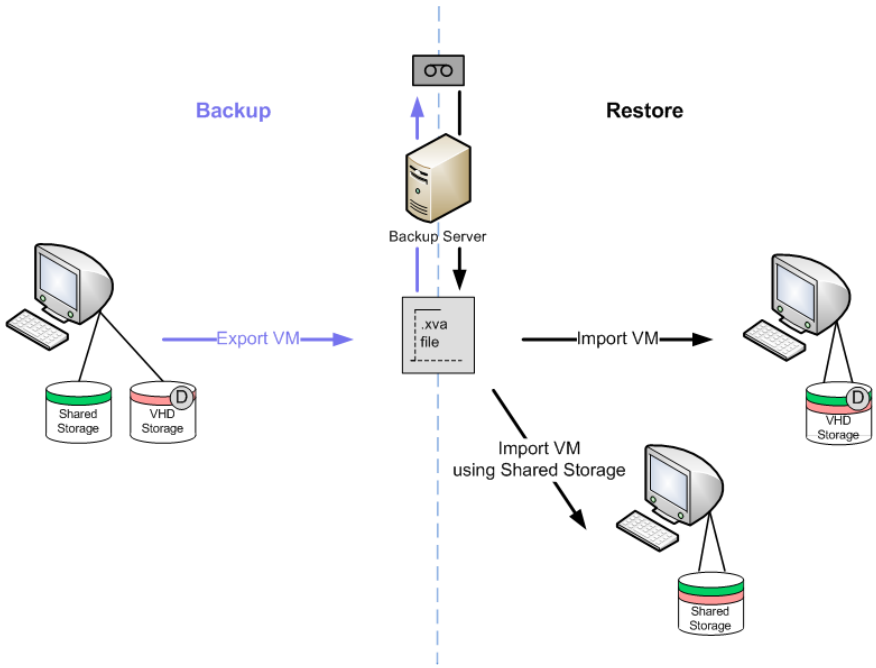

On the left side is a virtual machine with two disks: One disk uses the shared storage repository and the other the default storage repository.

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 86 (page 165).

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 86 Example layout

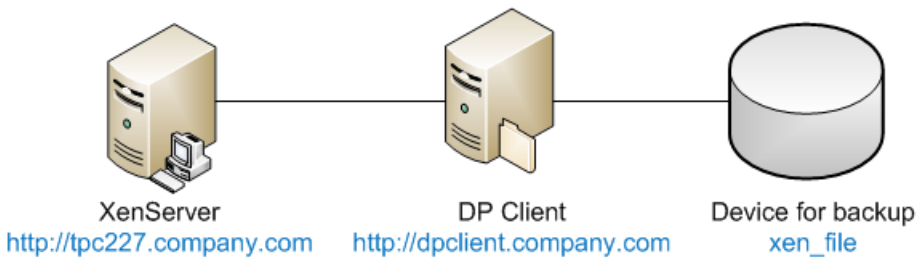

Installation of the integration

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 86 (page 165)) with the Data Protector Disk Agent installed and must be a member of a Data Protector Cell. For further details, refer to 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:

  - On Windows: *Data_Protector_home*\bin
  - On Linux: /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:

  - For Windows: http://pycurl.sourceforge.net/download/
  - For Linux: 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:

- On Windows: *Data_Protector_home*\bin\Python25\
- On Linux: /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 6.20, on the Data Protector 6.20 installation DVD-ROM in the directory:
  - For Windows: DP611_WIN_OTHER\Windows_other\Xen_sup
  - For Linux: DP611_SOL_LINUX/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 **Patches and utilities**.
  3. Select the Data Protector version with which you will be using the integration.

4. Select **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**
5. In the Software-Solutions section, click **Data Protector Release x.xx / 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:
   - On Windows: *`Data_Protector_home`*`\bin`
   - On Linux: `/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:
   - On Windows: `C:\tmp\backup\`
   - On Linux: `/tmp/backup/`

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, refer to "Updating configuration script DPxen_config.py for backup" (page 168) and "Updating configuration script DPxen_config.py for restore" (page 171).

  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, refer to "Integration concepts" (page 162).

`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 87 (page 169).

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 87 Obtaining the names of the virtual machines

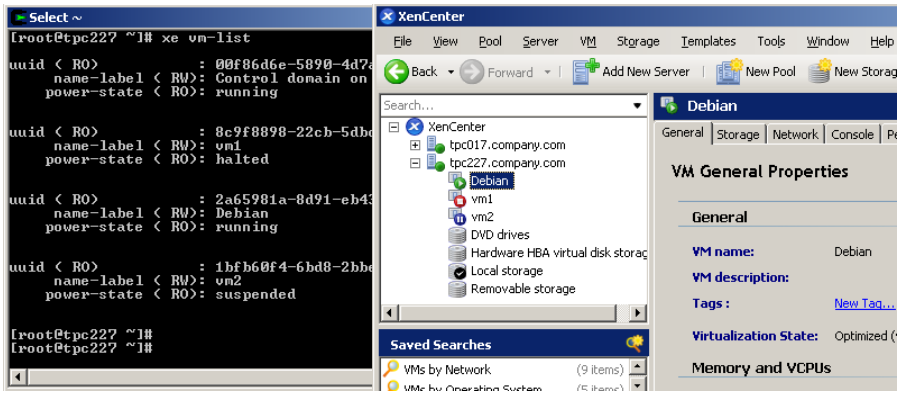

Once you have all the information for those parameters you can proceed to change the `DPxen_config.py` file or create a new configuration file.

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:
   - On Windows: `C:\tmp\backup`
   - On Linux: `/tmp/backup`

Figure 88 Temporary backup folder selection

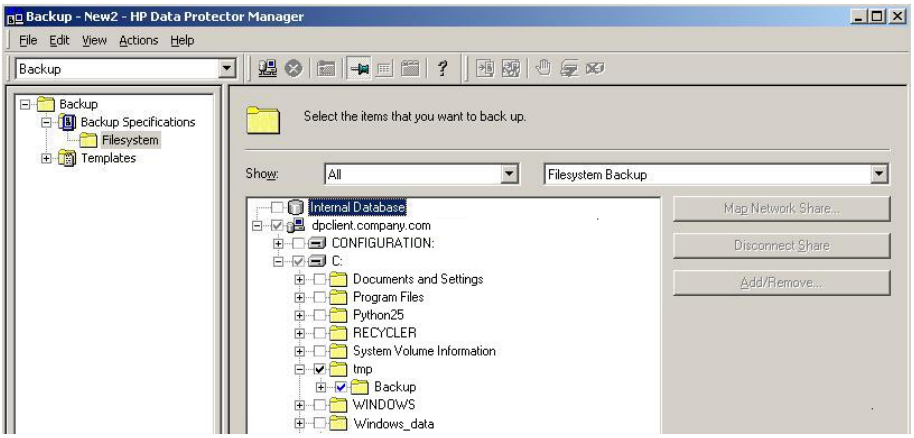

5. Click **Next**.
6. Select the backup device on which the backup will be performed and click **Next**.
7. In Backup Specification Options, click **Advanced**.
8. In the Backup Options window, make the following entries in the Pre-exec and Post-exec sections:

   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 165).

     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, refer to "Integration concepts" (page 162).

`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 89 (page 173)). Create a debug file.

```
Mode = "offline"
url = "http://tpc227.company.com"
username = "user"
password = "myPassword"
vmName = ["Debian", "vm2"]
debug = True
sr = "Shared Storage"
```

Figure 89 Available storage repositories

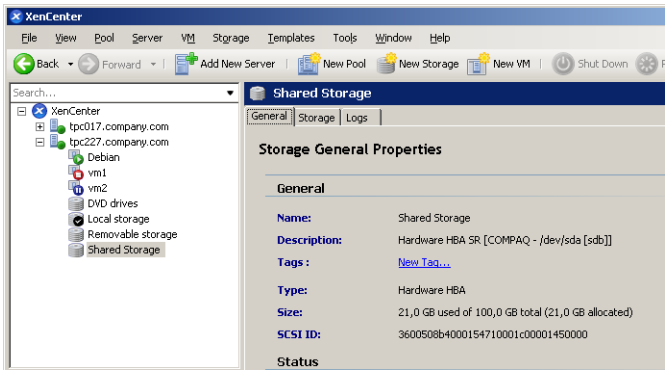

- Restore the Debian virtual machine, which was originally backed up using the `online` mode, into the "Removable storage" storage repository. Do not create a debug file.

```
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 90 (page 174).

Figure 90 Selection of files to restore

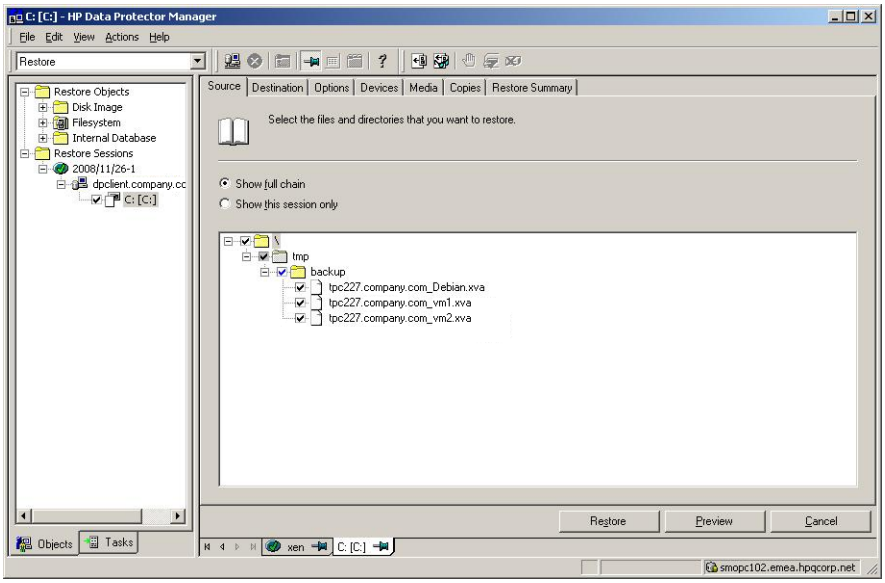

IMPORTANT: The Source tab shows the contents of the temporary backup folder when the backup session was performed. It does not show the current contents of the folder. It is important to ensure that the temporary backup folder is empty before performing a restore.

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, refer to 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, refer to http://www.hp.com/go/dataprotector.

Glossary

A

access rights — *See* user rights.

ACSLS — *(StorageTek specific term)* The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory — *(Windows specific term)* 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.

AES 256-bit encryption — 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.

AML — *(ADIC/GRAU specific term)* Automated Mixed-Media library.

AMU — *(ADIC/GRAU specific term)* Archive Management Unit.

application agent — A component needed on a client to back up or restore online database integrations.
See also Disk Agent.

application system — *(ZDB specific term)* 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.

archive logging — *(Lotus Domino Server specific term)* Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

archived redo log — *(Oracle specific term)* 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:

- 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.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also online redo log.

ASR set — 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 *Data_Protector_program_data*`\Config\Server\dr\asr` (Windows Server 2008), *Data_Protector_home*`\Config\Server\dr\asr` (other Windows systems), or `/etc/opt/omni/server/dr/asr` (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.

audit logs — Data files to which auditing information is stored.

audit report — User-readable output of auditing information created from data stored in audit log files.

auditing information — Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.

autochanger — *See* library.

autoloader — *See* library.

Automatic Storage Management (ASM) — *(Oracle specific term)* 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.

automigration *(VLS specific term)* The functionality that allows data backups to be first made to the VLS' virtual tapes and then migrated to physical tapes (one virtual tape emulating one physical tape) without using an intermediate backup application.
See also Virtual Library System (VLS) and virtual tape.

auxiliary disk 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

BACKINT *(SAP R/3 specific term)* 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.

backup API 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.

backup chain *See* restore chain.

backup device 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.

backup generation One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID 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.

backup object 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 (rawdisk).

A backup object is defined by:

- Client name: Hostname of the Data Protector client where the backup object resides.
- Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is located (drive on Windows and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.
- 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).
- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".

backup owner Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session 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.

backup set A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)* 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.

backup specification 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.

backup system *(ZDB specific term)* 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.

backup types *See* incremental backup, differential backup, transaction backup, full backup, and delta backup.

backup view 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.

BC *(EMC Symmetrix specific term)* Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also BCV.

BC Process *(EMC Symmetrix specific term)* A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also BCV.

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation 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.

Boolean operators The Boolean operators for the full text search functionality of the online 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.

boot volume/disk/partition 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.

BRARCHIVE *(SAP R/3 specific term)* 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.

BRBACKUP *(SAP R/3 specific term)* 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.

BRRESTORE *(SAP R/3 specific term)* An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

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.

BSM The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

C

CAP
(StorageTek specific term) Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also data protection.

CDB
The Catalog Database is a part of the IDB that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also MMDB.

CDF file
(UNIX 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.

cell
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.

Cell Manager
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.

centralized licensing
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.

Centralized Media Management Database (CMMDB)
See CMMDB.

Certificate Server
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.

Change Journal
(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.

Change Log Provider
(Windows specific term) A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

channel
(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:

- type 'disk'
- type 'sbt_tape'

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.

circular logging
(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.

client backup
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:

- 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.

- 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 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.

client or client system
Any system configured with any Data Protector functionality and configured in a cell.

cluster continuous replication
(Microsoft Exchange Server specific term) 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.

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.

See also Exchange Replication Service and local continuous replication.

cluster-aware application
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).

CMD script for Informix Server
(Informix Server specific term) 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.

CMMDB
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
See also MoM.

COM+ Class Registration Database
(Windows specific term) 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.

command device
(HP P9000 XP Disk Array Family specific term) 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.

Command View VLS
(VLS specific term) A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.
See also Virtual Library System (VLS).

command-line interface (CLI)
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency
See Disk Agent concurrency.

container
(HP P6000 EVA Disk Array Family specific term) Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file
(Oracle and SAP R/3 specific term) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set
(HP P6000 EVA Disk Array Family specific term) A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA.
See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

CRS
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 `root`.

CSM 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

data file *(Oracle and SAP R/3 specific term)* 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.

data protection 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.
See also catalog protection.

data replication (DR) group *(HP P6000 EVA Disk Array Family specific term)* 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.
See also copy set.

data stream Sequence of data transferred over the communication channel.

Data_Protector_home A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, and Windows Server 2008) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is `%ProgramFiles%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_program_data.

Data_Protector_program_data A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, and Windows Server 2008. Its default path is `%ProgramData%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_home.

database library 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.

database parallelism More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server 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.

Dbobject *(Informix Server specific term)* An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located on the Cell Manager in the directory *`Data_Protector_program_data`*`\db40` (Windows Server 2008), *`Data_Protector_home`*`\db40` (other Windows systems), or `/var/opt/omni/server/db40` (UNIX systems). You can create more DC directories and use a custom location. Up to 50 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

DCBF The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the filesystem settings.

delta backup A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also backup types.

device A physical unit which contains either just a drive or a more complex unit such as a library.

device chain 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.

device group *(EMC Symmetrix specific term)* 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.

device streaming 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.

DHCP server A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type.
See also incremental backup.

differential backup *(Microsoft SQL Server specific term)* A database backup that records only the data changes made to the database after the last full database backup.
See also backup types.

differential database backup A differential database backup records only those data changes made to the database after the last full database backup.

directory junction *(Windows specific term)* 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.

disaster recovery A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

disaster recovery operating system *See* DR OS.

Disk Agent 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.

Disk Agent concurrency The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk group *(Veritas Volume Manager specific term)* 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.

disk image (rawdisk) backup A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) 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.

disk quota 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.

disk staging 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).

distributed file media format 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.
See also virtual full backup.

Distributed File System (DFS) 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.

DMZ 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.

DNS server 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.

domain controller A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS 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.

drive 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.

drive index 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.

drive-based encryption Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the meta-data that is written to the medium.

E

EMC Symmetrix Agent A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file *(Informix Server specific term)* The Informix Server configuration file `ixbar.`*`server_id`* that resides in the directory *`INFORMIXDIR`*`/etc` (on Windows) or *`INFORMIXDIR`*`\etc` (on UNIX). *`INFORMIXDIR`* is the Informix Server home directory and *`server_id`* is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

encrypted control communication 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.

encryption key 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.

encryption KeyID-StoreID Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. `KeyID` identifies the key within the keystore. `StoreID` identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may several `StoreID`s used on the same Cell Manager.

enhanced incremental backup 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.

enterprise backup environment 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.
See also MoM.

Event Log (Data Protector Event Log) 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 *Data_Protector_program_data*\log\server\Ob2EventLog.txt (Windows Server 2008), *Data_Protector_home*\log\server\Ob2EventLog.txt (other 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.

Event Logs *(Windows specific term)* 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.

Exchange Replication Service *(Microsoft Exchange Server specific term)* 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.

exchanger Also referred to as SCSI Exchanger.
See also library.

exporting media A process that removes all data about backup sessions, such as systems, objects, and file names, 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.

Extensible Storage Engine (ESE) *(Microsoft Exchange Server specific term)* A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F

failover Transferring of the most important cluster data, called group (on Windows) or package (on UNIX) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

failover *(HP P6000 EVA Disk Array Family specific term)* 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.

FC bridge *See* Fibre Channel bridge.

Fibre Channel 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.

Fibre Channel bridge 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.

file depot A file containing the data from a backup to a file library device.

file jukebox device A device residing on disk consisting of multiple slots used to store file media.

file library device A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS) 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.

file tree walk *(Windows specific term)* The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

file version 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.

filesystem 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.

first-level mirror *(HP P9000 XP Disk Array Family specific term)* 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.

flash recovery area *(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).
See also recovery files.

fnames.dat The `fnames.dat` files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting 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.

free pool 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.

full backup A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also backup types.

full database backup 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.

full mailbox backup A full mailbox backup is a backup of the entire mailbox content.

full ZDB 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.
See also incremental ZDB.

G

global options file A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located on the Cell Manager in the directory *`Data_Protector_program_data`*`\Config\Server\Options` (Windows Server 2008), *`Data_Protector_home`*`\Config\Server\Options` (other Windows systems), or `/etc/opt/omni/server/options` (HP-UX, Solaris, and Linux systems).

group *(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.

GUI A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI that runs on Windows, Data Protector also provides a Java-based graphical user interface with the same look and feel, which runs on numerous platforms.

H

hard recovery *(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.

heartbeat 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.

Hierarchical Storage Management (HSM)
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.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file on the Cell Manager in the directory *Data_Protector_program_data*\Config\Server\holidays (Windows Server 2008), *Data_Protector_home*\Config\Server\holidays (other Windows systems), or /etc/opt/omni/server/Holidays (UNIX systems).

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP Business Copy (BC) P6000 EVA
(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.
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 StorageWorks P6000 EVA SMI-S Agent and HP StorageWorks 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 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 P6000 EVA SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses.
See also HP StorageWorks P6000 EVA SMI-S Agent and HP Command View (CV) EVA.

HP StorageWorks P6000 EVA SMI-S Agent
A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 EVA SMI-S Agent, the control over the array is established through HP SMI-S P6000 EVA Array provider, which directs communication between incoming requests and HP CV EVA.

See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP StorageWorks 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 Operations Manager
HP Operations Manager provides powerful capabilities for operations management of a large number of systems and applications in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for HP Operations Manager management servers on Windows, HP-UX, Solaris, and Linux. Earlier versions of HP Operations Manager were called IT/Operations, Operations Center, Vantage Point Operations, and OpenView Operations.

HP Operations Manager SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager, extending the managed domain. Through the Data Protector integration, which is implemented as an HP Operations Manager SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager.

I

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
The Data Protector Internal Database. IDB is an embedded database located on the Cell Manager and keeps information regarding which data was backed up, to which media it was backed up, how backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on.

IDB recovery file
An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

importing media
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.
See also exporting media.

incremental (re)-establish
(EMC Symmetrix specific term) 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.

incremental backup
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.
See also backup types.

incremental backup
(Microsoft Exchange Server specific term) 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.
See also backup types.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental restore
(EMC Symmetrix specific term) 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.

incremental ZDB 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.
See also full ZDB.

incremental1 mailbox backup An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet 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.

Information Store *(Microsoft Exchange Server specific term)* 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.
See also Key Management Service and Site Replication Service.

Informix Server *(Informix Server specific term)* Refers to Informix Dynamic Server.

initializing *See* formatting.

Installation Server 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.

instant recovery *(ZDB specific term)* 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.
See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

integration object A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Services (IIS) *(Windows specific term)* 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).

ISQL *(Sybase specific term)* A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

Java GUI Client The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities (the Cell Manager graphical user interface and the Manager-of-Managers (MoM) graphical user interface) and requires connection to the Java GUI Server to function.

Java GUI Server The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.

jukebox *See* library.

jukebox device 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

Key Management Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that provides encryption functionality for enhanced security.
See also Information Store and Site Replication Service.

keychain 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.

keystore All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS 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

LBO *(EMC Symmetrix specific term)* 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.

LDEV *(HP P9000 XP Disk Array Family specific term)* 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.
See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library 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.

lights-out operation or unattended operation 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.

LISTENER.ORA *(Oracle specific term)* An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing 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.

local and remote recovery 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.

local continuous replication *(Microsoft Exchange Server specific term)* 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.

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.

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.

See also cluster continuous replication and Exchange Replication Service.

lock name 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.

log_full shell script *(Informix Server UNIX specific term)* 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 `ALARMPROGRAM` configuration parameter defaults to the *`INFORMIXDIR`*`/etc/log_full.sh`, where *`INFORMIXDIR`* is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the `ALARMPROGRAM` configuration parameter to *`INFORMIXDIR`*`/etc/no_log.sh`.

logging level The logging level 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 influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files 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.

login ID *(Microsoft SQL Server specific term)* 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 syslogin.

login information to the Oracle Target Database *(Oracle and SAP R/3 specific term)* The format of the login information is *`user_name`*`/`*`password`*`@`*`service`*, where:

- *`user_name`* 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 SYSDBA or SYSOPER rights.
- *`password`* must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.
- *`service`* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database *(Oracle specific term)* The format of the login information to the Recovery (Oracle) Catalog Database is *`user_name`*`/`*`password`*`@`*`service`*, 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.

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.

media condition
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.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool
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.

media set
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.

media type
The physical type of media, such as DDS or DLT.

media usage policy
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`.

medium ID
A unique identifier assigned to a medium by Data Protector.

merging
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.
See also overwrite.

Microsoft Exchange Server
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.

Microsoft Management Console (MMC)
(Windows specific term) 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.

Microsoft SQL Server
A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy Service (VSS)
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.
See also shadow copy, shadow copy provider, replica, and writer.

mirror *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)*
See target volume.

mirror rotation *(HP P9000 XP Disk Array Family specific term)*
See replica set rotation.

mirror unit (MU) number
(HP P9000 XP Disk Array Family specific term) 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.
See also first-level mirror.

mirrorclone
(HP P6000 EVA Disk Array Family specific term) 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.

MMD
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.

MMDB
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 CDB.

MoM
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.

mount point
The access point in a directory structure for a disk or logical volume, for example `/opt` or `d:`. On UNIX, the mount points are displayed using the `bdf` or `df` command.

mount request
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.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping
(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.

O

OBDR capable device
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.

obdrindex.dat
See IDB recovery file.

object *See* backup object.

object consolidation 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.

object consolidation session 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.

object copy A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session 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.

object copying 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.

object ID *(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.

object mirror A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring 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.

object verification 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.

object verification session 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.

offline backup 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.

offline recovery 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. Recovery of the Cell Manager is always offline.

offline redo log *See* archived redo log.

ON-Bar *(Informix Server specific term)* 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:

- the `onbar` command
- Data Protector as the backup solution
- the XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

ONCONFIG *(Informix Server specific term)* 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 *`INFORMIXDIR`*`\etc` (on Windows) or *`INFORMIXDIR`*`/etc/` (on UNIX).

online backup 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.

online recovery Online recovery is performed when Cell Manager is accessible. In this case, most of the Data Protector] functionalities are available (Cell Manager runs the session, restore sessions are logged in the IDB, you can monitor the restore progress using the GUI, and so on).

online redo log *(Oracle specific term)* 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.

Oracle Data Guard *(Oracle specific term)* 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.

Oracle instance *(Oracle specific term)* Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)* A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired `ORACLE_SID`. The `ORACLE_SID` is included in the CONNECT DATA 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 objects, by default the owner is the user who starts the operation, unless a different owner is specified in the copy or consolidation specification.

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 Server 2008), *Data_Protector_home*\Config\Server\dr\p1s (other 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 StorageWorks 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 device 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.

phase 0 of disaster recovery Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.

phase 3 of disaster recovery Restoration of user and application data.

physical device A physical unit that contains either a drive or a more complex unit such as a library.

post-exec 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 and as shell scripts on UNIX.
See also pre-exec.

pre- and post-exec commands 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 and as shell scripts on UNIX.

pre-exec 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 and as shell scripts on UNIX.
See also post-exec.

prealloc list A subset of media in a media pool that specifies the order in which media are used for backup.

primary volume (P-VOL) *(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).

protection *See* data protection and also catalog protection.

public folder store *(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.

public/private backed up data When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID Redundant Array of Independent Disks.

RAID Manager Library *(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 StorageWorks P9000 XP Agent.

RAID Manager P9000 XP *(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.

rawdisk backup *See* disk image backup.

RCU *See* Remote Control Unit (RCU).

RDBMS Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)* 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.

RDS The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)* 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:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts

Recovery Catalog Database *(Oracle specific term)* An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files *(Oracle specific term)* 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.
See also flash recovery area.

Recovery Manager (RMAN) *(Oracle specific term)* 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.

RecoveryInfo 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.

recycle or unprotect 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.

redo log *(Oracle specific term)* 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.

Remote Control Unit (RCU) *(HP P9000 XP Disk Array Family specific term)* 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.

Removable Storage Management Database *(Windows specific term)* 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.

reparse point *(Windows specific term)* 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.

replica *(ZDB specific term)* 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 UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.
See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also replica and replica set rotation.

replica set rotation *(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.

restore chain All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.

restore session A process that copies data from backup media to a client.

resync mode *(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.

RMAN *(Oracle specific term)* *See* Recovery Manager.

RSM The Data Protector Restore Session Manager controls restore and object verification sessions. This process always runs on the Cell Manager system.

RSM *(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.

S

SAPDBA *(SAP R/3 specific term)* An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scanning 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.

Scheduler A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

secondary volume (S-VOL) *(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).

session *See* backup session, media management session, and restore session.

session ID 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.

session key 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.

shadow copy *(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 volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.
See also Microsoft Volume Shadow Copy Service and replica.

shadow copy provider *(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.

shadow copy set *(Microsoft VSS specific term)* A collection of shadow copies created at the same point in time.
See also shadow copy and replica set.

shared disks 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.

SIBF The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also Information Store and Key Management Service.

slot 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.

smart copy *(VLS specific term)* A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management.
See also Virtual Library System (VLS).

smart copy pool *(VLS specific term)* A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library.
See also Virtual Library System (VLS) and smart copy.

SMB *See* split mirror backup.

SMBF 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.

SMI-S Agent (SMISA) *See* HP StorageWorks P6000 EVA SMI-S Agent.

snapshot *(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions 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.
See also replica and snapshot creation.

snapshot backup *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation *(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions 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.
See also snapshot.

source (R1) device *(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.
See also target (R2) device.

source volume *(ZDB specific term)* A storage volume containing data to be replicated.

sparse file 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.

split mirror *(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.
See also replica and split mirror creation.

split mirror backup *(EMC Symmetrix specific term)* *See* ZDB to tape.

split mirror backup *(HP P9000 XP Disk Array Family specific term)* *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation *(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.

split mirror restore *(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.

sqlhosts file or registry *(Informix Server specific term)* An Informix Server connectivity information file (on UNIX) or registry (on Windows) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file *(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.

SRDF *(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.

SSE Agent (SSEA) *See* HP StorageWorks P9000 XP Agent.

sst.conf file The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Sun 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.

st.conf file The file `/kernel/drv/st.conf` 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.

stackers 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.

standalone file device A file device is a file in a specified directory to which you back up data.

Storage Group *(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.

storage volume *(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.

StorageTek ACS library *(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.

switchover *See* failover.

Sybase Backup Server API *(Sybase specific term)* 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.

Sybase SQL Server *(Sybase specific term)* 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.

SYMA *See* EMC Symmetrix Agent.

synthetic backup 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.

synthetic full backup 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.

System Backup to Tape *(Oracle specific term)* An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)* The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

System Recovery Data file *See* SRD file.

System State *(Windows specific term)* 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.

system volume/disk/partition 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.

SysVol *(Windows specific term)* A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

T

tablespace 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.

tapeless backup *(ZDB specific term)* *See* ZDB to disk.

target (R2) device *(EMC Symmetrix specific term)* 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.

target database *(Oracle specific term)* In RMAN, the target database is the database that you are backing up or restoring.

target system *(disaster recovery specific term)* 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.

target volume *(ZDB specific term)* A storage volume to which data is replicated.

Terminal Services *(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.

thread *(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.

TimeFinder *(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).

TLU Tape Library Unit.

TNSNAMES.ORA *(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.

transaction 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.

transaction backup 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.

transaction backup *(Sybase and SQL specific term)* A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup 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.

transaction log files Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction log table *(Sybase specific term)* A system table in which all changes to the database are automatically recorded.

transaction logs *(Data Protector specific term)* Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transportable snapshot *(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).

TSANDS.CFG file *(Novell NetWare specific term)* A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the `SYS:SYSTEM\TSA` directory on the server where `TSANDS.NLM` is loaded.

U

UIProxy The Java GUI Server (`UIProxy` service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.

unattended operation *See* lights-out operation.

user account (Data Protector user account) 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.

User Account Control (UAC) A security component in Windows Vista, Windows 7, and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas 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.

user group 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.

user profile *(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.

user rights 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.

user_restrictions file 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 *admin* and *operator*.

V

vaulting media 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.

verify 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.

Virtual Controller Software (VCS) *(HP P6000 EVA Disk Array Family specific term)* The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers.
See also HP Command View (CV) EVA.

Virtual Device Interface *(Microsoft SQL Server specific term)* This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk *(HP P6000 EVA Disk Array Family specific term)* 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.
See also source volume and target volume.

virtual full backup 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.

Virtual Library System (VLS) A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

virtual server 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.

virtual tape *(VLS specific term)* 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.
See also Virtual Library System (VLS) and Virtual Tape Library (VTL).

Virtual Tape Library (VTL) *(VLS specific term)* An emulated tape library that provides the functionality of traditional tape-based storage.
See also Virtual Library System (VLS).

VMware management client *(VMware (Legacy) integration specific term)* 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).

volser *(ADIC and STK specific term)* 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.

volume group 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.

volume mountpoint *(Windows specific term)* 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.

Volume Shadow Copy Service *See* Microsoft Volume Shadow Copy Service (VSS).

VSS *See* Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode *(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 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.

VxFS Veritas Journal Filesystem.

VxVM (Veritas Volume Manager) 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

Wake ONLAN Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting 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.

wildcard character 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.

Windows configuration backup 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.

Windows Registry A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server 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.

writer *(Microsoft VSS specific term)* 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

XBSA interface *(Informix Server specific term)* ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z

ZDB *See* zero downtime backup (ZDB).

ZDB database *(ZDB specific term)* 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).

ZDB to disk *(ZDB specific term)* 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.

ZDB to disk+tape *(ZDB specific term)* 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.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape *(ZDB specific term)* 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.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB) 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.
See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index

A

architecture
  Microsoft Hyper-V, Virtual Environment integration, , 133
  Virtual Environment integration, 20
  VMware (Legacy) integration, 78
audience, 10

B

backing up Microsoft Hyper-V
  backup options, 150
  backup specification, modifying, 150
  backup specifications, creating, 147
  scheduling backups, 151
  scheduling backups, example, 152
  starting backups, 152
backing up Microsoft Hyper-V, Virtual Environment integration, 147–153
backing up Virtual Environment
  backup methods, 29
backing up VMware
  backup methods, 138
  backup options, 57
  backup specification, modifying, 57
  backup specifications, creating, 54
  previewing backups, 59
  scheduling backups, 58
  scheduling backups, example, 59
  starting backups, 60
backing up VMware (Legacy), 104–110
  backup methods, 83
  backup options, 107
  backup specification, modifying, 107
  backup specifications, creating, 104
  backup types, 93
  differential backups, 93
  full backups, 93
  incremental backups, 93
  previewing backups, 108
  scheduling backups, 107
  scheduling backups, example, 107
  starting backups, 109
backing up VMware, Virtual Environment integration, 54–61
  differential backups, 31
  full backups, 31
  incremental backups, 31
backup
  XenServer integration, 168
backup methods
  Microsoft Hyper-V, Virtual Environment integration, , 138
  Virtual Environment integration, 29
  VMware (Legacy) integration, 83
backup options
  Microsoft Hyper-V integration, 150
  Virtual Environment integration, 57
  VMware (Legacy) integration, 107
backup specification
  XenServer integration, 170
backup specifications, creating
  Microsoft Hyper-V integration, 147
  Virtual Environment integration, 54
  VMware (Legacy) integration, 104
backup specifications, modifying
  Microsoft Hyper-V integration, 150
  Virtual Environment integration, 57
  VMware (Legacy) integration, 107
backup specifications, scheduling
  Microsoft Hyper-V integration, 151
  Virtual Environment integration, 58
  VMware (Legacy) integration, 107
backup types
  VMware (Legacy) integration, 93

C

checking configuration
  Virtual Environment integration, 48
  VMware (Legacy) integration, 103
concepts
  Virtual Environment integration, 28
  VMware (Legacy) integration, 77
  XenServer integration, 162
configuring Virtual Environment integration, 43–54
configuring Virtual Environment integration, Hyper-V, 143–147
configuring VMware
  checking configuration, 48
configuring VMware (Legacy), 96–104
  checking configuration, 103
conventions
  document, 15
creating backup specifications
  Microsoft Hyper-V integration, 147
  Virtual Environment integration, 54
  VMware (Legacy) integration, 104

D

differential backups
  VMware (Legacy) integration, 93
  VMware, Virtual Environment integration, 31
document
  conventions, 15
  related documentation, 10
documentation
  HP website, 10
  providing feedback, 17

E

examples, Microsoft Hyper-V integration
  scheduling backups, 152

examples, Virtual Environment integration
scheduling backups, 59
starting interactive backups, 60, 153
examples, VMware (Legacy) integration
scheduling backups, 107
starting interactive backups, 109

F
full backups
VMware (Legacy) integration, 93
VMware, Virtual Environment integration, 31

H
help
obtaining, 16
HP
technical support, 16

I
incremental backups
VMware (Legacy) integration, 93
VMware, Virtual Environment integration, 31
installing XenServer integration, 165
interactive backups
Microsoft Hyper-V integration, 152
Virtual Environment integration, 60
VMware (Legacy) integration, 109
introduction
Microsoft Hyper-V, Virtual Environment integration, 132
Virtual Environment integration, 20
VMware (Legacy) integration, 77
VSS integration, 138

M
main integration components
XenServer integration, 165
Microsoft Hyper-V backup
backup options, 150
backup specification, modifying, 150
backup specifications, creating, 147
scheduling backups, 151
scheduling backups, example, 152
starting backups, 152
Microsoft Hyper-V backup, Virtual Environment integration, 147–153
Microsoft Hyper-V integration
monitoring sessions, 157
Microsoft Hyper-V restore
finding information, 153
restore options, 156
using another device, 156
using GUI, 154
Microsoft Hyper-V restore, Virtual Environment integration, 153–157
Microsoft Hyper-V troubleshooting, Virtual Environment integration, 157–160
Microsoft Hyper-V, Virtual Environment integration
introduction, 132
modifying backup specifications
Microsoft Hyper-V integration, 150
Virtual Environment integration, 57
VMware (Legacy) integration, 107
monitoring sessions
Microsoft Hyper-V integration, 157
VMware (Legacy) integration, 128
VMware, Virtual Environment integration, 72

O
offline backup
XenServer integration, 162
online backup
XenServer integration, 162
online backups
Microsoft Hyper-V, Virtual Environment integration, 132
Virtual Environment integration, 20
VMware (Legacy) integration, 77

P
previewing backups
Virtual Environment integration, 59
VMware (Legacy) integration, 108

R
related documentation, 10
restore
XenServer integration, 171
restore from offline backup
XenServer integration, 163
restore from online backup
XenServer integration, 163
restore options
Microsoft Hyper-V integration, 156
Virtual Environment integration, 65
VMware (Legacy) integration, 116
restore processes
XenServer integration, 163
restoring Microsoft Hyper-V
finding information, 153
restoring Microsoft Hyper-V Server
restore options, 156
using another device, 156
using GUI, 154
restoring Microsoft Hyper-V, Virtual Environment integration, 153–157
restoring VMware
finding information, 61
restore options, 65
using another device, 72
using CLI, 66, 155
using GUI, 62
restoring VMware (Legacy), 110–128
finding information, 110
restore options, 116
using another device, 128
using CLI, 116
using GUI, 111
restoring VMware, Virtual Environment integration, 61–72

S

scheduling backups
- Microsoft Hyper-V integration, 151
- Virtual Environment integration, 58
- VMware (Legacy) integration, 107

starting backups
- Microsoft Hyper-V integration, 152
- Virtual Environment integration, 60
- VMware (Legacy) integration, 109

Subscriber's Choice, HP, 16

T

technical support
- HP, 16
- service locator website, 16

troubleshooting Virtual Environment integration, Microsoft Hyper-V, 157–160

troubleshooting Virtual Environment integration, VMware, 72–76

troubleshooting VMware (Legacy), 128–130

V

VEPA
- Virtual Environment integration, 28

Virtual Environment backup
- backup methods, 29

Virtual Environment integration
- architecture, 20
- backup, 54
- concepts, 28
- configuration, 43–54
- introduction, 20
- restore, 61–72

Virtual Environment integration configuration, 43–54

Virtual Environment integration configuration, Hyper-V, 143–147

Virtual Environment integration troubleshooting, VMware, 72–76

Virtual Environment integration, Hyper-V
- configuration, 143–147

Virtual Environment integration, Microsoft Hyper-V
- architecture, 133
- backup, 147–153
- restore, 153–157
- troubleshooting, 157–160

Virtual Environment integration, VMware
- troubleshooting, 72–76

VMware (Legacy) backup, 104–110
- backup methods, 83
- backup options, 107
- backup specification, modifying, 107
- backup specifications, creating, 104
- backup types, 93
- differential backups, 93
- full backups, 93
- incremental backups, 93
- previewing backups, 108
- scheduling backups, 107
- scheduling backups, example, 107
- starting backups, 109

VMware (Legacy) configuration, 96–104
- checking configuration, 103

VMware (Legacy) integration
- architecture, 78
- backup, 104–110
- concepts, 77
- configuration, 96–104
- introduction, 77
- monitoring sessions, 128
- restore, 110–128
- troubleshooting, 128–130

VMware (Legacy) restore, 110–128
- finding information, 110
- restore options, 116
- using another device, 128
- using CLI, 116
- using GUI, 111

VMware (Legacy) troubleshooting, 128–130

VMware backup
- backup methods, 138
- backup options, 57
- backup specification, modifying, 57
- backup specifications, creating, 54
- previewing backups, 59
- scheduling backups, 58
- scheduling backups, example, 59
- starting backups, 60

VMware backup, Virtual Environment integration, 54–61
- differential backups, 31
- full backups, 31
- incremental backups, 31

VMware configuration
- checking configuration, 48

VMware management client
- VMware (Legacy) integration, 97

VMware restore
- finding information, 61
- restore options, 65
- using another device, 72
- using CLI, 66, 155
- using GUI, 62

VMware restore, Virtual Environment integration, 61–72

VMware, Virtual Environment integration
- monitoring sessions, 72

Volume Shadow Copy Service
- VSS integration, 138

VSS integration
- introduction, 138

W

websites
- HP , 16
- HP Subscriber's Choice for Business, 16
- product manuals, 10

X

XenServer integration
- backup, 168

backup specification, 170
concepts, 162
installation, 165
integration scripts, 167
main integration components, 165
offline backup, 162
online backup, 162
restore, 171
restore from offline backup, 163
restore from online backup, 163
restore processes, 163
specifying restore, 173
updating DPxen_config.py for backup, 168
updating DPxen_config.py for restore, 171