

HP Data Protector 6.20

Granular Recovery Extension User Guide for VMware vSphere

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.....	5
About the guide.....	6
Intended audience.....	6
Documentation set.....	6
Guides.....	6
Online Help.....	8
Documentation map.....	9
Abbreviations.....	9
Map.....	9
Integrations.....	10
Document conventions and symbols.....	11
Data Protector graphical user interface.....	11
General information.....	12
HP technical support.....	12
Subscription service.....	12
HP websites.....	12
Documentation feedback.....	13
1 Introduction.....	14
Backup.....	14
Restore and recovery.....	14
2 Installation.....	16
Prerequisites.....	16
Considerations.....	17
HP Data Protector Granular Recovery Extension environment.....	17
VMware vCenter Server system.....	18
Component on the VMware vCenter Server system.....	18
Mount proxy system.....	19
Components on the mount proxy system.....	19
HP Data Protector Granular Recovery Extension users.....	20
3 Configuration.....	21
Meeting Data Protector configuration requirements for Granular Recovery Extension.....	21
VMware vCenter Server system.....	22
Importing and configuring VMware clients.....	22
Mount proxy system.....	23
Configuring a Granular Recovery Extension VMware user group.....	24
Configuring Granular Recovery Extension for a VMware vSphere user group.....	24
Adding users to the GRE VMware group.....	25
Adding an Inet user account to the Data Protector Administrator group.....	27
Configuring Granular Recovery Extension Administrators using HP Data Protector.....	28
Configuring your systems for VMware vSphere purposes.....	29
Configuring Windows Firewall exceptions.....	29
Setting VMware Granular Recovery Extension Agent in Firewall exceptions.....	29
Configuring vCenter Server options.....	29
4 Backup.....	31
Considerations.....	31
5 Restore and recovery.....	32
Considerations.....	32
Restore and recovery flow.....	33

Operator's tasks.....	35
Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI.....	35
Requesting restores.....	37
Monitoring restore requests.....	39
Recovering virtual machines.....	40
Administrator's tasks.....	43
Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI.....	43
Configuring mount proxy systems.....	44
Handling restore requests.....	44
Triggering restore sessions.....	45
Rejecting restore requests.....	46
Monitoring restore requests.....	46
Stopping restore sessions.....	47
Removing virtual machine disks.....	47
6 Troubleshooting.....	49
Before you begin.....	49
Debugging.....	50
Enabling debugging option.....	50
Known issues and workarounds.....	51
Mounting virtual machine disks.....	51
Issues after removing this extension.....	51
VMware VirtualCenter Management Webservices service is not running.....	52
Remote installation of VMware Granular Recovery Extension Web Plug-In ends unexpectedly.....	52
VMware Granular Recovery Extension tab is missing.....	53
VMware Granular Recovery Extension tab is missing with vCenter Server plug-in disabled.....	53
HP Data Protector cannot add any Granular Recovery Extension component.....	54
Local installation workaround.....	55
Monitor displays request removed by administrator.....	55
Overwritten files issues.....	56
Opening the extension with script errors.....	56
Glossary.....	58
Index.....	88

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 the guide

Intended audience

This guide is intended for VMware vCenter Server users and Data Protector backup administrators responsible for planning, setting up, and maintaining backups and recovery of VMware Virtual Environment. It assumes you are familiar with:

- Basic Data Protector functionality
- VMware administration (vCenter and vSphere).

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.

								Integration Guides							ZDB		GRE		MO					
	Help	GS	Concepts	Install	Trouble	DR	PA	MS	O/S	IBM	Var	VSS	VirtEnv	OMU	OMW	Concept	Admin	IG	SPS	VMware	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 11)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	<ul style="list-style-type: none"> 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
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none"> File and directory names System output Code Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none"> 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.

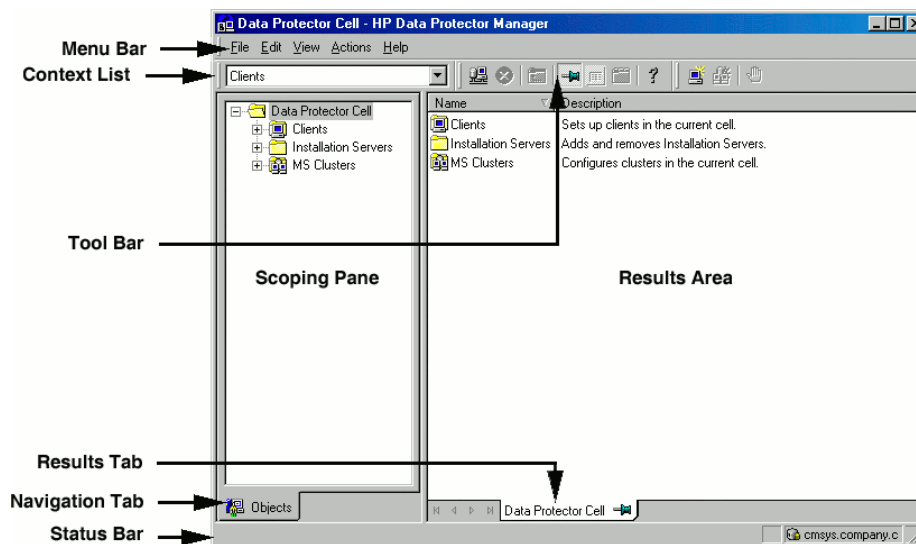


TIP: Provides helpful hints and shortcuts.

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.

1 Introduction

This chapter describes Granular Recovery Extension for VMware vSphere.

Backup

The HP Data Protector Granular Recovery Extension for VMware vSphere uses backups performed by the Data Protector Virtual Environment Integration component to perform restore and recovery of the single items of VMware virtual machines.

Restore and recovery

The benefits of this extension are the following:

- **Recovery granularity**

The smallest item that can be restored is a disk, but from that restore, operators can recover individual files. This ability to recover items at file level helps minimize the use of disk space when dealing with large disks (100 GB or more). You can restore a set of virtual disks, mount them, and then only recover what you need.

- **Space-saving approach**

Recovering items at file level is useful, especially when recovering large sized disks, avoiding extensive disk space usage. Only restoring a set of files of virtual disks, mounting these virtual machine disks and then recovering what is essential on the target system. You specify just the necessary data, specific files or folders to the recovery session, this avoids excessive use of disk space.

- **Integration into VMware vCenter Server**

The extension is fully integrated into VMware vCenter Server, and has the same look and feel. This allows operators to request a restore and then recover individual items from virtual disks. Operators with Data Protector `start restore` right can also restore data by using Data Protector.

- **Operator-directed restore and recovery**

Operators (VMware vCenter users) can be allowed to restore virtual disks and recover their own files. This enables operators to:

- request restore (the restore is handled by Granular Recovery Extension Administrators).
- recover single items independently or with minimal interference from Data Protector backup administrators (no need to be Data Protector backup administrators themselves or have Data Protector administrator rights).

For detailed description of user roles and tasks, see [“HP Data Protector Granular Recovery Extension users”](#) (page 20).

For detailed configuration steps, see [“Configuring Granular Recovery Extension Administrators using HP Data Protector”](#) (page 28).

- **No Data Protector related expertise, low learning curve**

Operators without any Data Protector knowledge, or expertise can request restores. After the Granular Recovery Extension administrator handles and performs the restore, operators can start recovery sessions. This way, operators can perform restore and recovery without ever starting Data Protector or planning any backups, or dealing with devices, managing disk spaces and so on.

- **Clear overview of restore and recovery objects**

HP Data Protector Granular Recovery Extension for VMware vSphere enables selection of objects at file level, exactly the data needed. This increases the clarity, optimizing the overview of what is being restored or recovered.

- **Quickness and ease of use**

This extension is quicker and easier than using whole virtual machine restore and recovery solution. For example, it avoids additional configuration steps: start two virtual machines, copy the whole disk, configure network shares, IP names and addresses, and so on.

- **Monitoring restore progress**

HP Data Protector Granular Recovery Extension for VMware vSphere provides monitoring functionality. Operators without Data Protector backup administrators rights can monitor multiple restore requests, and recovery requests. Operators with Data Protector backup administrators rights handles these requests, have privileges to monitor multiple restore sessions, reject restore requests.

- **Recovery to different locations**

Granular Recovery Extension Administrators can specify different locations for the restore and recovery of virtual disks.

- Items can be first restored to different temporary locations:
 - different specified mount proxies, and different locations on a mount proxy.
If there is a shortage of disk space on a mount proxy, they can add new virtual machine.
- Finally items can be recovered to the following target locations:
 - virtual machines in online mode (with the same operating systems as the original backup).
The virtual machine disks must be online for the recovery to be successful.

- **Security**

The security of the virtual machines is based on the current VMware vCenter Server permissions. VMware virtual machine users can only access disks to which they are granted access in the VMware vCenter Server system. Access to secured data is granted to a limited number of users, so only certain users (Granular Recovery Extension administrators) can access this extension in the VMware vCenter Administration component.

In addition, access is restricted to requested virtual machines. This means that operators can mount and browse only disks on virtual machines to which they are granted access rights.

- **Users identification**

When using the extension, users' rights are checked, identifying whether they are Granular Recovery Extension administrators or operators. In addition, virtual machine operating systems credentials, user name and password must be specified before a recovery session can begin.

- **Specific access points**

There are two entry points to this extension (according to the roles):

- For administrative tasks, locate the HP Data Protector Granular Recovery Extension in the VMware vCenter Administration components.
- For operational tasks, locate the HP Data Protector Granular Recovery Extension tab in the VMware vCenter inventory components, under VMs and Templates.

2 Installation

This chapter describes how to install HP Data Protector Granular Recovery Extension for VMware vSphere.

Prerequisites

Data Protector :

- For detailed Data Protector installation, see the *HP Data Protector Installation and Licensing Guide*.

Data Protector Virtual Environment integration:

- Install and configure the Data Protector Virtual Environment integration as described in the *HP Data Protector Integration Guide for Virtualization Environments*, the chapter about the Data Protector VMware Virtual Infrastructure integration.

For detailed procedure, see the “[Importing and configuring VMware clients](#)” (page 22).

- Install VMware tools 4.0 or newer version on each virtual machine on which you plan to perform restore sessions.

For details on installing VMware tools on the virtual machine, see the vSphere Client Help, index: Upgrading VMware Tools. To install the VMware tools on a physical disk, download the VMware tools package from:

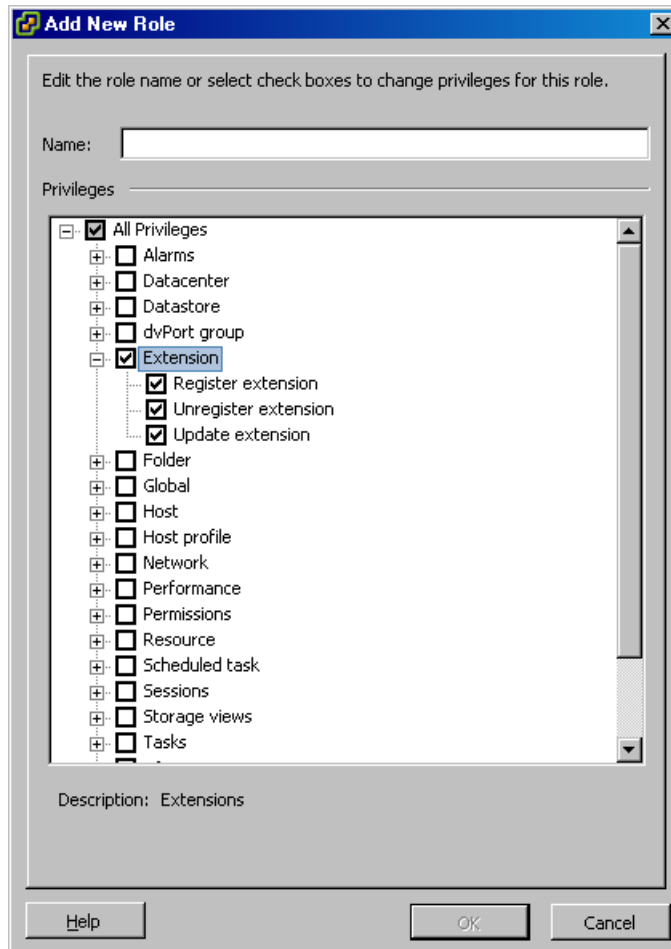
<http://www.vmware.com/download>

vCenter:

- **Non-administrator installation vCenter**

To be able to install VMware Granular Extension Web Plug-In component as a non-administrator installation vCenter user, ensure you have the vCenter Extension privilege set, with the Register extension, Unregister extension, and Update extension roles assigned.

Figure 2 Granting the Extension privileges to a vCenter user



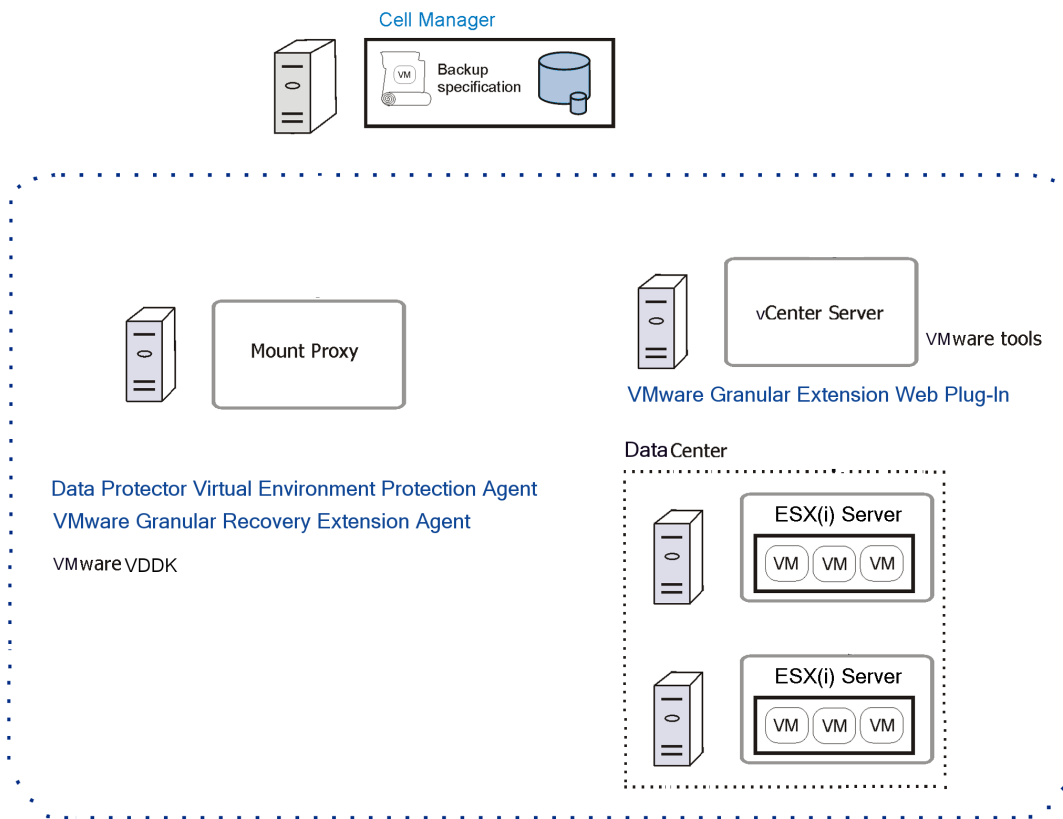
Considerations

- Remote installation of the HP Data Protector Granular Recovery Extension for VMware vSphere is recommended. If you cannot, for any reason, install this extension remotely, see [“Local installation workaround”](#) (page 55).

HP Data Protector Granular Recovery Extension environment

In the [“Installing HP Data Protector Granular Recovery Extension”](#) (page 18), the HP Data Protector components are colored blue, and VMware components are black.

Figure 3 Installing HP Data Protector Granular Recovery Extension



HP Data Protector Granular Recovery Extension restores data by means of the Data Protector Virtual Environment integration; this extension is a recovery solution. To use the extension, install and configure the following systems:

- Data Protector cell and client, imported with the Virtual Environment Integration agent enabled on it
- VMware vCenter Server
- Mount proxy

For detailed Data Protector concepts, see the *HP Data Protector Concepts Guide*.

VMware vCenter Server system

HP Data Protector Granular Recovery Extension is integrated into the VMware vCenter Server. The access to virtual machines is made through vCenter vSphere Client. This extension alters the VMware vCenter Server interface by adding new tabs and pages with the HP Data Protector Granular Recovery Extension icon.

Component on the VMware vCenter Server system

The web plug-in component must be installed on VMware vCenter Server system to enable communication with HP Data Protector Granular Recovery Extension. To use the HP Data Protector Granular Recovery Extension integrated into the vCenter Server, install the following Data Protector component on the vCenter Server system:

- VMware Granular Extension Web Plug-In

For detailed configuration procedure, see [“Importing and configuring VMware clients”](#) (page 22).

Mount proxy system

HP Data Protector Granular Recovery Extension for VMware vSphere requires a mount proxy system as a temporary restore location between the original location and target location on the VMware vCenter Server system. Any supported system, also a virtual machine, can be used as a mount proxy. It is recommended to allocate a dedicated system to be the mount proxy.

Once the restore session is finished the files of virtual machines disks are located in a folder marked with the request ID. This folder contains a collection of virtual machine files, for example, vdmk, tmp files and so on. The virtual machine disks are not mounted yet. The mount session starts when VMware vCenter users browse through files with this extension integrated in the vCenter environment.

The default retention period of the restored collections of files of virtual machines is 30 days. The Granular Recovery Extension administrator can change this setting.

The drive letter is dismounted automatically after the page for the selection of files is closed, in other words after VMware vCenter users browse through files, select files and start the recovery session. This dismount is necessary because only 23 drive letters are available.

After that, only the selected data is recovered to the specified virtual machine disk target location in the VMware vCenter environment.

The mount proxy system should have sufficient disk space for the restored data. The Granular Recovery Extension administrator can adjust the disk space on demand by attaching additional disks or by adding another mount proxy.

For detailed description of data flow, see [“Restore and recovery flow”](#) (page 33).

For detailed list of tasks and description of user roles, see [“HP Data Protector Granular Recovery Extension users”](#) (page 20).

Components on the mount proxy system

To use the mount proxy system for HP Data Protector Granular Recovery Extension, install the following Data Protector components on the mount proxy:

- Virtual Environment Integration
- VMware Granular Recovery Extension Agent

For detailed configuration procedure, see [“Mount proxy system”](#) (page 23).

HP Data Protector Granular Recovery Extension users

There is a distinction between two user roles:

- **Administrators**

The administrators are VMware vCenter Server users that have the Data Protector `start restore` user right assigned. These users can perform additional administrative tasks, for detailed procedures, see:

- [“Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI” \(page 43\).](#)
- [“Configuring mount proxy systems” \(page 44\).](#)
- [“Handling restore requests” \(page 44\).](#)
- [“Triggering restore sessions” \(page 45\).](#)
- [“Rejecting restore requests” \(page 46\).](#)
- [“Monitoring restore requests” \(page 46\).](#)
- [“Removing virtual machine disks” \(page 47\).](#)
- [“Stopping restore sessions” \(page 47\).](#)

- **Operators**

The operators are VMware vCenter Server users. These users can perform operational tasks, for detailed procedures, see:

- [“Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI” \(page 35\).](#)
- [“Requesting restores” \(page 37\).](#)
- [“Monitoring restore requests” \(page 39\).](#)
- [“Recovering virtual machines” \(page 40\).](#)

For details on configuring Data Protector `start restore` user rights, see [“Configuring Granular Recovery Extension Administrators using HP Data Protector” \(page 28\).](#)

3 Configuration

This section describes the configuration steps that you need to follow.

Meeting Data Protector configuration requirements for Granular Recovery Extension

NOTE: If any Data Protector component is already installed on the vCenter Server, you can skip this step.

Install a Data Protector component. During the installation procedure, select one of the HP Data Protector components, for example **Disk Agent**.

VMware vCenter Server system

Importing and configuring VMware clients

With the Data Protector Virtual Environment integration, it is not necessary to install any Data Protector components on VMware clients (vCenter Server systems), but to make them Data Protector clients, they must be imported into the Data Protector cell.

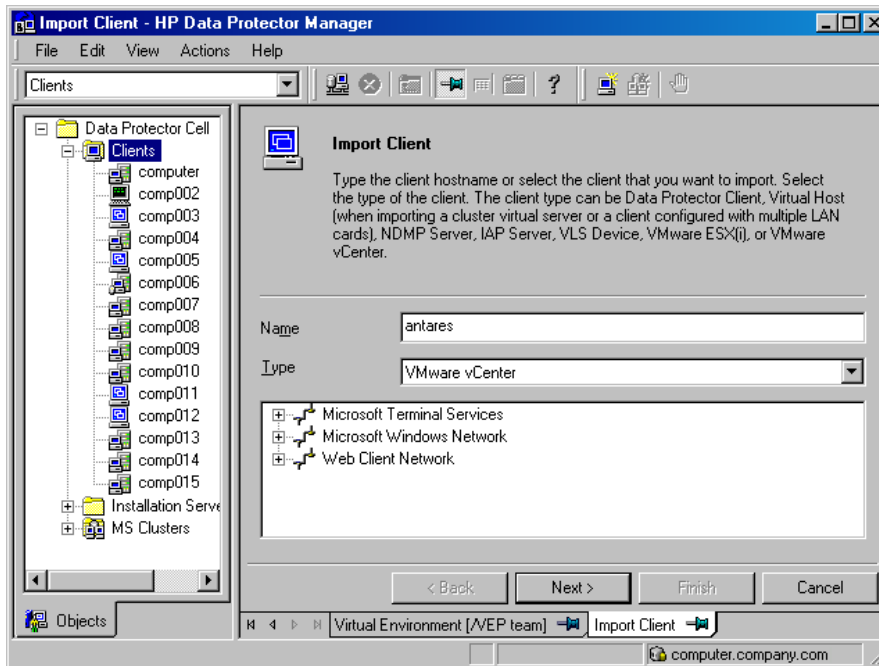
NOTE: To install a Data Protector component on a VMware client (for example, you want to be able to perform a filesystem backup on the client), install the Data Protector component (such as Disk Agent) first and then import the client as a VMware vCenter Server. Otherwise, you will have to re-import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

Procedure

To import a client into a Data Protector cell:

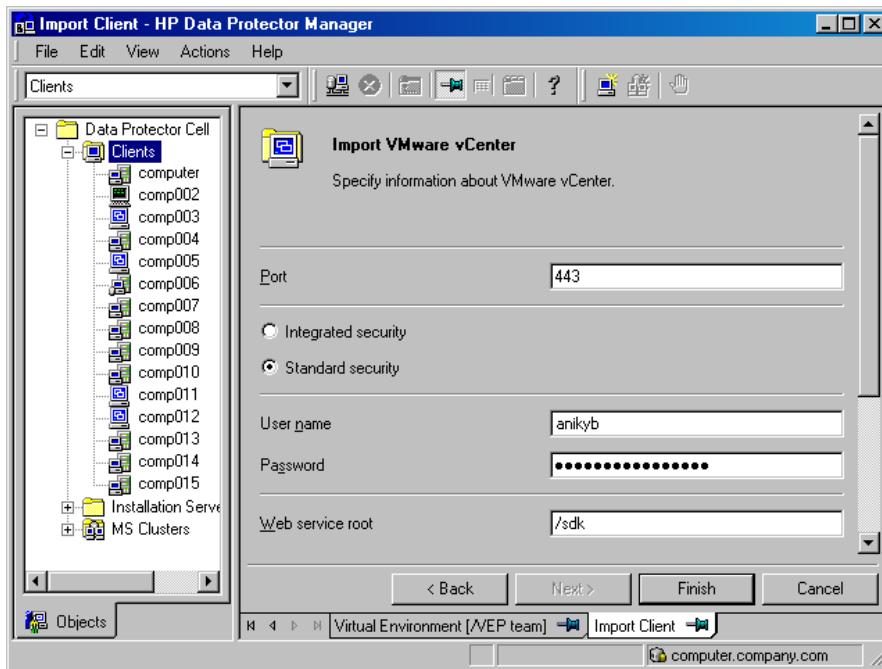
1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell**, right-click **Clients**, and click **Import Client**.
3. In the Import client page, enter the client name in the **Name** option, select the appropriate client type **VMware vCenter** from the **Type**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

Figure 4 Importing a VMware client (name and type)



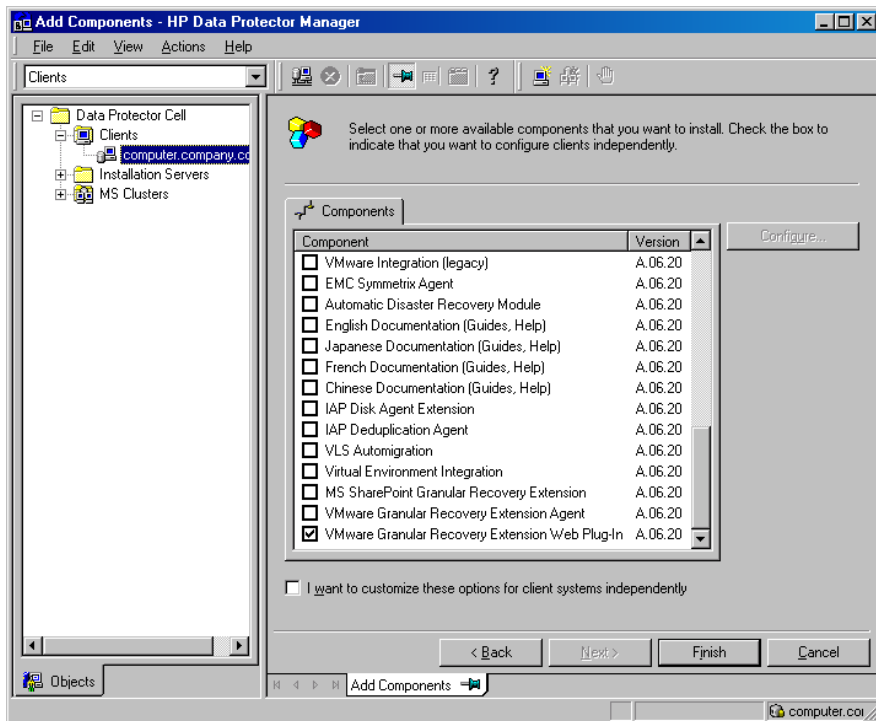
4. Specify login credentials:
Port: Specify the port that VMware vSphere is using. By default, VMware uses the port 443.
User name and **Password:** Specify an operating system user account that has the following VMware vSphere privileges:
 - **Web service root:** Optionally, change the web service entry point URI. Default: /sdk.Click **Finish**.

Figure 5 Importing a VMware client (login credentials)



5. Install the VMware Granular Extension Web Plug-In component remotely.

Figure 6 Adding Data Protector component on VMware vCenter Server system

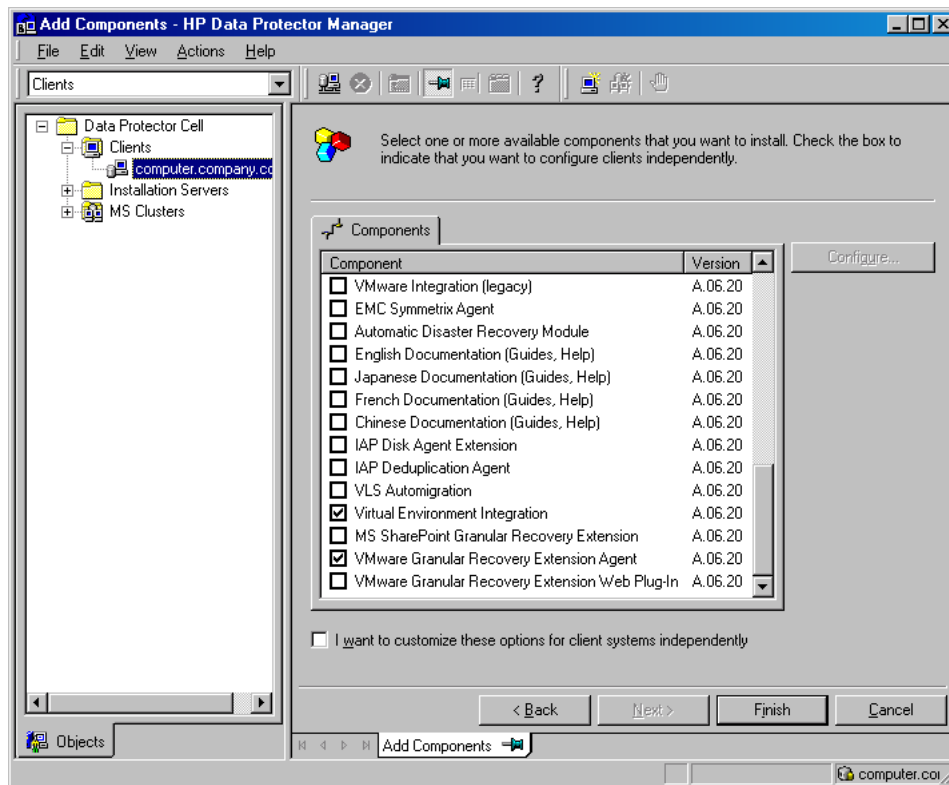


If you cannot, for any reason, install this extension remotely, see the [“Local installation workaround”](#) (page 55).

Mount proxy system

Add the Virtual Environment Integration and the VMware Granular Recovery Extension Agent components, and click **Finish**.

Figure 7 Adding Data Protector components on mount proxy system



If you cannot, for any reason, install this extension remotely, see the “[Local installation workaround](#)” (page 55).

Configuring a Granular Recovery Extension VMware user group

Configuring Granular Recovery Extension for a VMware vSphere user group

Prerequisite

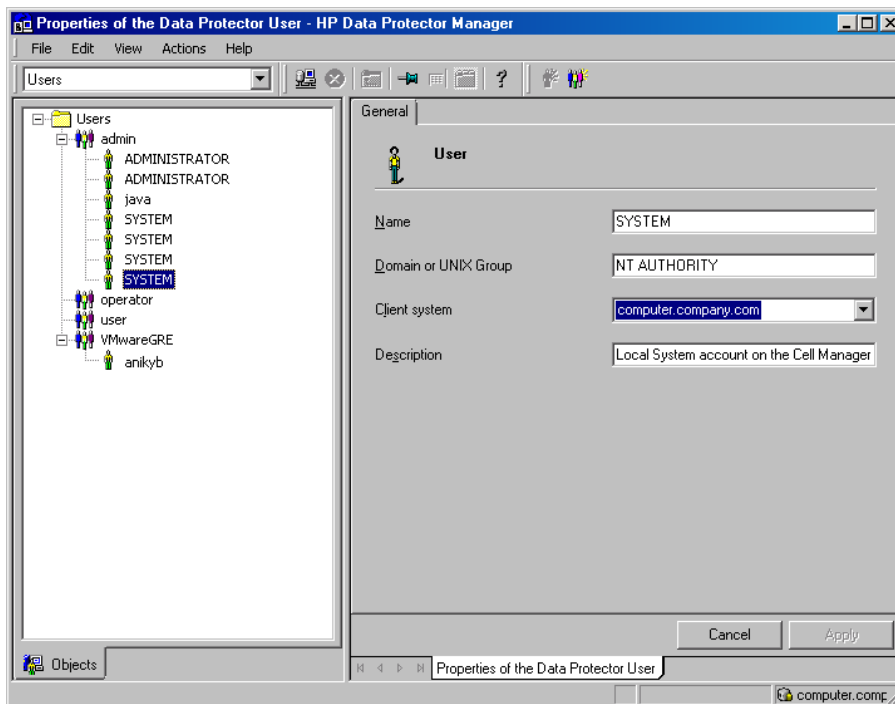
You must have Data Protector User configuration rights.

Procedure

To create a Data Protector group called GRE VMware to meet the requirements of this extension:

1. Open the HP Data Protector GUI (Data Protector Manager).
2. In the Context List, click **Users**.
3. In the Scoping Pane, right-click **Users**

Figure 8 GRE VMware vSphere user group



4. Click **Add User Group** to open the wizard.
5. Under General, type the **Name**: GRE VMware and **Description** of the new group: Granular Recovery Extension for VMware administrators.
6. Click **Next**
7. Set the specific start restore user right for the new group.
8. Click **Finish** to exit the wizard.

The new empty GRE VMware user group is added to Data Protector. To add new users, see [“Adding users to the GRE VMware group” \(page 25\)](#).

If you do not want a GRE VMware user group, you can have a Granular Recovery Extension administrator. For detailed procedure see [“Configuring Granular Recovery Extension Administrators using HP Data Protector” \(page 28\)](#).

Adding users to the GRE VMware group

To configure users for the Granular Recovery Extension for VMware administrators group, add users to the group.

Prerequisite

You need to have the User configuration right to be able to add users.

Procedure

When adding users to GRE VMware group:

1. In the Context List, click **Users**.
2. In the Scoping Pane, expand **Users**.
3. Right-click the user group to which you want to add a user.
4. Click **Add/Delete Users** to open the wizard.
5. In the Add/Delete Users dialog, enter the specific user properties.

When entering **Name** and **Group/Domain**, make sure you enter information about an existing user on your network.

To ensure GRE VMware Administrators have access to the administrative entry point of the extension, specify the following information:

Type: Windows

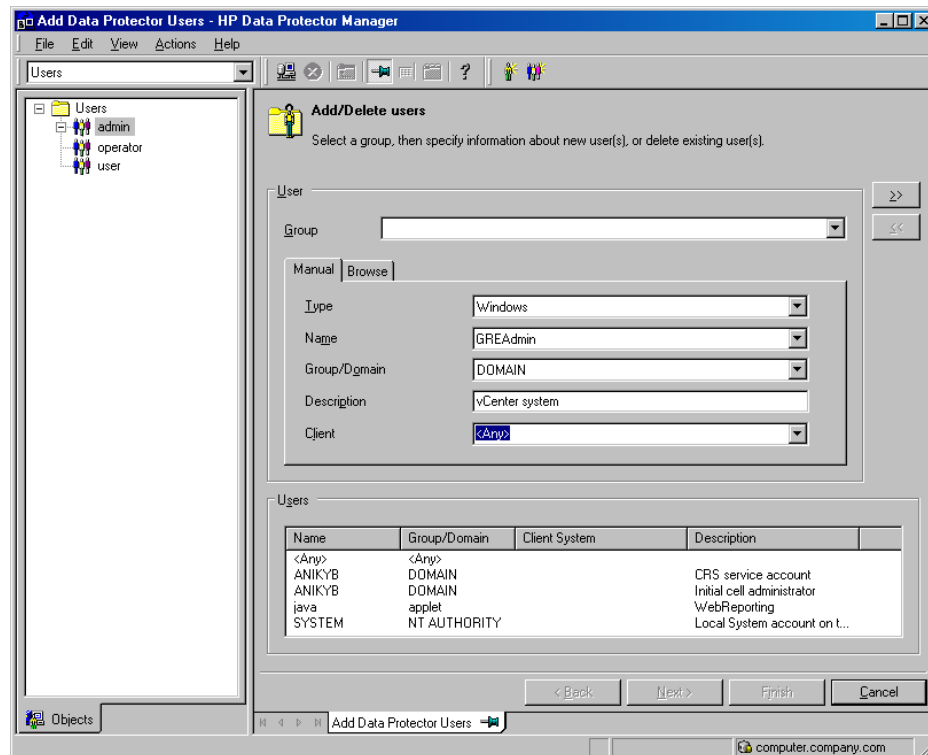
Name: *username*

Group/Domain: GRE VMware user group, VCENTER

Client: *VCenterSystemName*

These GRE VMware Administrators must be added to the vSphere permission tab. Set the user's role to Administrator. For detailed procedure, see [“Configuring vCenter Server options”](#) (page 29).

Figure 9 Adding user domain information



6. Click the arrow button >> to add the user to the user list.
7. Click **Finish** to exit the wizard.

The user is added to the Granular Recovery Extension for VMware administrators group and has the `start restore` user right assigned to the group.



TIP: You can also delete a user by selecting the user in the user list and clicking <<.

For a detailed description of user roles and tasks, see [“HP Data Protector Granular Recovery Extension users”](#) (page 20).

For details of the procedure for configuring Data Protector to meet this extension's user rights requirements, see [“Configuring Granular Recovery Extension Administrators using HP Data Protector”](#) (page 28).

NOTE: Granular Recovery Extension administrators should be added to this group, but you do not need to add Granular Recovery Extension operators. The Granular Recovery Extension operators can perform their own tasks without ever opening Data Protector.

Adding an Inet user account to the Data Protector Administrator group

To ensure the proper functioning of the VMware Granular Recovery Extension Agent, and the VMware Granular Extension Web Plug-In extension components, you must add an Inet user account to the Data Protector Administrator's group on the following systems:

1. In the Context List, click **Users**.
2. In the Scoping Pane, expand **Users**.
3. Right-click the **admin** group to which you want to add a user.
4. Click **Add/Delete Users** to open the wizard.
5. In the Add/Delete Users dialog, enter specific user properties:

When entering **Name** and **Group/Domain**, make sure you enter information about an existing user on your network:

Mount proxy system:

Type: Windows

Name: SYSTEM

Group/Domain: NT AUTHORITY

Client: *MountProxySystemName*

vCenter Server system:

Type: Windows

Name: the account under which the VMware vCenter Server runs (by default is SYSTEM).

Group/Domain: the account group/domain under which the VMware vCenter Server runs (by default, NT AUTHORITY).

Client: *VCenterSystemName*.

6. Click the arrow button **>>** to add the user to the user list.
7. Click **Finish** to exit the wizard.

The user accounts are added to the Data Protector administrator's group and have all the user rights assigned to them.

NOTE: The Data Protector Inet uses SYSTEM, NT AUTHORITY as in Data Protector Local system account (as in Windows operating system).

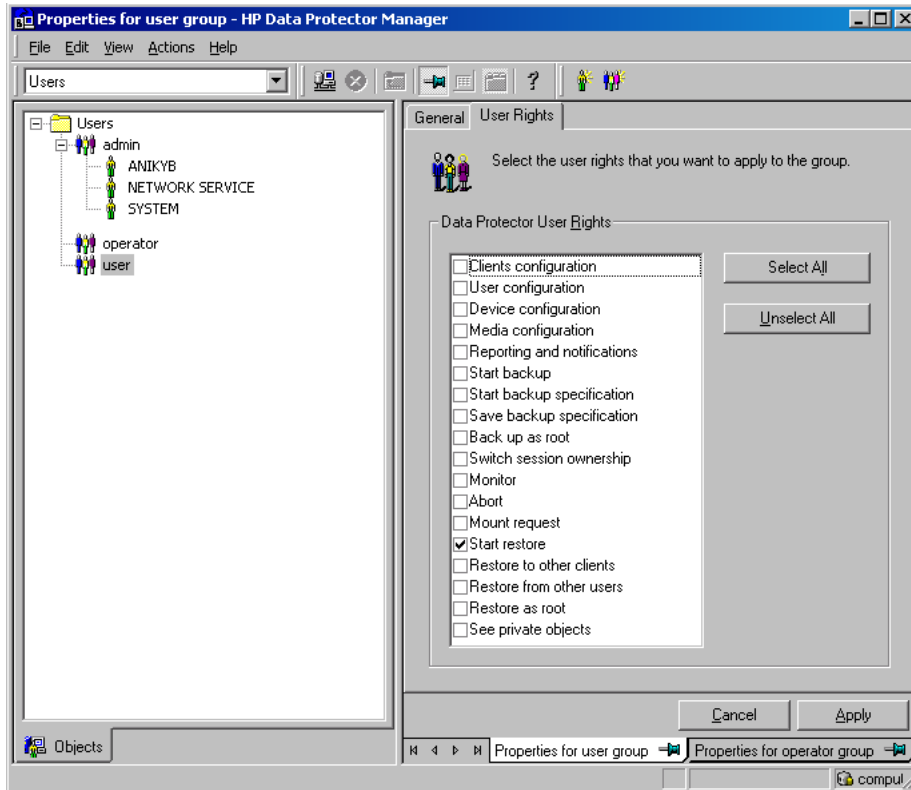
Configuring Granular Recovery Extension Administrators using HP Data Protector

Procedure

To configure the free access of Granular Recovery Extension administrators to the administrative tasks of this extension:

1. Open the HP Data Protector GUI (**Data Protector Manager**).
2. In the Context list, select **Users**.
3. Right-click the GRE VMware user group to be modified.
4. Click **Properties** and then click the **User Rights** tab.
5. Ensure the Data Protector user account is assigned Data Protector Start restore user right.

Figure 10 Data Protector user rights



NOTE: In the Granular Recovery Extension, administrators are Data Protector users assigned the Data Protector Start restore user rights. For a detailed description of user roles and tasks, see [“HP Data Protector Granular Recovery Extension users”](#) (page 20).

The specified user rights are assigned to the user group and to all users belonging to the group. It is recommended to create a user group called VMware GRE vSphere for this extension's administrators.

For details of the procedure, see [“Configuring a Granular Recovery Extension VMware user group”](#) (page 24).

Configuring your systems for VMware vSphere purposes

Configuring Windows Firewall exceptions

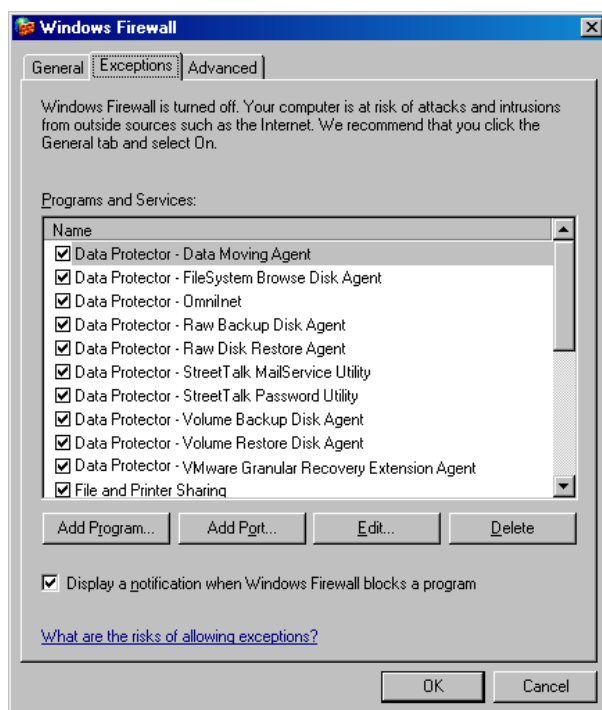
Setting VMware Granular Recovery Extension Agent in Firewall exceptions

Procedure

To ensure communication between the Mount proxy system component and the extension on the vCenter Server:

1. Check your Windows Firewall Exceptions list.
2. Make sure the VMware Granular Recovery Extension Agent (vmwaregre-agent.exe) is in the list of Windows Firewall exceptions both on the mount proxy and on the vCenter Server systems:

Figure 11 Windows Firewall Exceptions tab



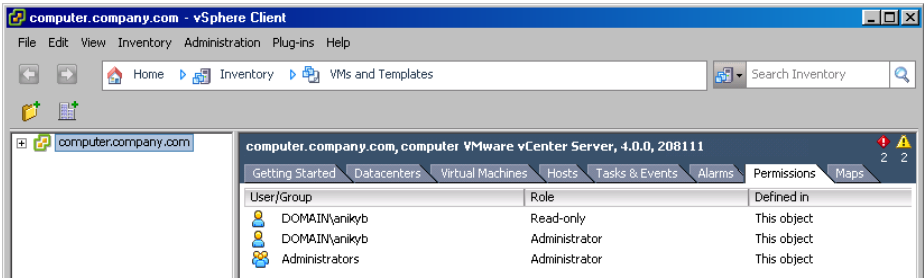
NOTE: Only inbound firewall rules are automatically created during the installation of the Granular Recovery Extension for VMware. You must manually create any outbound firewall rules. This is required for communication between this extension's components and the vCenter Server. For the required port ranges, see the HP Data Protector online Help index: "firewall support".

Configuring vCenter Server options

Procedure

Set the Administrator permission role to the user account on vCenter Server. Read only does not allow all the necessary permissions. For reference, see "[vCenter Server role options](#)" (page 30).

Figure 12 vCenter Server role options



4 Backup

Back up VMware vSphere data files and folders of virtual machine disks as described in the Data Protector backup solution documentation.

This extension supports full, incremental and differential backups. For details on Data Protector backup types, see HP Data Protector Help, index: “backup types”.

For more information on the HP Data Protector backup solutions, see the *HP Data Protector Integration Guide for Virtualization Environments*.

NOTE: The Granular Recovery Extension for VMware vSphere uses the same procedure for recovery of all VMware vSphere data. The procedure does not depend on the backup type.

Considerations

- The Granular Recovery Extension for VMware vSphere does not support cross-platform backup, restore, and recovery. This means that you must perform the backup, restore, and recovery using the same Windows virtual machine's platform.
- Use the Data Protector Virtual Environment integration to back up virtual machine disks' VMware vSphere. The Data Protector VMware (Legacy) integration is not supported.

NOTE: It is recommended to setup a dedicated mount proxy system, this extension will need to allocate extendable disk space (disk space that can be made larger) for the temporary restore location of virtual machine disks.

5 Restore and recovery

Considerations

The following considerations apply if the virtual machine is **not** on the network:

- When restoring an NTFS file system, the following are not restored or recovered:
 - Access control list ACL
 - Alternative streams (for example, thumbnails for graphical files, parsing information for program sources, spell check and formatting data for documents).
 - File attributes
 - Symbolic links

In this case, if symbolic links point to directories, the recovered symbolic links are turned into directories.
- Other considerations:
 - If two operators connect remotely using the vSphere interface with the same user account and browse one partition of a virtual machine disk simultaneously, once the first partition is mounted the second partition is dismounted. Parallel mounting of partitions is not supported.
 - This extension displays the whole restore chain. When browsing through files, the objects displayed contain the whole restore chain: a full backup of the object and any number of related incremental backups.
 - You may notice slow performance when browsing too many files and sub files in a folder.
 - If you backed up virtual machine disks using the Data Protector options described below, the Granular Recovery Extension for VMware vSphere cannot restore or recover the data:
 - If the virtual machine disks used the Microsoft Volume Shadow Copy (VSS) quiescence functionality, the **Use quiescence snapshots** option selected for an application-consistent backup (vStorage Image backup method).
For detailed information about quiescence functionality, see the HP Data Protector Integration Guide for Virtualization Environments guide.
 - If the virtual machines had user-created snapshots configured with snapshot handling, and the following options enabled:
 - **Mixed Incremental and differential possible in same backup chain**
 - **Enable changed block tracking**
 - **Use change block tracking**

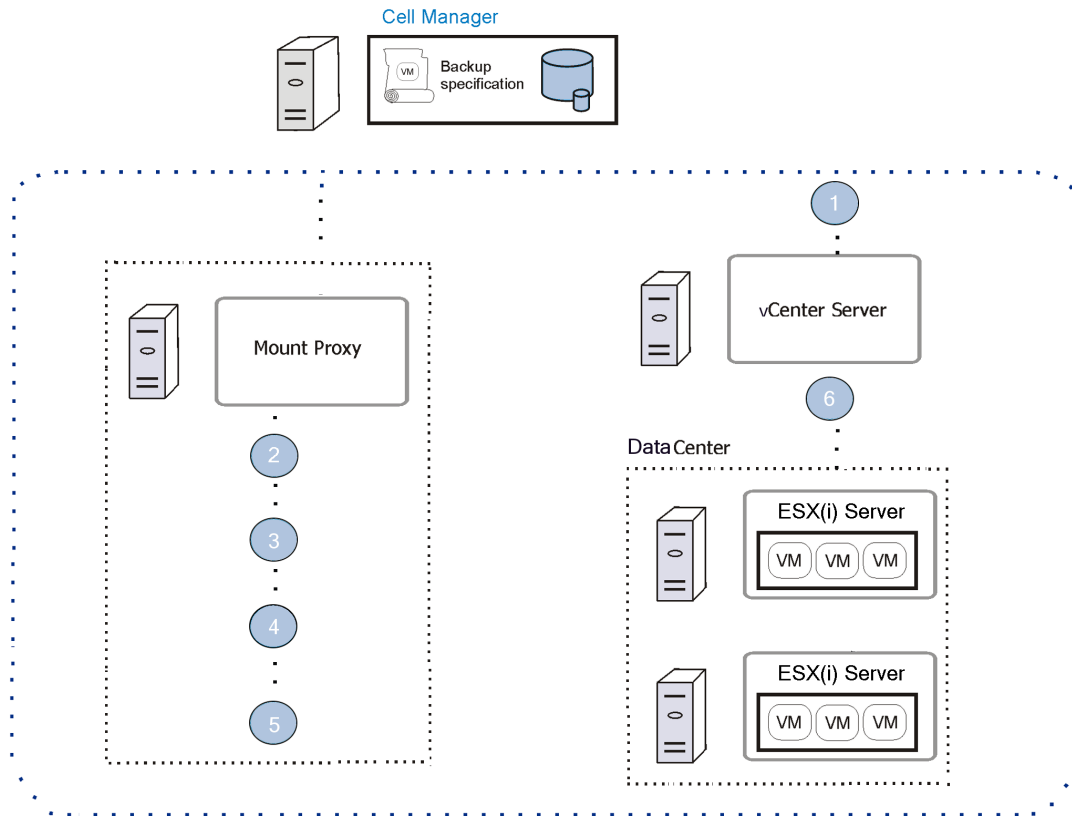
NOTE: To enhance the display of the Granular Recovery Extension interface, use Windows Internet Explorer 8 or newer version.

For a list of supported environments, see the latest support matrices at <http://www.hp.com/support/manuals>.

Restore and recovery flow

The basic procedure for recovering virtual machine filesystems is as follows:

Figure 13 Restore and recovery flow



1. Request restore

Operators (VMware vCenter users) request Granular Recovery Extension Administrators to perform restores.

For detailed configuration steps on assigning the Data Protector `Start restore` user rights, see [“Configuring Granular Recovery Extension Administrators using HP Data Protector”](#) (page 28).

At this stage, VMware Granular Recovery Extension Agent starts the restore session in conjunction with the Virtual Environment Protection Agent to restore the files on the mount proxy system.

2. Handle requests

Allowing restore requests enables Granular Recovery Extension Operators to browse through and select virtual machine files and folders for recovery.

Once a Granular Recovery Extension Administrator clicks **Start restore** the restore session is initiated, and the virtual disk files are copied to a temporary location on the mount proxy system.

3. Mount virtual disks

Granular Recovery Extension Operators browse through files with this extension and the virtual machine's disks are mounted to a drive letter on the mount proxy system. Once the virtual machine's disks are mounted, the recovery session is possible.

The mounted virtual machine disks are dismounted from the drive letter when the page is closed. This makes the drive letters available for new mounting.

4. **Recover**

Granular Recovery Extension Operators select files and start the recovery session. At this stage, the recover granularity is visible. The individual files are recovered to any online virtual machine on the vCenter Server system.

5. **Remove restored data, files of virtual machine disks**

The restored disks are deleted automatically after the retention period expires.

If no more disk space is available on the mount proxy, the administrator must remove virtual machine disks manually.

6. **Removal of virtual machine disks**

The retention period for the restored files of virtual machines' disks on the mount proxy is 30 days, after that the files are no longer available for browsing recovery. After the retention period, the restored files are deleted automatically from the temporary restore location on the mount proxy. It is an administrative task to change the default value according to the requirements of Granular Recovery Extension Operators, taking their needs into account.

NOTE: Change retention period options on the **Settings** page.

Changes only affect new requests.

Operator's tasks

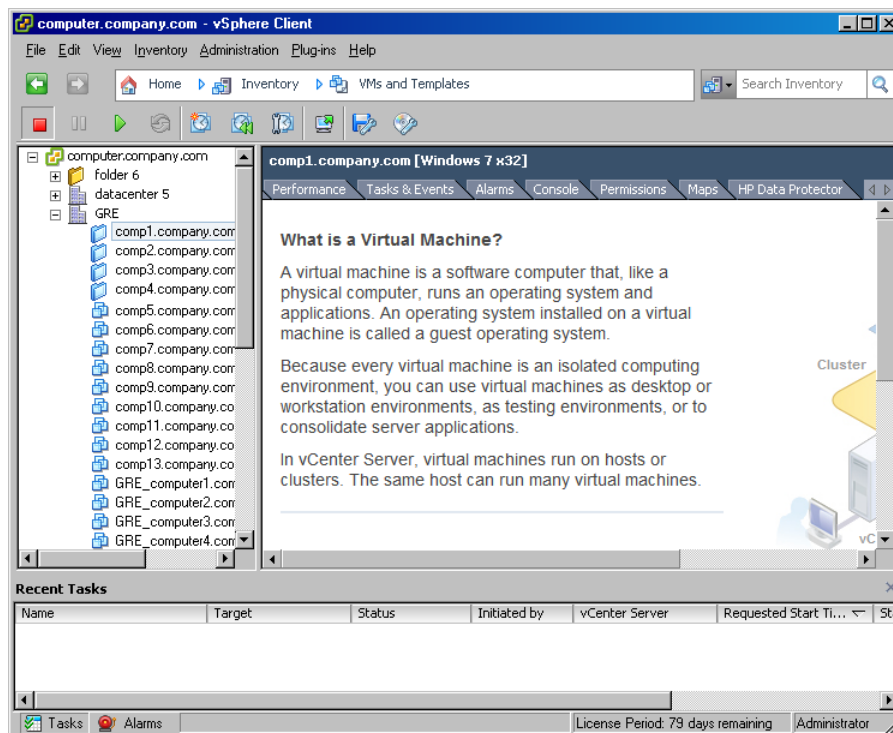
Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI

Procedure

To open the HP Data Protector Granular Recovery Extension for VMware vSphere GUI:

1. Connect remotely to the vCenter Server system using VMware vSphere Client interface, and specify the user account and password.
2. The VMware vSphere home page is displayed. Under the Inventory tab select **VMs and templates**, and expand the virtual machines browse tree. Locate the virtual machine and select it.
3. Locate the HP Data Protector tab and click it.

Figure 14 Locating the extension



NOTE: After clicking the HP Data Protector tab, you receive Security Alert notifications, if you do not have the vSphere certificate installed on the system. Click **Yes** to open this extension. For detailed information on how to obtain a valid vSphere certificate, see the VMware end-user documentation, or Help index: "About vCenter Server Certificates" or "Replacing vCenter Server Certificates".

The **HP Data Protector Granular Recovery Extension** is displayed in operator's mode.

Figure 15 HP Data Protector Granular Recovery Extension open

Getting StartedSummaryResource AllocationPerformanceTasks & EventsAlarmsConsolePermissionsMapsHP Data Protector

HP Data Protector Granular Recovery Extension

View:Request RestoreMonitor RequestsRecover FilesSettings

User Manual

Select the backup of Virtual Machine disks from which you want to recover files. Specify the retention time you need to recover files (after this period, disks will be deleted). Click **Request Restore** to submit restore request.

Restore Request Options

Request Restore Refresh

Source Data Center:Data Center

Source VM:computer:company.com

Include Backups In The Last:

Please select...

Backup Start Time:

Please select...

Disk Name

Disk Size (GB)

Virtual Disks:No Data Available

Disk Retention Time:

30

Days

36 Restore and recovery

Requesting restores

Procedure

To request restores:

1. In the **HP Data Protector Granular Recovery Extension for VMware vSphere** page, click **Request Restore**.

Figure 16 Loading request restore

Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions Maps HP Data Protector

hp HP Data Protector Granular Recovery Extension

View: Request Restore Monitor Requests Recover Files Settings User Manual

Select the backup of Virtual Machine disks from which you want to recover files. Specify the retention time you need to recover files (after this period, disks will be deleted). Click **Request Restore** to submit restore request.

Restore Request Options Request Restore Refresh

Source Data Center: Data Center

Source VM: computer.company.com Please wait...

Include Backups In The Last: 30 days

Backup Start Time: Please select...

Virtual Disks: No Data Available

Disk Retention Time: 30 Days

Command Status

2. The Restore Request Options page is displayed.

Figure 17 Restore page loaded

Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions Maps HP Data Protector

hp HP Data Protector Granular Recovery Extension

View: Request Restore Monitor Requests Recover Files Settings User Manual

Select the backup of Virtual Machine disks from which you want to recover files. Specify the retention time you need to recover files (after this period, disks will be deleted). Click **Request Restore** to submit restore request.

Restore Request Options Request Restore Refresh

Source Data Center: Data Center 1

Source VM: computer.company.com

Include Backups In The Last: Please select...

Backup Start Time: Please select...

Virtual Disks: No Data Available

Disk Retention Time: 30 Days

Command Status

In the **Include backups in the last:** drop-down list, select the period of the backup. The **Backup start time** drop-down list is displayed; select one.


3. Select the virtual disk to be restored.

Figure 18 Selecting virtual disks

The screenshot shows the HP Data Protector Granular Recovery Extension web interface. The top navigation bar includes tabs: Getting Started, Summary, Resource Allocation, Performance, Tasks & Events, Alarms, Console, Permissions, Maps, and HP Data Protector. The main header displays the HP logo and the title 'HP Data Protector Granular Recovery Extension'. Below the header, there are four tabs: 'Request Restore' (selected), 'Monitor Requests', 'Recover Files', and 'Settings'. A 'User Manual' link is visible on the right. The main content area contains instructions: 'Select the backup of Virtual Machine disks from which you want to recover files. Specify the retention time you need to recover files (after this period, disks will be deleted). Click **Request Restore** to submit restore request.' Below this, there are 'Request Restore' and 'Refresh' buttons. The 'Restore Request Options' section includes: 'Source Data Center: Data Center', 'Source VM: computer.company.com', 'Include Backups In The Last: 30 days', and 'Backup Start Time: 2/17/2011 1:46:12 PM Full'. A table for 'Virtual Disks' is shown with columns 'Disk Name' and 'Disk Size (GB)'. The table lists three disks: 'scsi0:0' (6.00 GB) with a checked checkbox, 'scsi0:1' (0.20 GB) with an unchecked checkbox, and 'scsi0:2' (0.29 GB) with an unchecked checkbox. Below the table, 'Disk Retention Time' is set to '30 Days'. At the bottom, there is a 'Command Status' section.

4. Specify the Retention period. This period starts at the moment of the restore. After that the virtual disks are no longer available. Click **Request Restore**.
The default retention period of the restored data is 30 days.
5. Click **Monitor Requests** to display this request status. For future reference, note your request ID to easily identify your request.

Figure 19 Request ID



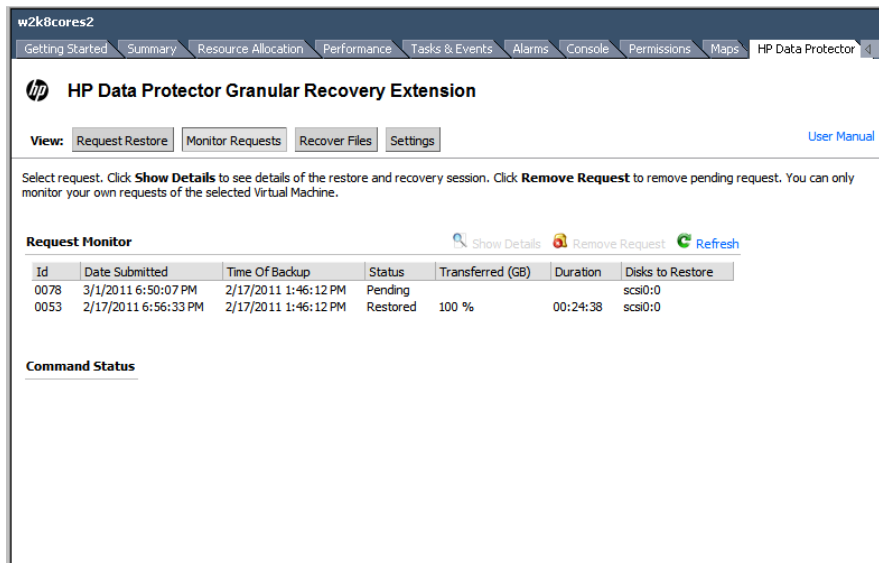
Monitoring restore requests

Procedure

To monitor restore requests:

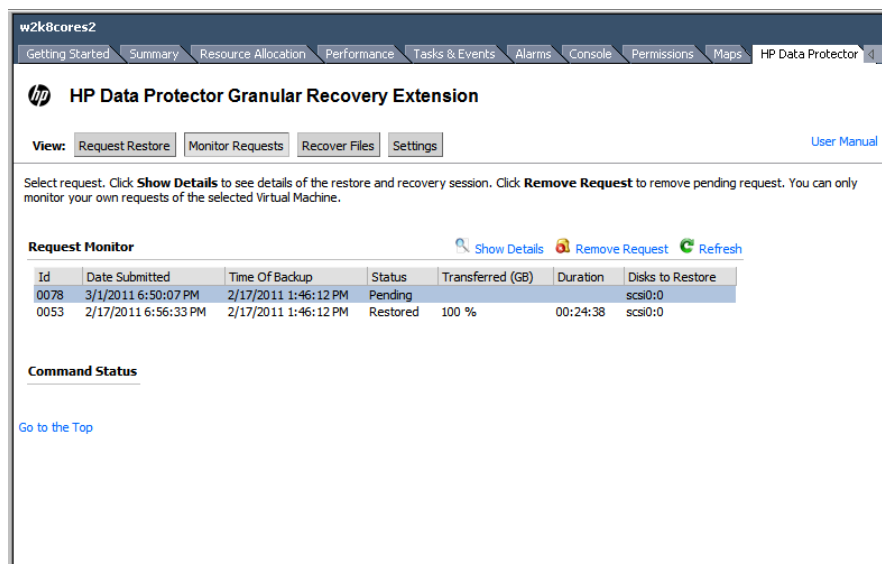
1. In the **HP Data Protector Granular Recovery Extension for VMware vSphere** page, click **Monitor Requests**.

Figure 20 Viewing restore requests



2. Identify your request, and select it.
3. You can:
 - Click **Refresh** to update the status request information.
 - Click **Show Details** to display information about this request status.
 - Click **Remove request** to delete the pending requests status, that are no longer needed.

Figure 21 Viewing restore requests details



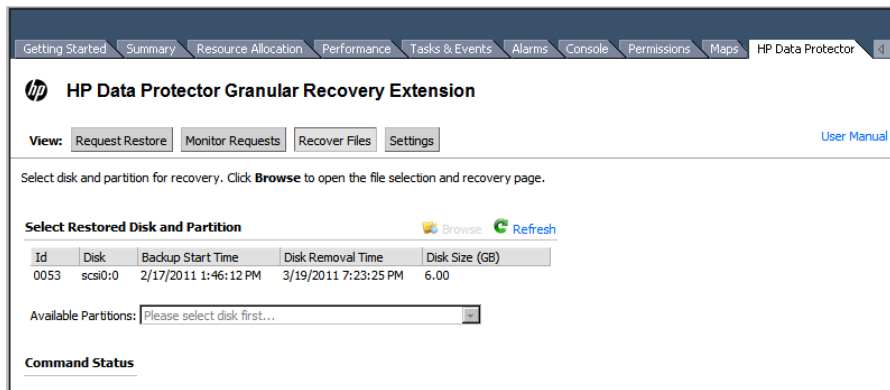
Recovering virtual machines

Procedure

To recover files of virtual machine disks:

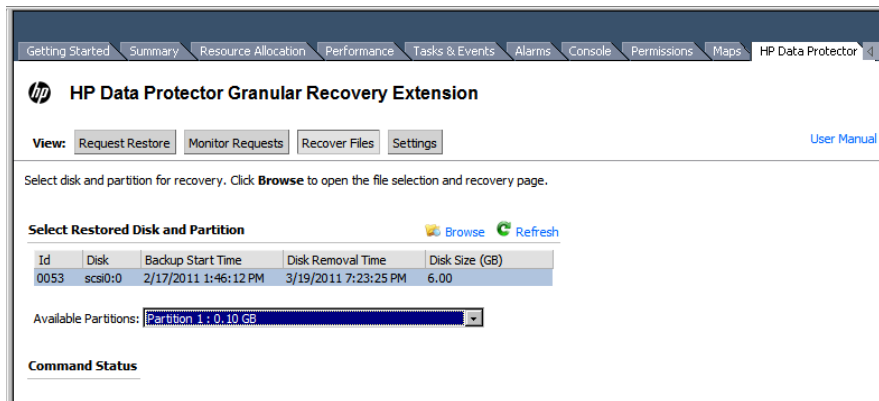
1. In the **HP Data Protector Granular Recovery Extension for VMware vSphere** page, click **Recover files**.

Figure 22 Recover files loaded



2. Select the virtual machine disk that contains the files for recovery. The available partitions are now displayed in the drop-down list. Select a partition, and click **Browse**. The Select Files to Recover page is displayed. Under the Virtual Machine OS Credentials, select the virtual machine, and specify the user name and password.

Figure 23 Selecting virtual disks



Under the Recovery Options, type the target recovery location. If the file already exists on the target system, select the recovery option to be used:

- overwrite: deletes the original files and folders and saves the latest files and folders.
- rename: keeps the original files and folders, and saves the recovered files and folders with the following naming convention: *pathname.timestamp*.
- skip: keeps the original files and folders intact.

Figure 24 Setting recovery options part

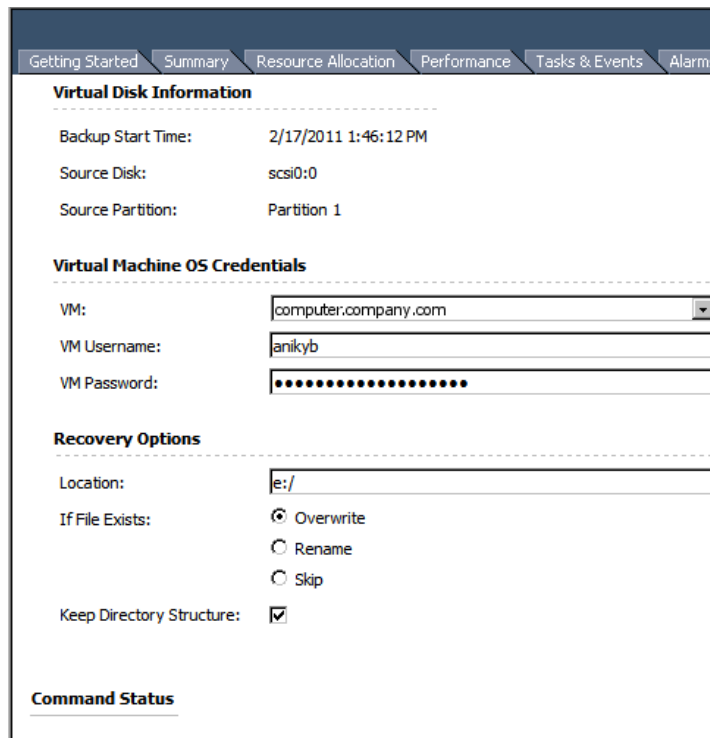
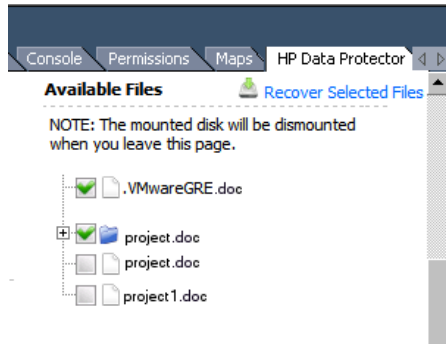


Figure 25 Granular recovery, selecting files for recovery



To maintain the original directory file structure from the source virtual machine disk on the target system, select **Keep Directory Structure**.

Under **Available Files**, select files and folders to be recovered and click **Recover Selected Files**.

NOTE: The files and folders displayed on this page depend on the backup you selected in the Request Restore procedure. For instance, if you selected an incremental backup, the whole restore chain is displayed. See the Data Protector online Help index: backup types, or restore chain, for details about Data Protector restore and backup options.

Administrator's tasks


Opening the HP Data Protector Granular Recovery Extension for VMware vSphere GUI

Procedure

To open the HP Data Protector Granular Recovery Extension for VMware vSphere GUI:

1. Connect remotely to the vCenter Server system using the VMware vSphere Client interface, and enter the user account and password. Specify the user credentials of the Data Protector user to which the `start restore` right is assigned.
2. Connect to the vSphere Client Home page.
3. In the Administration tab, locate and click **HP Data Protector (Admin)**.

Figure 26 Locate HP Data Protector (Admin)



NOTE: After clicking the HP Data Protector tab, if you do not have vSphere certificate installed on the system, you receive Security Alert notifications. This happens after clicking the. Clicking **Yes** opens this extension. For detailed information on how to obtain a valid vSphere certificate, see the VMware end-user documentation, or Help index: "About vCenter Server Certificates" or "Replacing vCenter Server Certificates".

The HP Data Protector Granular Recovery Extension for VMware vSphere GUI is displayed on the Handle Requests page. For details of the procedure, see "[Handling restore requests](#)" (page 44).

NOTE: You can open this extension using an alternative path. This administrative entry can be accessed by selecting **View, Administration, HP Data Protector (Admin)**.

Configuring mount proxy systems

Procedure

To specify the mount proxy systems used as a target location for restoring virtual machine disks:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Settings**.

Figure 27 Administrative settings

HP Data Protector Granular Recovery Extension

View: [Handle Requests](#) [Monitor Requests](#) [Cleanup](#) [Settings](#) [User Manual](#)

NOTE: Ensure to click **Save** to save the options before leaving this page.

Select host that will be your temporary restore location and restore path. You can specify multiple restore paths.

Mount Proxy Information [Save](#) [Refresh](#)

Host:

Restore Paths: [Add](#)

Enter Path in Windows Format

[×](#)

[×](#)

Retention Time Options

Restored Disks Will Be Removed After: Days

Failed Request Will Be Removed After: Days

Rejected Request Will Be Removed After: Days

Debugging Options

Enable Debugging: ☒

2. There is a list of available mount proxy systems (Data Protector clients with a Virtual Environment Integration and the VMware Granular Recovery Extension Agent components installed) in the Host drop-down list. Specify the restore path for each mount proxy system:

NOTE: If you do not specify a restore path for a mount proxy system, this system will not be available for restore.

- a. In the Restore Paths text box, type the location on the mount proxy system, where you want to restore the virtual machine disks.
- b. Click **Add** to add the specified path to the list of restore paths. Notice that you can add more restore paths to the list.
- c. Click **Save**.

You can delete the specified restore path by clicking the red cross to the right of the corresponding text box.

Handling restore requests

Procedure

When you receive restore requests from operators, you can perform two tasks for detailed procedures, see:

- “Triggering restore sessions” (page 45).
- “Rejecting restore requests” (page 46).

Triggering restore sessions

Procedure

To trigger the restore sessions:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Handle Requests**.

Figure 28 Managing restore requests

The screenshot shows the HP Data Protector Granular Recovery Extension GUI. The breadcrumb navigation at the top indicates the path: Home > Administration > HP Data Protector (Admin). Below the navigation bar, the title 'HP Data Protector Granular Recovery Extension' is displayed. Under the 'View:' section, the 'Handle Requests' tab is selected, with other tabs being 'Monitor Requests', 'Cleanup', and 'Settings'. A message states: 'Select request. Click **Start Restore** to start restore session of Virtual Machine disks. Click **Reject** to reject request.' Below this, there are three buttons: 'Start Restore' (green), 'Reject' (red), and 'Refresh' (blue). The main section is titled 'Browse Requests' and contains a table with the following data:




Id	Status	Date Submitted	Submitter
0061	Pending	1/31/2011 7:02:04 PM	ADMINISTRATOR
0065	Pending	2/2/2011 9:43:47 AM	ADMINISTRATOR
0067	Pending	2/2/2011 11:14:02 AM	ADMINISTRATOR
0069	Pending	2/4/2011 2:07:06 PM	ADMINISTRATOR
0073	Pending	2/4/2011 4:37:16 PM	ADMINISTRATOR
0074	Pending	2/4/2011 6:41:59 PM	ADMINISTRATOR
0075	Pending	2/4/2011 6:42:05 PM	ADMINISTRATOR
0076	Pending	2/18/2011 11:18:51 AM	ADMINISTRATOR
0077	Pending	2/18/2011 2:46:17 PM	ADMINISTRATOR
0078	Pending	2/22/2011 3:12:38 PM	ADMINISTRATOR
0079	Pending	2/22/2011 3:16:19 PM	ADMINISTRATOR
0072	Pending	3/1/2011 11:30:46 AM	ADMINISTRATOR
0074	Pending	3/1/2011 11:31:41 AM	ADMINISTRATOR
0075	Pending	3/1/2011 11:32:46 AM	ADMINISTRATOR
0076	Pending	3/1/2011 11:49:46 AM	ADMINISTRATOR
0077	Pending	3/1/2011 12:22:29 PM	ADMINISTRATOR
0078	Pending	3/1/2011 6:50:07 PM	ANIKYB
0079	Pending	3/1/2011 7:26:24 PM	ADMINISTRATOR
0072	Rejected	3/1/2011 11:01:36 AM	ADMINISTRATOR

2. In the Browse Requests list, click the pending restore request. The restore request details are displayed on the right.
3. From the Restore location drop-down list, select one of the specified restore paths. The available disk space on the target restore location and the cumulative size of the disks selected for restore are displayed.

NOTE: To change the mount proxy system, or specify an additional restore path, see “Configuring mount proxy systems” (page 44).

4. After you specified the restore options for the selected restore request, click **Start Restore**.

Figure 29 Starting operators restore requests

 [Start Restore](#)  [Reject](#)  [Refresh](#)

Request Details

Request ID: 0078

Source Data Center: Data Center 1

Source ESX: computer.company.com

Source VM: computer.company.com

Time Of Backup: 2/17/2011 1:46:12 PM

	Disk Name	Disk Size (GB)
Virtual Disks:	<input checked="" type="checkbox"/> scsi0:0	6.00

Disks Retention Time: 30 Days

Estimated Size Needed on Disk: 6.00 GB

Mount Proxy: computer.company.com

Restore Location:

Available Disk Space on Mount Proxy: 0 GB

Rejecting restore requests

Procedure

To reject operators restore requests:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Handle Requests**.
2. In the Browse Requests list, click the restore request with the Pending status you want to reject.
3. Click **Reject**.

Monitoring restore requests

Procedure

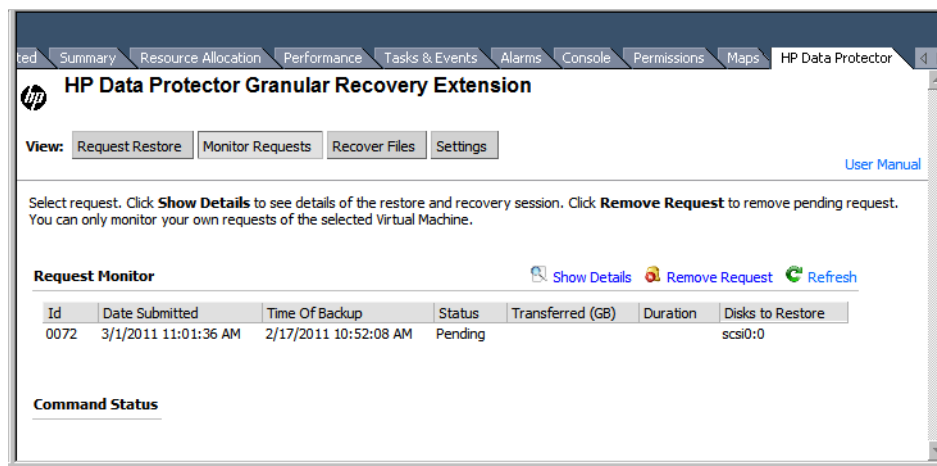
To monitor requests:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Monitor Requests**.

The Request Monitor list contains all the pending and triggered restore requests. Restore requests can have the following statuses: Restoring (In Progress), Restored, Recovering, or Failed.

2. To see a detailed session report for a specific session, select this session from the list, and then click **Show Details**.
3. To end the running session (the session with the In Progress status), select the session you want to end, and then click **Stop Restore**.

Figure 30 Viewing operator restore requests



Stopping restore sessions

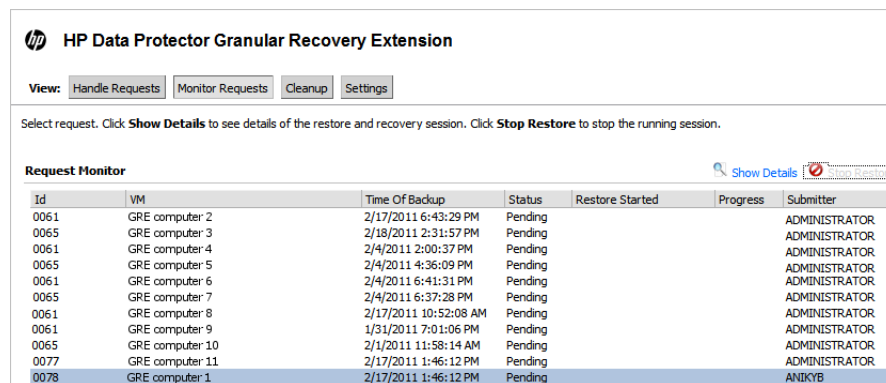
Procedure

To stop operator's restore sessions:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Handle Requests**.
2. In the Browse Requests list, click the restore request with the Restoring (In Progress) status you want to stop, and click **Stop Restore**.

The restore session is stopped.

Figure 31 Stopping pending sessions



Removing virtual machine disks

Procedure

When the specified retention period is expired, the virtual machine disks are deleted automatically. In certain cases, you can perform this removal manually. Under normal circumstances this is not recommended.

To delete virtual machine disks manually:

1. In the HP Data Protector Granular Recovery Extension for VMware vSphere GUI, click **Cleanup**.

Figure 32 Removing disks

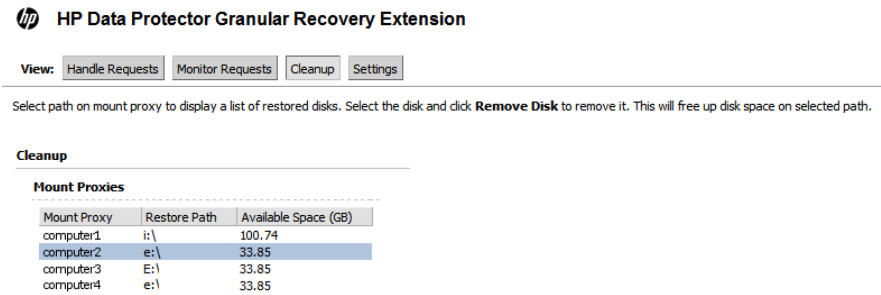
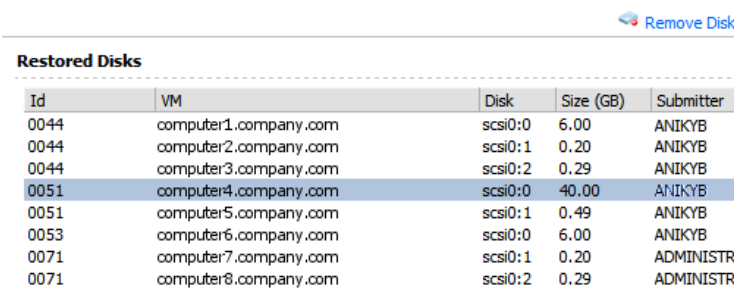


Figure 33 Clean up



The Cleanup list contains the restored virtual disks.

2. To remove the restored virtual disk from the mount proxy system, select it from the Cleanup list and then click **Remove Disk**.

6 Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Granular Recovery Extension for VMware vSphere.

- For Virtualization Environments troubleshooting information, see the troubleshooting section of the VMware part about the Data Protector Virtual Environment integration, in the HP Data Protector Integration Guide for Virtualization Environments.
- For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

Before you begin

- To enable debugging, see “Enabling debugging option” (page 50).
- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*
- For an up-to-date list of supported versions, platforms, and other information, see <http://www.hp.com/support/manuals>

Debugging

Enabling debugging option

1. If you encounter an issue when using this extension, the information in the log files can help you determine the problem. To enable the debugging option, locate the Settings page by clicking the **Settings** button. The Settings page is displayed.
2. Locate the Debugging option, select **Enable Debugging**. When you close the vSphere client interface this option is cleared, and the debugging option for this extension disabled.

For detailed descriptions of the Data Protector log files, see the HP Data Protector Troubleshooting Guide, index: "Contents of log files".

Figure 34 Debugging option enabled



Location of log files:

Windows Server 2008 x64 Edition:

Data_Protector_program_data\tmp

Data_Protector_program_data\log

Windows Server 2003 x64 Edition:

Data_Protector_home\tmp

Data_Protector_home\log

For a list of known issues and workarounds, see "Known issues and workarounds" (page 51).

Known issues and workarounds

Mounting virtual machine disks

Problem

When performing recovery of files, the following message is displayed after selecting the partition:

```
[EXCEPTION] boost: filesystem: status: The volume does not contain a
recognized file system. Please make sure that all required file system
drivers are loaded and that the volume is not corrupted: "\\?\M:\\"
ProxyGetAllNodesForPath.
```

Then after mounting the virtual disks manually, the following is displayed in the command line interface:

```
The volume does not contain a recognized file system. Please make sure
that all required file system drivers are loaded and that the volume
is not corrupted
```

This is an unsupported configuration. The root cause of the problem is that this extension and VMware VDDK do not support GPT disk layout and dynamic disks. This is a known vmware-mount limitation. This issue may occur when a partition contains no filesystem at all, or contains other non-supported filesystem such as a Linux filesystem,- for example from a dual-boot machine.

Action

Use a supported file system, for example NTFS or FAT system formats.

Issues after removing this extension

Problem

You removed the VMware Granular Recovery Extension Agent component and the post-install script. When you remove the mount proxy agent system, the removal tries to delete the driver. If the removal command is not successful, the driver stays in the stopped state. Even after repeating the installation procedure the driver stays in a stopped state.

When you try again to install the VMware Granular Recovery Extension Agent component, the HP Data Protector Manager displays the following message in the monitor session:

```
Status: SERVICE_NAME: vstor2-mntapi10-shared
[Critical] computer.company.com Post-installation script for the VMware
Granular Recovery Extension Agent failed with the output
Data_Protector_home\bin
perl -I "..\lib\perl" vmwgre_ag.pl -install
```

In the command line interface this message is displayed:

```
Installation of driver failed: System error 2 has occurred.
```

Action

Check the driver status by running the command:

```
sc query vstor2-mntapi10-shared
```

Check if the service state is RUNNING. If not, remove the service and run the command:

```
sc delete vstor2-mntapi10-shared
```

Check if the service is removed and start a remote installation of the VMware Granular Recovery Extension Agent component.

NOTE: If the vdkm driver is already installed on the system and running then this driver is used, it is not removed or reinstalled. The upgrade procedure is successful.

VMware VirtualCenter Management Webservices service is not running

Problem

The debug log file contains the following message:

```
File not found - C:\Program Files  
(x86)\VMware\Infrastructure\tomcat\webapps\VMWareGRE\register.xml
```

There is a problem with VMware VirtualCenter Management Webservices service. Check that the target system has VMware VirtualCenter Management Webservices running.

Action

Check if the C:\Program Files

(x86)\VMware\Infrastructure\tomcat\webapps\VMWareGRE folder exists. If it exists, remove it. The VMware VirtualCenter Management Webservices service should create a new folder automatically. Run the script: *Data_Protector_home\bin\vmwgre_wp.cmd*

If folder does not exist, VMware VirtualCenter Management Webservices is probably not started. Start it, check that the folder is created and run the script:

Data_Protector_home\bin\vmwgre_wp.cmd.

Remote installation of VMware Granular Recovery Extension Web Plug-In ends unexpectedly

Problem

The remote installation of VMware Granular Recovery Extension Web Plug-In end unexpectedly. In the debug.log file on the target system the following message is displayed:

```
Log on to Virtual Center computer.company.com could not be performed.  
Web service error: No connection could be made because the target machine  
actively refused it.
```


Or:

```
Log on to Virtual Center computer.company.com could not be performed.  
Details: Web service error: Cannot complete login due to an incorrect  
user name or password.
```

Action

Ensure the user credential information and the correct port number are specified on the following HP Data Protector client import settings page:

Figure 35 User credentials during the import client procedure



VMware Granular Recovery Extension tab is missing

Problem

You connected with vSphere Client to vCenter. When you select a virtual machine in the VMs and Templates view, there is no HP Data Protector tab, the extension is missing. The root cause is the firewall is probably preventing communication.

Action

On the vCenter server, configure the Windows Firewall. Select the Exceptions tab and **Add port** of the VMware vCenter Server-Web Services HTTPS (default 8443) to the exceptions list, and restart the vSphere Client interface.

VMware Granular Recovery Extension tab is missing with vCenter Server plug-in disabled

Problem

You connected with vSphere Client to vCenter. When you select a virtual machine in the VMs and Templates view, there is no HP Data Protector tab, and the extension is missing. The root cause is an installation that ended abnormally.

Action

1. On the vCenter server, remove the extension from the system.
2. Remotely install the extension again. For details, see [“VMware vCenter Server system” \(page 22\)](#).

3. Connect with the vSphere Client interface to a vCenter Server system, and click **Plug-ins**. The Plug-in Manager window is displayed.
Under the Plug-in Name column, locate VMWareGRE, right click it, and click **Enable**.

Figure 36 Extension disabled





Figure 37 Extension enabled



HP Data Protector cannot add any Granular Recovery Extension component

Problem

HP Data Protector cannot add components. The probable root cause is that HP Data Protector was installed on your system in the following order:

1. Data Protector cell and client, was imported with the Virtual Environment Integration agent enabled on it.
2. The HP Data Protector client was imported.

Action

1. Re-import the HP Data Protector client.
2. Import the Virtual Environment Integration agent and select VMware as the client type.

Local installation workaround

Problem

You cannot install the extension remotely on your system.

Action

Install the extension on your local system.

For details of the importing procedure, see [“Importing and configuring VMware clients”](#) (page 22).

Local installation workaround

1. To install this extension on a Data Protector cell system, follow these steps:
Install the VMware Granular Extension Web Plug-In component or the VMware Granular Recovery Extension Agent component by changing the directory to:
Insert the Windows installation DVD-ROM by changing the directory to:

32-bit Windows:

`\Windows_other\i386`

64-bit Windows:

`\Windows_other\x8664`

2. Run the VMware Granular Extension Web Plug-In component installation:

The `\` denotes continuation of the command line.

```
csetup.exe /quiet ADDLOCAL=core,vmwaregre_webplugin \  
INSTALLATIONTYPE="Client" \  
CELLNAME="CellManagerSystemName" \  
CELLCLIENTNAME="VirtualCenterSystemName" \  
INSTALLDIR="InstallationPath"\  
INETPORT=InetPort OPTDNSCHECK=1 OPT_MSG=1 \  
PUSHADDUPG=1
```

Run the VMware Granular Recovery Extension Agent component installation:

```
csetup.exe /quiet ADDLOCAL=core,vmwaregre_agent \  
INSTALLATIONTYPE="Client" \  
CELLNAME="CellManagerSystemName" \  
CELLCLIENTNAME="TargetSystemName" \  
INSTALLDIR="InstallationPath" \  
INETPORT=InetPort OPTDNSCHECK=1 OPT_MSG=1 \  
PUSHADDUPG=1
```

Monitor displays request removed by administrator

Problem

A message, similar to the following is displayed when monitoring your request:

```
0030 scsi0:1 2/17/2011 10:35:07 AM Removed by administrator 3.00
```

Your administrator has removed this disk.

Action

- To recover your files, ensure that you request a new restore for this removed disk.
For details of the procedure, see [“Requesting restores”](#) (page 37).

Overwritten files issues

Problem

A message, similar to the following is displayed when recovering items with the **Overwrite** option selected:

The \ denotes continuation of the command line.

```
[Failed] c:\vix_27-2\incremental-21-2\incremental\ \
Username\CheckVix\vixlibs\arp.ico
Source:\incremental-21-2\incremental\Username\CheckVix\ \
vixlibs\arp.ico
```

You do not have access rights to this file.

```
[Failed] c:\overwrite_incr-21-2\incr24-2\incremental-21-2\ \
incremental\Username\CheckVix\vixlibs\vix.h
Source:\incr24-2\incremental-21-2\incremental\Username\ \
CheckVix\vixlibs\vix.h
```

```
[5] Access is denied.
```

The item already exists on the target system. This item cannot be overwritten due to the file security option. If the source location contains NTFS filesystem and the target virtual machine disk are on the network, the Granular Recovery Extension for VMware recovers all security information associated with the items. This information cannot be overwritten.

Action

If the item already exists in the target location, perform one of the following instead:

- Recover these items to another location.
- Select the **Skip** recovery option.
- Select the **Rename** recovery option.
- Change the file permissions on the target location manually before starting recovery.

Opening the extension with script errors

Problem

After you install the extension and attempt to open the extension interface for the first time, the following Internet Explorer pop-up windows are displayed:

Figure 38 Blocked website message displayed by Internet Explorer


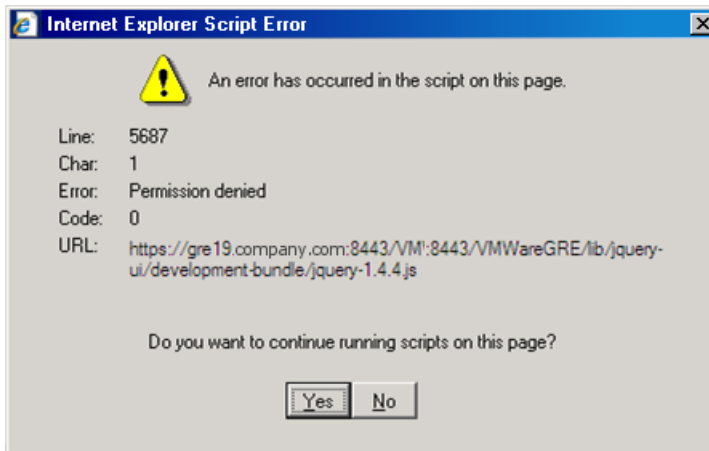


Figure 39 Script error message displayed by Internet Explorer



Additional pop-up windows may appear, containing similar error messages:

An error has occurred in the script on this page. `Offset.setoffset` is null or not an object. Code 0

Object does not support this property or method.

Action

To resolve the problem, lower the security settings and add the web service entry point URI to the list of trusted sites:

1. When the pop-up window displayed on [Figure 38 \(page 57\)](#) appears, click **Add**. The Internet Explorer Script Error is displayed. See “[Script error message displayed by Internet Explorer](#)” ([page 57](#)).
2. When the pop-up window displayed on [Figure 39 \(page 57\)](#) appears, click **Yes** to continue running these scripts. When this error message dialog box is displayed again, click **Yes** to proceed.

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: <ul style="list-style-type: none">• 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 <code>Data_Protector_program_data\Config\Server\dr\asr</code> (Windows Server 2008), <code>Data_Protector_home\Config\Server\dr\asr</code> (other Windows systems), or <code>/etc/opt/omni/server/dr/asr</code> (UNIX systems) as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.
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	<p>A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk).</p> <p>A backup object is defined by:</p> <ul style="list-style-type: none"> • 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	<p>(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).</p> <p>See also application system, target volume, and replica.</p>
backup types	See incremental backup, differential backup, transaction backup, full backup, and delta backup.
backup view	<p>Data Protector provides different views for backup specifications:</p> <p>By Type - according to the type of data available for backups/templates. Default view.</p> <p>By Group - according to the group to which backup specifications/templates belong.</p> <p>By Name - according to the name of backup specifications/templates.</p> <p>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.</p>
BC	<p>(EMC Symmetrix specific term) Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.</p> <p>See also BCV.</p>
BC Process	<p>(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.</p> <p>See also BCV.</p>
BCV	<p>(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.</p> <p>See also BC and BC Process.</p>
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	<p>(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.</p> <p>See also BRBACKUP and BRRESTORE.</p>
BRBACKUP	<p>(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.</p> <p>See also BRARCHIVE and BRRESTORE.</p>
BRRESTORE	<p>(SAP R/3 specific term) An SAP R/3 tool that can be used to restore files of the following type:</p> <ul style="list-style-type: none">• 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 <p>You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.</p> <p>See also BRBACKUP and BRARCHIVE.</p>
BSM	The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

C

CAP	<i>(StorageTek specific term)</i> Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.
catalog protection	Defines how long information about backed up data (such as file names and file versions) is kept in the IDB. <i>See also</i> 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. <i>See also</i> MMDB.
CDF file	<i>(UNIX specific term)</i> 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. <i>See also</i> MoM.
Centralized Media Management Database (CMMDB)	<i>See</i> 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	<i>(Windows specific term)</i> 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	<i>(Windows specific term)</i> A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.
channel	<i>(Oracle specific term)</i> 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: <ul style="list-style-type: none"> • 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	<i>(Microsoft Exchange Server and Lotus Domino Server specific term)</i> 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: <ul style="list-style-type: none"> • 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	<p>(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.</p> <p>A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.</p> <p>See also Exchange Replication Service and local continuous replication.</p>
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	<p>The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended</p> <p>See also MoM.</p>
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	<p>(VLS specific term) A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.</p> <p>See also Virtual Library System (VLS).</p>
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	<p>(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.</p> <p>See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.</p>
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 blob space, dspace, 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. <i>See also</i> incremental backup.
differential backup	<i>(Microsoft SQL Server specific term)</i> A database backup that records only the data changes made to the database after the last full database backup. <i>See also</i> backup types.
differential database backup	A differential database backup records only those data changes made to the database after the last full database backup.
directory junction	<i>(Windows specific term)</i> Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.
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	<i>See</i> 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	<i>(Veritas Volume Manager specific term)</i> The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.
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. <i>See also</i> 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	<i>(Informix Server specific term)</i> The Informix Server configuration file <code>ixbar.server_id</code> that resides in the directory <code>INFORMIXDIR/etc</code> (on Windows) or <code>INFORMIXDIR\etc</code> (on UNIX). <code>INFORMIXDIR</code> is the Informix Server home directory and <code>server_id</code> is the value of the <code>SERVERNUM</code> 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. <code>KeyID</code> identifies the key within the keystore. <code>StoreID</code> identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may several <code>StoreIDs</code> used on the same Cell Manager.
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. <i>See also MoM.</i>

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 <code>Data_Protector_program_data\log\server\Ob2EventLog.txt</code> (Windows Server 2008), <code>Data_Protector_home\log\server\Ob2EventLog.txt</code> (other Windows systems), or <code>/var/opt/omni/server/log/Ob2EventLog.txt</code> (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	(<i>Windows specific term</i>) Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.
Exchange Replication Service	(<i>Microsoft Exchange Server specific term</i>) The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.
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)	(<i>Microsoft Exchange Server specific term</i>) 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	(<i>HP P6000 EVA Disk Array Family specific term</i>) An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations. See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.
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	(<i>Windows specific term</i>) 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	<i>(HP P9000 XP Disk Array Family specific term)</i> A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used. See also primary volume and mirror unit (MU) number.
flash recovery area	<i>(Oracle specific term)</i> 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 <code>fnames.dat</code> 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 <code>Data_Protector_program_data\Config\Server\Options</code> (Windows Server 2008), <code>Data_Protector_home\Config\Server\Options</code> (other Windows systems), or <code>/etc/opt/omni/server/options</code> (HP-UX, Solaris, and Linux systems).
group	<i>(Microsoft Cluster Server specific term)</i> 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	<i>(Microsoft Exchange Server specific term)</i> 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 <i>Data_Protector_program_data\Config\Server\holidays</i> (Windows Server 2008), <i>Data_Protector_home\Config\Server\holidays</i> (other Windows systems), or <i>/etc/opt/omni/server/Holidays</i> (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	<i>(HP P6000 EVA Disk Array Family specific term)</i> 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	<i>(HP P9000 XP Disk Array Family specific term)</i> 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	<i>(HP P6000 EVA Disk Array Family specific term)</i> 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	<i>(HP P9000 XP Disk Array Family specific term)</i> 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	<i>(HP P6000 EVA Disk Array Family specific term)</i> 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.
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	(<i>Microsoft Exchange Server specific term</i>) The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. See also Key Management Service and Site Replication Service.
Informix Server	(<i>Informix Server specific term</i>) 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	(<i>ZDB specific term</i>) A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery. 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)	(<i>Windows specific term</i>) Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).
ISQL	(<i>Sybase specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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	<i>(EMC Symmetrix specific term)</i> A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.
LDEV	<i>(HP P9000 XP Disk Array Family specific term)</i> A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. <i>See also</i> HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.
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	<i>(Oracle specific term)</i> 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	<i>(Microsoft Exchange Server specific term)</i> Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. 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. <i>See also</i> 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	(<i>Informix Server UNIX specific term</i>) A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the <code>INFORMIXDIR/etc/log_full.sh</code> , where <code>INFORMIXDIR</code> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code> .
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	(<i>Microsoft SQL Server specific term</i>) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.
login information to the Oracle Target Database	(<i>Oracle and SAP R/3 specific term</i>) The format of the login information is <code>user_name/password@service</code> , where: <ul style="list-style-type: none">• <code>user_name</code> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.• <code>password</code> must be the same as the password specified in the Oracle password file (<code>orapwd</code>), which is used for authentication of users performing database administration.• <code>service</code> is the name used to identify an SQL*Net server process for the target database.
login information to the Recovery Catalog Database	(<i>Oracle specific term</i>) The format of the login information to the Recovery (Oracle) Catalog Database is <code>user_name/password@service</code> , where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <code>service</code> 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	(<i>Lotus Domino Server specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text <code>.edb</code> file and a streaming native internet content <code>.stm</code> file.
Main Control Unit (MCU)	(<i>HP P9000 XP Disk Array Family specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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)	(<i>Windows specific term</i>) An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.
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	(<i>HP P9000 XP Disk Array Family specific term</i>) A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family. See also first-level mirror.
mirrorclone	(<i>HP P6000 EVA Disk Array Family specific term</i>) A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.
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	(<i>HP P6000 EVA Disk Array Family specific term</i>) 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.
○	
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	(<i>Windows specific term</i>) 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	(<i>Informix Server specific term</i>) A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: <ul style="list-style-type: none"> • 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	(<i>Informix Server specific term</i>) An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file in the directory <i>INFORMIXDIR/etc</i> (on Windows) or <i>INFORMIXDIR/etc/</i> (on UNIX).

online backup	<p>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.</p> <p>In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.</p> <p>See also zero downtime backup (ZDB) and offline backup.</p>
online recovery	<p>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).</p>
online redo log	<p>(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.</p> <p>See also archived redo log.</p>
Oracle Data Guard	<p>(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.</p>
Oracle instance	<p>(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.</p>
ORACLE_SID	<p>(Oracle specific term) A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <code>ORACLE_SID</code>. The <code>ORACLE_SID</code> is included in the <code>CONNECT DATA</code> parts of the connect descriptor in a <code>TNSNAMES.ORA</code> file and in the definition of the TNS listener in the <code>LISTENER.ORA</code> file.</p>
original system	<p>The system configuration backed up by Data Protector before a computer disaster hits the system.</p>
overwrite	<p>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.</p> <p>See also merging.</p>
ownership	<p>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.</p> <p>If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.</p> <p>If a modified backup specification is started by a user, the user is the owner unless the following is true:</p> <ul style="list-style-type: none"> • 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. <p>If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.</p> <p>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.</p> <p>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>

P

P1S file	<p>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</p>
-----------------	--

	<p>backup medium and on Cell Manager into the directory <i>Data_Protector_program_data\Config\Server\dr\pls</i> (Windows Server 2008), <i>Data_Protector_home\Config\Server\dr\pls</i> (other Windows systems), or <i>/etc/opt/omni/server/dr/pls</i> (UNIX systems) with the filename <i>recovery.pls</i>.</p>
package	<p>(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.</p>
pair status	<p>(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:</p> <ul style="list-style-type: none"> • 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	<p>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.</p>
parallelism	<p>The concept of reading multiple data streams from an online database.</p>
phase 0 of disaster recovery	<p>Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.</p>
phase 1 of disaster recovery	<p>Installation and configuration of DR OS, establishing previous storage structure.</p>
phase 2 of disaster recovery	<p>Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.</p>
phase 3 of disaster recovery	<p>Restoration of user and application data.</p>
physical device	<p>A physical unit that contains either a drive or a more complex unit such as a library.</p>
post-exec	<p>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.</p>
pre- and post-exec commands	<p>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.</p>
pre-exec	<p>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.</p>
prealloc list	<p>A subset of media in a media pool that specifies the order in which media are used for backup.</p>
primary volume (P-VOL)	<p>(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</p>

	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	(<i>Microsoft Exchange Server specific term</i>) 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: <ul style="list-style-type: none"> • 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	(<i>HP P9000 XP Disk Array Family specific term</i>) 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	(<i>HP P9000 XP Disk Array Family specific term</i>) 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	(<i>EMC Symmetrix specific term</i>) A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.
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	(<i>Oracle specific term</i>) A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about: <ul style="list-style-type: none"> • 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	(<i>Oracle specific term</i>) An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.
recovery files	(<i>Oracle specific term</i>) Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area.
Recovery Manager (RMAN)	(<i>Oracle specific term</i>) An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

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	<i>(Oracle specific term)</i> Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.
Remote Control Unit (RCU)	<i>(HP P9000 XP Disk Array Family specific term)</i> An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.
Removable Storage Management Database	<i>(Windows specific term)</i> A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.
reparse point	<i>(Windows specific term)</i> A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.
replica	<i>(ZDB specific term)</i> An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on 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. <i>See also</i> snapshot, snapshot creation, split mirror, and split mirror creation.
replica set	<i>(ZDB specific term)</i> A group of replicas, all created using the same backup specification. <i>See also</i> replica and replica set rotation.
replica set rotation	<i>(ZDB specific term)</i> 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. <i>See also</i> 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	<i>(HP P9000 XP Disk Array Family VSS provider specific term)</i> One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the 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. <i>See also</i> 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)	<i>See</i> 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	(<i>Windows specific term</i>) 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	(<i>SAP R/3 specific term</i>) 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)	(<i>HP P9000 XP Disk Array Family specific term</i>) 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	(<i>Microsoft VSS specific term</i>) 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	(<i>Microsoft VSS specific term</i>) 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	(<i>Microsoft VSS specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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	(<i>VLS specific term</i>) 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	<i>(VLS specific term)</i> 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 LDB 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	<i>(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)</i> 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	<i>(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)</i> 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	<i>(EMC Symmetrix specific term)</i> 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	<i>(ZDB specific term)</i> 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	<i>(EMC Symmetrix Disk Array and HP P9000 XP Disk Array Family specific term)</i> 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	<i>(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</i> 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	<i>(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)</i> 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	<i>(Informix Server specific term)</i> 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	<i>(disaster recovery specific term)</i> 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	<i>(EMC Symmetrix specific term)</i> 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 <code>/usr/kernel/drv/sst.conf</code> 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 <code>/kernel/drv/st.conf</code> is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.
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	<i>(Microsoft Exchange Server specific term)</i> 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	<i>(ZDB specific term)</i> 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	<i>(StorageTek specific term)</i> 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	<i>(Sybase specific term)</i> An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.
Sybase SQL Server	<i>(Sybase specific term)</i> The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.
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	<i>(Oracle specific term)</i> An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases	<p>(<i>Sybase specific term</i>) The four system databases on a newly installed Sybase SQL Server are the:</p> <ul style="list-style-type: none"> • master database (master) • temporary database (tempdb) • system procedure database (sybsystemprocs) • model database (model).
System Recovery Data file	See SRD file.
System State	<p>(<i>Windows specific term</i>) The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.</p>
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	<p>(<i>Windows specific term</i>) A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.</p>
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	<p>(<i>EMC Symmetrix specific term</i>) An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.</p> <p>See also source (R1) device.</p>
target database	<p>(<i>Oracle specific term</i>) In RMAN, the target database is the database that you are backing up or restoring.</p>
target system	<p>(<i>disaster recovery specific term</i>) A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.</p>
target volume	(<i>ZDB specific term</i>) A storage volume to which data is replicated.
Terminal Services	<p>(<i>Windows specific term</i>) 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.</p>
thread	<p>(<i>Microsoft SQL Server specific term</i>) 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.</p>
TimeFinder	<p>(<i>EMC Symmetrix specific term</i>) 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).</p>
TLU	Tape Library Unit.
TNSNAMES.ORA	<p>(<i>Oracle and SAP R/3 specific term</i>) 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.</p>

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	<i>(Sybase and SQL specific term)</i> 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	<i>(Sybase specific term)</i> A system table in which all changes to the database are automatically recorded.
transaction logs	<i>(Data Protector specific term)</i> 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	<i>(Microsoft VSS specific term)</i> 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	<i>(Novell NetWare specific term)</i> 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 <code>SYS:SYSTEM\TSA</code> directory on the server where <code>TSANDS.NLM</code> is loaded.

U

UIProxy	The Java GUI Server (<code>UIProxy</code> 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	<i>(Windows specific term)</i> 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 <i>admin</i> and <i>operator</i> .
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)	(<i>HP P6000 EVA Disk Array Family specific term</i>) The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. See also HP Command View (CV) EVA.
Virtual Device Interface	(<i>Microsoft SQL Server specific term</i>) This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.
virtual disk	(<i>HP P6000 EVA Disk Array Family specific term</i>) A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. 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	(<i>VLS specific term</i>) An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. See also Virtual Library System (VLS) and Virtual Tape Library (VTL).
Virtual Tape Library (VTL)	(<i>VLS specific term</i>) An emulated tape library that provides the functionality of traditional tape-based storage. See also Virtual Library System (VLS).
VMware management client	(<i>VMware (Legacy) integration specific term</i>) The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).
volser	(<i>ADIC and STK specific term</i>) A VOLUME SERIAL number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.
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	(<i>Windows specific term</i>) An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.
Volume Shadow Copy Service	See Microsoft Volume Shadow Copy Service (VSS).
VSS	See Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode	<i>(HP P9000 XP Disk Array Family VSS provider specific term)</i> One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching the disks. <i>See also</i> 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	<i>(Microsoft VSS specific term)</i> A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.
X	
XBSA interface	<i>(Informix Server specific term)</i> ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).
Z	
ZDB	<i>See</i> zero downtime backup (ZDB).
ZDB database	<i>(ZDB specific term)</i> A part of the IDB, storing ZDB-related information such as source volumes, replicas, and security information. The ZDB database is used in zero downtime backup, instant recovery, and split mirror restore sessions. <i>See also</i> zero downtime backup (ZDB).
ZDB to disk	<i>(ZDB specific term)</i> A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process. <i>See also</i> zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.
ZDB to disk+tape	<i>(ZDB specific term)</i> A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using

the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore.

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

- access points
 - administrative, operational, 15
- administrative
 - monitoring requests, 46
 - restore requests, for operators, 44
- administrator's
 - tasks, roles, 20
- administrators
 - mount proxy, configuring, 44
- allowing
 - restore sessions, for operators, 44
- audience, 6

B

- backup, 31
- backup solutions, 14

C

- clear overview
 - restore, recovery, 15
- components
 - mount proxy, system, 19
- configuring, 21
 - Data Protector, 21
 - Data Protector user rights, 28
 - firewall, 29
 - mount proxy systems, 44
 - user groups, 24
 - Virtual Center, 30
 - VMware vCenter Server, 22
- considerations
 - before installing, 17
- conventions
 - document, 11

D

- Data Protector, 24
- debugging
 - before you begin, 50
 - enabling, 50
- document
 - conventions, 11
 - related documentation, 6
- documentation
 - HP website, 6
 - providing feedback, 13

E

- entry point
 - operators, 35

F

- firewall, 29

G

- granular recovery
 - starting, administrators, 43
- granularity, 14

H

- handling
 - restore requests, administrators, 44
- help
 - obtaining, 12
- HP
 - technical support, 12
- HP Data Protector
 - add component, 54
- HP Data Protector Granular Recovery Extension
 - administrators, operators, 20
 - environment, 17
 - installing, 16
 - users, roles, 20
 - users, system, 20
 - users, tasks, 20

I

- identification
 - users, 15
- importing
 - configuring, VMware vCenter Server, 22
- installation, 16
- installing
 - HP Data Protector Granular Recovery Extension, 16
- integration
 - VMware vCenter Server, 14

K

- known issues
 - workarounds, 51

L

- local installation
 - workaround, 55
- low learning
 - curve, 14

M

- monitoring
 - restore requests, 15, 46
 - restore requests, operators, 39
- mount proxy
 - system, 19
 - system, adding Data Protector components, 23
- mounting
 - virtual machine disks, 51

O

- opening

- extension, administrative entry point, [43](#)
- extension, administrators, [43](#)
- HP Data Protector Granular Recovery Extension for VMware vSphere GUI, [35](#)

- operative
 - entry point, [35](#)

- operator's
 - tasks, roles, [20](#)

- operators
 - tasks, [35](#)

- options, [30](#)

P

- prerequisites
 - before installing, [16](#)

Q

- quickness
 - ease, [15](#)

R

- recover granularity capability, [14](#)

- recovering
 - virtual machines, operators, [40](#)

- recovery
 - different locations, [15](#)
 - granularity, [14](#)
 - restore, [32](#)
 - restore flow, [33](#)

- rejecting
 - restore requests, for operators, [46](#)

- related documentation, [6](#)

- remote installatio
 - VMware Granular Recovery Extension Web Plug-In, [52](#)

- removal
 - components, [51](#)
 - disks, [47](#)

- removing
 - virtual machine disks, [47](#)

- requesting restores
 - operators, [37](#)

- restore requests
 - handling, administrative, [44](#)

- restoring (in progress)
 - stop, status, [47](#)

S

- security, [15](#)

- space-saving
 - approach, [14](#)

- stopping
 - restore sessions, administrative, [47](#)

- Subscriber's Choice, HP, [12](#)

T

- technical support
 - HP, [12](#)
 - service locator website, [12](#)

- triggering

- restore sessions, administrative, [45](#)
- troubleshooting, [49](#)
 - before you begin, [49](#)
 - request removed by administrator, [55–56](#)

U

- user rights
 - configuring, [28](#)

- user-directed
 - restore, recovery, [14](#)

V

- vCenter users
 - requesting restores, [37](#)

- virtual disks, [14](#)

- VMware Granular Recovery Extension tab
 - missing, [53](#)

- VMware users
 - tasks, [35](#)

- VMware vCenter Server
 - system, [18](#)

- VMware VirtualCenter Management Webservices service
 - not running, [52](#)

W

- websites
 - HP, [12](#)
 - HP Subscriber's Choice for Business, [12](#)
 - product manuals, [6](#)