

HP Data Protector A.06.10

Integration guide for Microsoft applications

SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service
[Build 500]



T B D

Part number: TBD
First edition: TBD Month 2008



i n v e n t

Legal and notice information

© Copyright 2004, 2008 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 US trademark of Sun Microsystems, Inc.

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

UNIX is a registered trademark of The Open Group.

Printed in the US

Contents

Publication history	13
About this guide	15
Intended audience	15
Documentation set	15
Guides	15
Online help	18
Documentation map	19
Abbreviations	19
Map	20
Integrations	21
Document conventions and symbols	22
Data Protector graphical user interface	23
General Information	24
HP technical support	24
Subscription service	25
HP websites	25
Documentation feedback	25
1 Integrating Microsoft SQL Server and Data Protector	27
Introduction	27
Integration concepts	28
Parallelism	29
Configuring the integration	30
Prerequisites	30
Before you begin	31
Data Protector SQL Server configuration file	31
Configuring users	32
Configuring SQL Server	33
Using the Data Protector GUI	34
Using the Data Protector CLI	36
Changing and checking configuration	36
Using the Data Protector GUI	36

Using the Data Protector CLI	37
Backup	38
Creating backup specifications	38
SQL Server Specific Backup Options	41
Object specific options	42
Scheduling backups	43
Scheduling example	43
Starting backup sessions	44
Using the Data Protector GUI	44
Restore	44
Restore methods	44
Before you begin	44
Restoring using the Data Protector GUI	45
Restore options	48
Restoring to another SQL Server instance or/and another SQL Server	50
Restoring using the Data Protector CLI	51
Disaster recovery	51
Recovering the master database	52
Recovering user databases	53
Performance tuning	54
Monitoring sessions	58
Troubleshooting	58
Before you begin	59
Checks and verifications	59
Problems	60

2 Integrating Microsoft SharePoint Portal Server and Data Protector 65

Introduction	65
Integration concepts	66
Backup and restore flow	69
Configuring the integration	69
Prerequisites	70
Before you begin	70
Configuring SPS users	70
Configuring SPS farms	70
Checking the configuration	73
Backup	73
Before you begin	74
Creating backup specifications	74
Modifying backup specifications	77
Scheduling backup specifications	77

Scheduling example	77
Previewing backup sessions	78
Using the Data Protector GUI	78
Using the Data Protector CLI	79
What happens during the preview?	79
Starting backup sessions	79
Before you begin	79
Using the Data Protector GUI	79
Using the Data Protector CLI	79
Preparing for disaster recovery	80
Restore	81
Before you begin	81
Considerations	81
Restoring using the Data Protector GUI	81
Restoring using the Data Protector CLI	86
Restoring using another device	88
Using the Data Protector GUI	88
Using the Data Protector CLI	88
Disaster recovery	89
Monitoring sessions	89
Troubleshooting	89
Before you begin	90
Checks and verifications	90
Problems	90

3 Integrating Microsoft Exchange Server and Data Protector 93

Introduction	93
Integration concepts	94
Configuring the integration	95
Prerequisites	95
Limitations	96
Before you begin	96
Backup	97
Configuring Exchange Server Backup	98
Creating backup specifications	98
Exchange Server specific backup options	103
Scheduling backups	104
Scheduling example	104
Starting backup sessions	104
Using the Data Protector GUI	104
Restore	105
Restoring using the GUI	106

Restoring to another client	110
Restoring using the CLI	111
Troubleshooting	112
Before you begin	112
Checks and verifications	112
Problems	113

4 Integrating Microsoft Exchange Single Mailbox and Data Protector 115

Introduction	115
Integration concepts	116
Configuring the integration	117
Prerequisites	117
Limitations	118
Before you begin	118
Cluster-aware clients	118
Configuring Exchange Server users	118
Configuring Exchange servers	118
Checking the configuration	119
Backup	120
Creating backup specifications	120
Modifying backup specifications	124
Scheduling backup specifications	124
Scheduling example	124
Previewing backup sessions	125
Using the Data Protector GUI	125
Using the Data Protector CLI	125
What happens during the preview?	126
Starting backup sessions	126
Using the Data Protector GUI	126
Using the Data Protector CLI	126
Restore	127
Before you begin	127
Restoring using the Data Protector GUI	127
Restoring using the Data Protector CLI	133
Restore examples	135
Monitoring sessions	136
Performance tuning	136
Troubleshooting	137
Before you begin	138
Checks and verifications	138
Problems	139

5 Integrating Microsoft Volume Shadow Copy Service with Data Protector 143

Introduction	143
Integration concepts	147
Backup	148
Data consistency	149
Restore	149
Restoring components	150
Restoring files	150
Microsoft Exchange Server 2007 writer concepts	151
Backup	151
Restore	151
Prerequisites and limitations	152
Prerequisites	152
Common prerequisites	152
Limitations	152
Common limitations	152
Configuring the integration	153
Configuring the Data Protector VSS cluster-aware integration	153
Configuring a VSS cluster-aware client	154
Configuring backups for a cluster-aware VSS client	154
Microsoft Exchange Server writer specific configuration	155
Microsoft Exchange Server 2007 configuration in LCR environment	155
Writers specifics	155
Backing up writers data	159
Overview	159
Overview	160
Creating backup specifications using GUI	160
VSS specific backup options	163
Microsoft Exchange Server writer backup specifics	163
Microsoft Data Protection Manager 2006 writer specifics	167
Scheduling the backup	169
Running an interactive backup	170
Restoring writers data	171
Restore procedure	172
Restore options	177
MSDE writer restore specifics	177
Microsoft Exchange Server writer restore specifics	179
Rollforward recovery from the loss of one or more databases	179
Point-in-time restore after loss of a log file	180
Microsoft Exchange Server 2007 writer restore specifics	181
Microsoft Data Protection Manager 2006 writer restore specifics	184

Restore the DPM server first	184
Restore the DPM clients directly	187
Monitoring a VSS backup and restore	187
Monitoring current sessions	187
Viewing previous sessions	188
Troubleshooting	189
Before you begin	189
Checks and verifications	190
Backup problems	190
Restore problems	192
Glossary	195
Index	251

Figures

1 Data Protector graphical user interface	24
2 Data Protector SQL Server integration architecture	29
3 Database parallelism = 4, Overall Concurrency = 10	30
4 SQL Server users	33
5 Configuring SQL Server	35
6 Checking configuration	37
7 Selecting a blank template	39
8 Selecting backup objects	40
9 Saving a backup specification	41
10 Application specific options	42
11 Object properties	43
12 Restore objects	46
13 Selecting object specific options	46
14 Restore options	47
15 Adjusted local device	56
16 Advanced options	57
17 Block sSizes not adjusted	58
18 SPS integration architecture	67
19 Selecting a front-end Web server system	72
20 Configuring an SPS farm	73
21 Selecting SPS objects	75
22 Application specific options	76
23 Scheduling a backup specification	78
24 Selecting SPS objects for restore	82

25	Specifying the restore destination for an SQL Server database	83
26	Specifying the restore destination for an index server	84
27	SPS restore options	85
28	Data Protector Exchange Server integration architecture	95
29	Path system variable	97
30	Selecting a blank template	99
31	Client name and application database	100
32	Backup objects	101
33	Backup devices	102
34	Application specific options	103
35	Unmounting a database	105
36	Restore objects	107
37	Selecting a backup version	108
38	Restore options	110
39	Exchange Single Mailbox integration architecture	116
40	Configuring the Exchange Server	119
41	Selecting a template	121
42	Selecting Exchange items for backup	122
43	Exchange Single Mailbox specific backup options	123
44	Scheduling a backup specification	125
45	Selecting Exchange items for restore	128
46	Version properties	129
47	Advanced properties	129
48	Selecting the destination Exchange Server system	130
49	Restored mailbox and public folders content with the restore to new folder option selected.	131
50	Example of backup with concurrency 1	137
51	Example of backup with concurrency 2	137
52	Specifying the Exchange Server administrator's mailbox	141

53	Actors of the traditional backup model	145
54	Actors of the Data Protector VSS integration backup model	145
55	Local or network VSS backup	149
56	Selecting backup objects (Microsoft Exchange Server 2007 CCR)	161
57	Additional options for Microsoft Exchange Server 2007 in CCR environment	162
58	Selecting Microsoft Exchange Server 2003 storage groups	165
59	Selecting a replicated Microsoft Exchange Server 2007 storage group	166
60	Selecting Microsoft Data Protection Manager database and replicas	169
61	Scheduling a backup	170
62	Restore objects	173
63	Restore to different location options (Exchange Server 2007 Writer)	176
64	MSDE writer	178
65	Selecting Microsoft Exchange Server 2007 stores for rollforward recovery	180
66	Selecting Microsoft Exchange Server 2003 stores for point-in-time restore	181
67	Restoring the Microsoft Data Protection Manager 2006 client	186
68	Monitoring a current session	188
69	Viewing a previous session	189

Tables

1	Edition history	13
2	Document conventions	22
3	SQL Server online backup types	27
4	Legend	29
5	SQL Server backup options	41
6	Object specific options	43
7	SQL Server restore options	48
8	Backup types	65
9	Legend	67
10	SPS objects	68
11	Backup types	74
12	SPS backup options	77
13	What must be backed up	80
14	SPS restore options	86
15	Exchange Server online backup types	93
16	Legend	95
17	Application specific options	103
18	Exchange Server restore options	109
19	Exchange Single Mailbox backup types	115
20	Exchange Single Mailbox backup types	120
21	Exchange Single Mailbox specific backup options	123
22	Exchange Single Mailbox restore options	131
23	Writer description	156
24	VSS specific backup options	163

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
B6960-90108	October 2004	Data Protector Release A.05.50
B6960-96007	July 2006	Data Protector Release A.06.00
TBD	TBD	Data Protector Release A.06.10

About this guide

This guide describes how to configure and use Data Protector with Microsoft applications.

Intended audience

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

- Basic Data Protector functionality
- Database administration

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

Documentation set

Other documents and online Help provide related information.

Guides

Data Protector guides are available in printed format and in PDF format. Install the PDF files during the Data Protector setup procedure by selecting the `English documentation` and `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 Business Support Center website:

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

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 four guides:

- *HP Data Protector integration guide for Microsoft applications: SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service*

This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server, Microsoft SQL Server, and Volume Shadow Copy Service.

- *HP Data Protector integration guide for Oracle and SAP*

This guide describes the integrations of Data Protector with Oracle, SAP R3, and SAP DB.

- *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, and Lotus Notes/Domino Server.

- *HP Data Protector integration guide for Sybase, Network Node Manager, Network Data Management Protocol, and VMware*

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

- *HP Data Protector integration guide for HP Service Information Portal*

This guide describes how to install, configure, and use the integration of Data Protector with HP Service Information Portal. It is intended for backup administrators. It discusses how to use the OpenView applications for Data Protector service management.

- *HP Data Protector integration guide for HP Reporter*

This manual describes how to install, configure, and use the integration of Data Protector with HP Reporter software. It is intended for backup administrators. It discusses how to use the applications for Data Protector service management.

- *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 software software 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 software software and HP Service Navigator on Windows.

There are two versions of the guide:

- for OVO 7.1x, 7.2x
- for OVO 7.5

- *HP Data Protector software integration guide for HP Performance Manager software and HP Performance Agent software*

This guide provides information about how to monitor and manage the health and performance of the Data Protector environment with HP Performance Manager (PM) software and HP Performance Agent (PA) software on Windows, HP-UX, Solaris and Linux.

- *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 StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array, EMC Symmetrix Remote Data Facility and TimeFinder, and HP StorageWorks Disk Array XP. 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, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases. The guide also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.
- *HP Data Protector MPE/iX system user guide*
This guide describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.
- *HP Data Protector Media Operations user guide*
This guide provides tracking and management of offline storage media. It is intended for network administrators responsible for maintaining and backing up systems. 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 A.06.10. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at <http://www.hp.com/support/manuals>
There are also four other *Product announcements, software notes and references*, which serve a similar purpose for the following:
 - HP Operations Manager software software UNIX integration
 - HP Operations Manager software software Windows integration
 - HP Service Information Portal and HP Reporter software
 - HP Performance Manager software and HP Performance Agent integration
 - HP Media Operations

Online help

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

You can access the online help from the top-level directory on the installation DVD without installing Data Protector:

- **Windows:** Unzip `DP_help.zip` and 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
Help	Online Help
IG-IBM	Integration guide—IBM applications
IG-MS	Integration guide—Microsoft applications
IG-O/S	Integration guide—Oracle, SAP R/3, and SAP DB/MaxDB
IG-OMU	Integration guide—HP Operations Manager software software, UNIX
IG-OMW	Integration guide—HP Operations Manager software software, Windows
IG-PM/PA	Integration guide—Performance Manager and Performance Agent software
IG-Report	Integration guide—HP Reporter software
IG-SIP	Integration guide—HP Service Information Portal
IG-Var	Integration guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and licensing guide
MO GS	Media Operations getting started guide

Abbreviation	Guide
MO RN	Media Operations product announcements, software notes, and references
MO UG	Media Operations user guide
MPE/iX	MPE/iX system user guide
PA	Product announcements, software notes, and references
Trouble	Troubleshooting guide
ZDB Admin	ZDB administrator's guide
ZDB Concpt	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.

Integrations

Look in these guides for details of the following integrations:

Integration	Guide
HP Operations Manager software software	IG-OMU, IG-OMW
HP Performance Manager software	IG-PM/PA
HP Performance Agent software	IG-PM/PA
HP Reporter Light	IG-OMW
HP Reporter software	IG-R
HP Service Information Portal	IG-SIP
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX System	MPE/iX
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var

Integration	Guide
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG
Sybase	IG-Var
Symmetrix (EMC)	all ZDB
VMware	IG-Var

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: Table 2 on page 22	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
<i>Italic text</i>	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 text</i>	<ul style="list-style-type: none"> • Code variables • Command variables
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 or the Data Protector Java GUI. Refer to the online Help for information about the Data Protector graphical user interface.




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/service_locator
- <http://www.hp.com/support/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 AppRM.DocFeedback@hp.com. All submissions become the property of HP.

1 Integrating Microsoft SQL Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Microsoft SQL Server integration. It describes the concepts and methods you need to understand to back up and restore the Microsoft SQL Server (**SQL Server**) database objects.

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

Table 3 SQL Server online backup types

Full database backup	Includes all data regardless of the changes made after the last backup.
Transaction log backup	Uses fewer resources than database backups, so can be created more frequently. By applying transaction log backups, you can recover the database to a specific point in time.
Differential database backup	Records only changes made to the database since the last full database backup. By creating differential backups more frequently than full database backups, you can conserve the media used for backup.

Data Protector offers different restore types, depending on your needs. You can select point-in-time restore, full database restore, as well as restore your SQL Server data to a new location, to another SQL Server, or another SQL Server instance. For detailed information, see [“Restore options”](#) on page 48.

This chapter provides information specific to this integration. For general Data Protector procedures and options, see the online Help.

Integration concepts

Data Protector integrates with SQL Server through the Data Protector `sql_bar.exe` executable, installed on SQL Server. It implements multiple virtual devices for backup and restore and transforms SQL Server Virtual Device Interface (VDI) commands from SQL Server into Data Protector backup or restore streams.

The VDI architecture allows the Data Protector General Media Agent to access data directly in the SQL Server memory, provided the devices are directly attached to SQL Server. Therefore, high backup and restore speed is achieved.

You can perform interactive and scheduled full database backups, differential database backups, and transaction log backups. Full and differential backups, combined with regular transaction log backups, prevent data loss if a disk failure occurs. Furthermore, transaction log backups are needed to perform point-in-time restore.

You can back up the whole server or particular databases listed below:

User databases	Contain user data.
Master	Controls user databases and the SQL Server operation. Keeps track of user accounts, configurable environment variables, and system error messages.
Model	Provides a template or prototype for new user databases.
Distribution	A system database used by SQL Server replication components, such as Distribution Agent, to store data, including transactions, snapshot jobs, synchronization status, and replication history information.
MsdB	Provides storage for scheduling and backup information.

For more information about system databases, see the SQL Server documentation.

Data Protector restores databases so that the last differential backup is applied to the most recent full backup. Then the transaction log backups are applied according to the specified restore options.

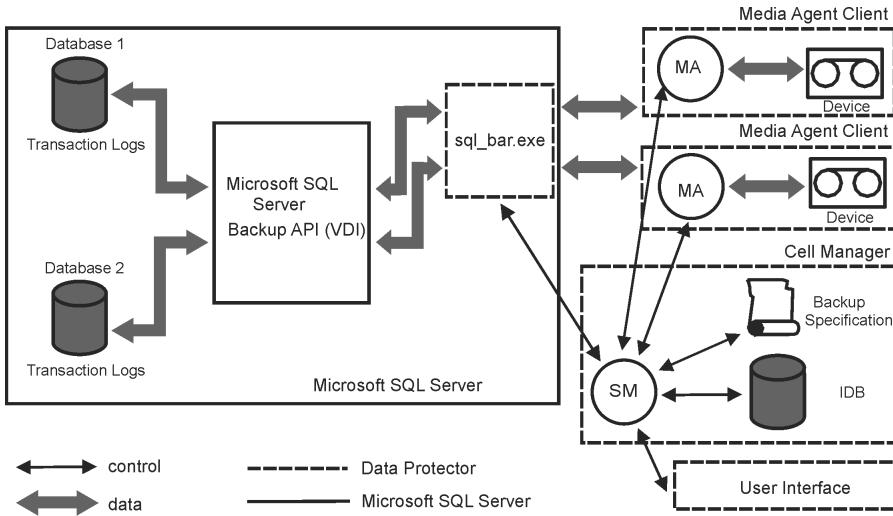


Figure 2 Data Protector SQL Server integration architecture

Table 4 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
Backup API or VDI	SQL Server VDI, the backup interface introduced with SQL Server 7.0.
MA	Data Protector General Media Agent.

Parallelism

You can back up more than one SQL Server database at a time or back up a single database using multiple streams.

Parallelism types used with SQL Server are:

- Database parallelism
More than one database is backed up if the number of available devices allows to perform backups in parallel.
The allocation of streams to available devices is done automatically.
- Number of concurrent streams

This is a number of devices used to back up a particular database or a server. Can be specified by the user or calculated automatically.

 **NOTE:**

SQL Server cannot back up multiple streams to one device.

Figure 3 on page 30 shows a session with four SQL Server databases backed up using a different number of concurrent streams per database.

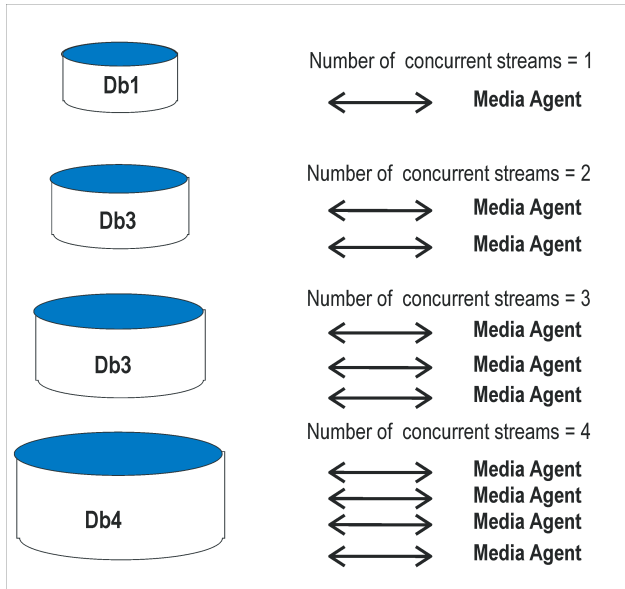


Figure 3 Database parallelism = 4, Overall Concurrency = 10

Configuring the integration

Prerequisites

- You need a license to use the SQL Server integration. See the *HP Data Protector installation and licensing guide* for information.
- Ensure that you correctly installed and configured SQL Server.

- For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or <http://www.hp.com/support/manuals>.
- For information on installing, configuring, and using SQL Server, see the SQL Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing the Data Protector SQL Server integration, see the *HP Data Protector installation and licensing guide*.
Every SQL Server to be used with Data Protector must have the MS SQL Integration component installed.
- Install SQL Server on the application system. Install user databases on the disk array source volumes (system databases can be installed anywhere). If the system databases are also installed on a disk array, they *must* reside on *different* source volumes than user databases.

If SQL Server is installed on the backup system as well, its databases *must* reside on source volumes *different* from the source volumes used for this integration. Drive letters/mount points assigned to those volumes must also be different from the drive letters/mount points assigned to the volumes used for this integration.

Before you begin

- Configure devices and media for use with Data Protector. See the online Help index: “configuring devices” and “creating media pools” for instructions. See also “[Performance tuning](#)” on page 54 for advanced options.
- On the SQL Server client, restart the Data Protector `Inet` service under a Windows domain user account that has Windows administrative rights on the client. Ensure that this user has appropriate SQL Server permissions for running backups and restores if you plan to use the Integrated authentication type. For more information on authentication types, see “[Configuring SQL Server](#)” on page 33. For information on changing the user account under which the Data Protector `Inet` service is running, see the online Help index: “Inet, changing account”.
- To test whether SQL Server and Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore. See the online Help for instructions.

Data Protector SQL Server configuration file

Data Protector stores integration parameters for every configured SQL Server on the Cell Manager in:

HP-UX, Solaris:

```
/etc/opt/omni/server/integ/config/MSSQL/client_name%instance_name
```

Windows:

```
Data_Protector_home\Config\Server\Integ\Config\MSSQL\client_name%instance_name
```

Configuration parameters are the username and password of the SQL Server user, who must have permissions to run backups and restores within SQL Server (assuming the standard security is used). They are written to the Data Protector SQL Server configuration file during configuration of the integration.

The syntax of the configuration file is:

```
Login='user';  
Password='encoded_password';
```



IMPORTANT:

To avoid backup problems, ensure that the syntax of your configuration file matches the examples.

Examples

- **SQL Server authentication:**

```
Login='sa';  
Domain='';  
Password='jsk74yh80fh43kdf';
```

- **Windows authentication:**

```
Login='Administrator';  
Domain='IPR';  
Password='dsjf08m80fh43kdf';
```

- **Integrated authentication:**

```
Login='';  
Domain='';  
Password='kf8u3hdgtfh43kdf';
```

Configuring users

Add the Windows domain user account under which the Data Protector `Inet` service on the SQL Server client is running to the Data Protector `admin` or `operator` group.

For information on adding users to Data Protector groups, see the online Help index: “adding users”.

Configuring SQL Server

Configuration is performed during the creation of the first backup specification and consists of setting the username and password for the Data Protector services. Thereafter, the services are able to connect to SQL Server and operate under the specified account. The specified login information is saved to the Data Protector SQL Server instance configuration file on the Cell Manager.

NOTE:

Ensure that the user account to be used has appropriate SQL Server permissions for running backups and restores. Check the permissions using SQL Server Enterprise Manager.

You can change configuration by following instructions described in “Changing and checking configuration” on page 36.

Prerequisites

- SQL Server must be online during configuration.
- Configuration must be performed for every SQL Server instance separately.

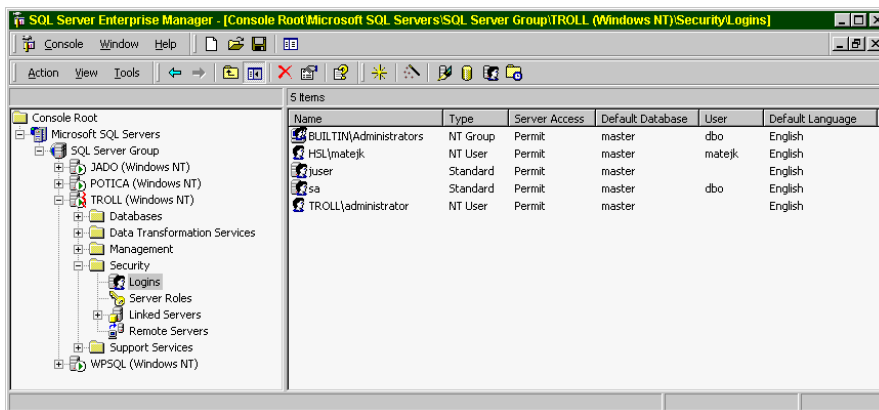


Figure 4 SQL Server users

Configure SQL Server instances using the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup dialog box**, select the **Blank Microsoft SQL Server Backup** template.
Click **OK**.
4. In **Client**, select the SQL Server system. For cluster environments, select the virtual server of the SQL Server resource group.
In **Application database**, specify the name of the SQL Server instance.
Click **Next**.

5. In the Configure MS SQL dialog box, specify the user account that Data Protector should use to connect to the SQL Server instance:
 - **SQL Server authentication:** SQL Server user account. Specify a username and password.
 - **Windows authentication:** Windows domain user account (preferred option). Specify a username, password, and the domain.
 - **Integrated authentication:** Windows domain user account under which the Data Protector `Inet` service on the SQL Server system is running.

See [Figure 5](#) on page 35.

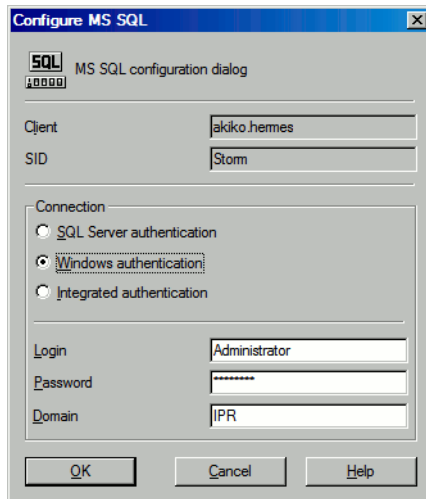


Figure 5 Configuring SQL Server

 **NOTE:**

It is recommended that the SQL Server system administrator configures the integration.

For details about security, see the SQL Server documentation.

Click **OK** to confirm the configuration.

6. The SQL Server instance is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#) on page 40.

Using the Data Protector CLI

From the `Data_Protector_home\bin` directory, run:

```
sql_bar config [-appsrv:SQL_Server_client]
[-instance:instance_name] [-dbuser:SQL_Server_user
-password:password | -dbuser:Windows_user -password:password
-domain:domain]
```

Parameter description

`-appsrv:SQL_Server_client`

Client system on which the SQL Server instance is running. This option is not required if you run the command locally.

`-instance:instance_name`

SQL Server instance name. If you omit this option, the default SQL Server instance is configured.

`-dbuser:SQL_Server_user -password:password`

SQL Server user account

`-dbuser:Windows_user -password:password -domain:domain`

Windows domain user account

NOTE:

If no user account is specified, Data Protector will connect to the SQL Server instance with the Windows domain user account under which the Data Protector `Inet` service on the SQL Server system is running.

Changing and checking configuration

You can check and change configuration using the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS SQL Server**. Click a backup specification for which you want to change the configuration.
3. In the **Source** property page, right-click the SQL Server name and select **Configure**.

4. Configure SQL Server as described in “[Configuring SQL Server](#)” on page 33.
5. Right-click SQL Server and select **Check Configuration**. See [Figure 6](#) on page 37.

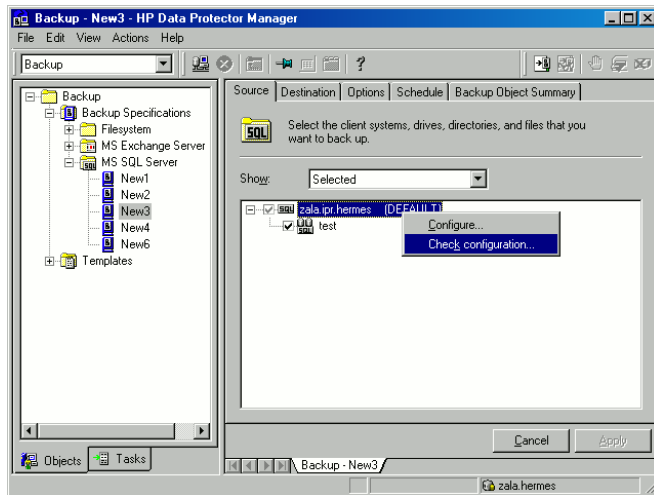


Figure 6 Checking configuration

Using the Data Protector CLI

To change the configuration, run the command for configuring SQL Server instances again, entering different data.

To check configuration, run:

```
sql_bar chkconf [-instance:instance_name]
```

If the optional parameter `-instance:instance_name` is not specified, the default instance is checked.

If the integration is not properly configured, the command returns:

```
*RETVAl*8523
```

To get the information about the existing configuration, run:

```
sql_bar getconf [-instance:instance_name]
```

If `-instance:instance_name` is not specified, Data Protector returns configuration for the default instance.

Backup

To run an online backup of an existing SQL Server backup specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

For information on starting interactive backups using the CLI, see the `omnib` man page.

Considerations

- A transaction log backup is not possible if the `Recovery model` option on SQL Server is *not* set to `Bulk-Logged` or `Full`. In this case, Data Protector performs a differential or full backup.

To configure a backup, create a Data Protector SQL Server backup specification.

Creating backup specifications

Create a backup specification, using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.

3. In the **Create New Backup** dialog box, select the **Blank Microsoft SQL Server Backup** template. See [Figure 7](#) on page 39.

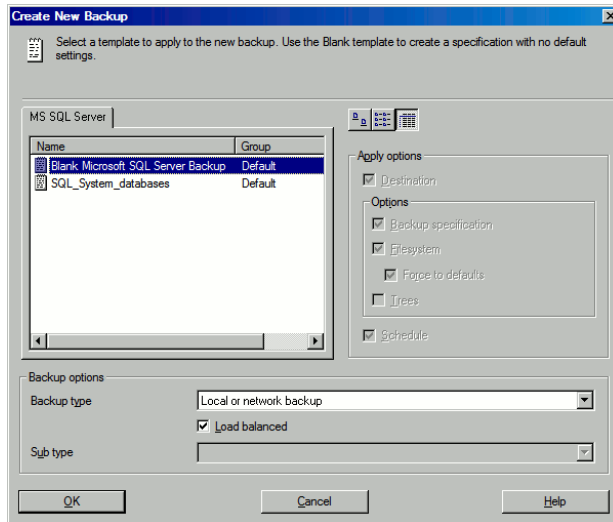


Figure 7 Selecting a blank template

Click **OK**.

4. In **Client**, select SQL Server. For cluster environments, select the virtual server of the SQL Server resource group.

In **Application database**, specify the name of the SQL Server instance.

Click **Next**.

5. If the client is not configured, the **Configure Microsoft SQL** dialog box appears. Configure it as described in [“Configuring SQL Server”](#) on page 33.

6. Select the databases to be backed up.

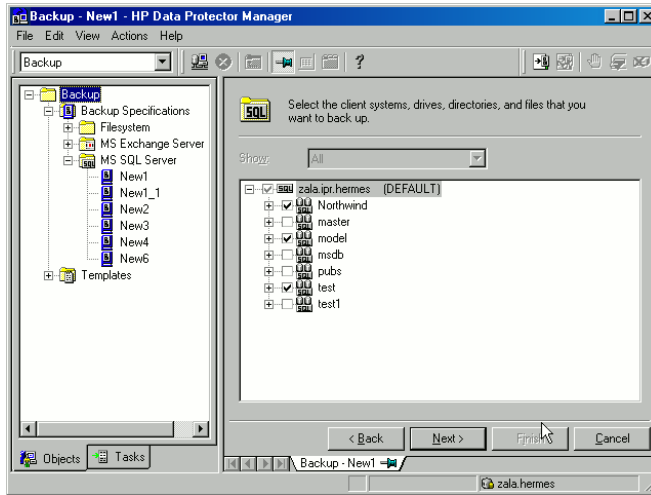


Figure 8 Selecting backup objects

Click **Next**.

7. Select the devices. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on options, press **F1**.

To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror/Remove mirror**. Select separate devices for each mirror. The minimum number of devices for mirroring equals the number of devices used for backup.

For more information on object mirroring, see the online Help.

Click **Next**.

8. Select backup options.

For information on **Backup Specification Options** and **Common Application Options**, see the online Help.

For information on **Application Specific Option**, see [“SQL Server Specific Backup Options”](#) on page 41.

Click **Next**.

9. Optionally, schedule the backup. For information on scheduler, press **F1**.
10. Save the backup specification, specifying a name and backup specification group. You start the backup specification by clicking **Start Backup**.




Figure 9 Saving a backup specification

Table 5 SQL Server backup options

Concurrent streams	Sets the number of concurrent streams (devices) used for backup. Available for a whole server backup.
Fast direct mode	Used with locally connected devices to optimize performance. Must be combined with special device settings (see “Performance tuning” on page 54 for details).
Check database integrity	Performs data integrity validation before backup. If the check fails, the session completes with warnings.
Pre-exec	Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server before backup. Resides in the <code>Data_Protector_home\bin</code> directory. Only the filename must be provided in the backup specification.
Post-exec	Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server after backup. Resides in the <code>Data_Protector_home\bin</code> directory. Only the filename must be provided in the backup specification.



NOTE:

Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

SQL Server Specific Backup Options

SQL Server specific backup options are specified by clicking the **Advanced** tab in the **Application Specific Options** group box.

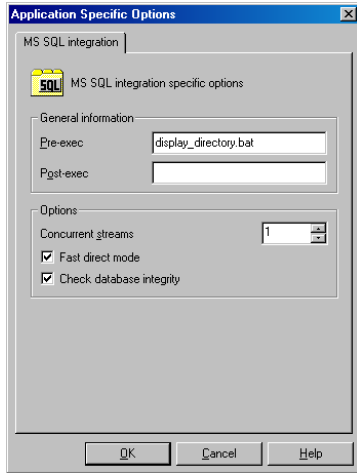


Figure 10 Application specific options

Object specific options

If you selected one or more databases for backup (as opposed to a whole server backup), you can set backup options on a single database level by going to the **Backup Specification Summary** property page and double-clicking an object.



NOTE:

If you selected a whole server backup, the same options as in the **Application Specific Options** windows are displayed.

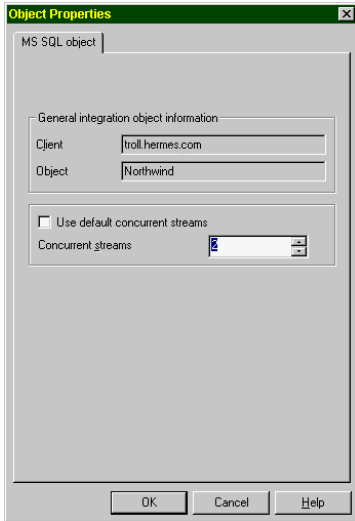


Figure 11 Object properties

Table 6 Object specific options

Use default concurrent streams	The number of concurrent streams is defined by Data Protector and all available devices are used.
Concurrent scripts	Sets the number of concurrent streams (devices). VDI supports up to 32 virtual devices per database.

Scheduling backups

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

Scheduling example

To schedule a database backup at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.

2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**.
Click **OK**.
3. Repeat [Step 1](#) on page 43 and [Step 2](#) on page 44 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

Starting backup sessions

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

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS SQL Server**. Right-click the backup specification you want to start and select **Start Backup**.
3. Select **Backup type** and **Network load**. For information on these options, click **Help**. Click **OK**.

Restore

Data Protector offers different restore types depending on your needs. You can select point-in-time restore, full database restore, as well as restore your SQL Server data to a new location, to another SQL Server or another SQL Server instance. For detailed information, see [“Restore options”](#) on page 48.

Restore methods

You can restore SQL Server databases using:

- The Data Protector GUI. See [“Restoring using the Data Protector GUI”](#) on page 45. TBD
- The Data Protector CLI. See [“Restoring using the Data Protector CLI”](#) on page 51.

Before you begin

- Before starting restore, verify that the database is not being in use.

Restoring using the Data Protector GUI



NOTE:

On SQL Server 2000 and higher, there is no need to create an empty database before restore, because the database and its files are generated automatically.

Proceed as follows using the Data Protector Manager:

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects, MS SQL Server**, and then select the MS SQL Server from which you want to restore. A list of backed up objects is displayed in the Results Area.

3. Select the restore objects. See [Figure 12](#) on page 46.

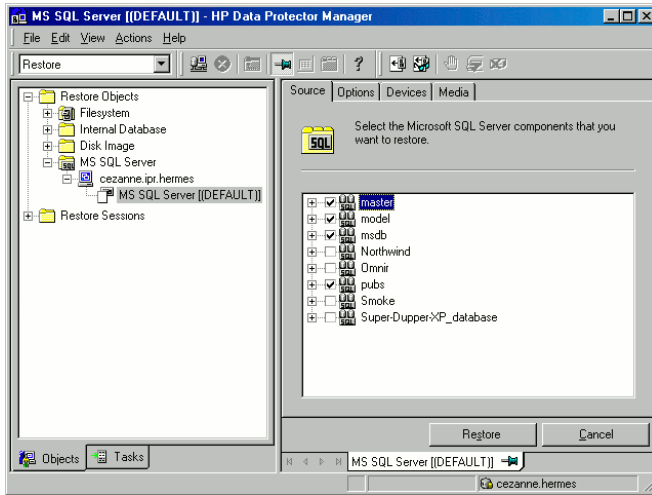


Figure 12 Restore objects

To select backup object specific options, right-click the object and select **Properties**.

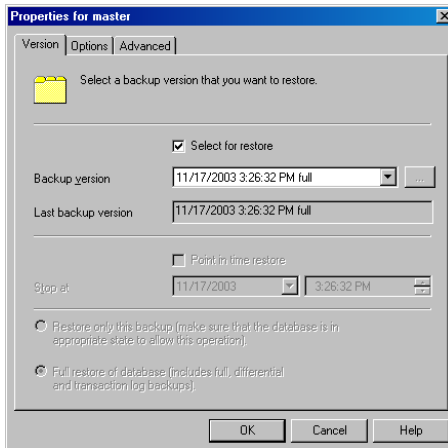


Figure 13 Selecting object specific options

You can select the version (backup date) from which you want to restore and choose SQL Server specific restore options. See [“Restore options”](#) on page 48 for details.

Click **OK**.

4. If you want to restore your data to another client or instance, specify new locations for the databases in the **Options** property page. See “[Restore options](#)” on page 48.

 **NOTE:**

When you click **Options**, the cell is browsed for running SQL Server instances that can become target instances for restore. If no instances are found, **Restore to another instance** is disabled and the message **There are no instances on this client system** is displayed.

Select one of the following **Restore actions**:

- **Restore data** (default). Select to restore the whole database.
- **Restore and display file list only**. Select if you do not know the original filenames. In this case, the files backed up in a particular session are displayed.
- **Restore and display headers only**. Select if you need specific details about backup. SQL Server header information is displayed.

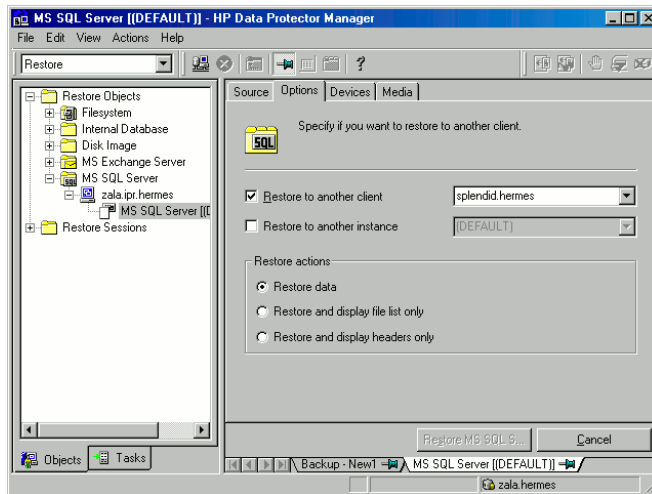


Figure 14 Restore options

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

The **Automatic device selection** option is selected by default, but it is recommended to select the **Original device selection** option.

 **IMPORTANT:**

If you decide to select the **Automatic device selection** option, ensure that the number of available devices is equal to or greater than the number of devices that were used for backup.

 **NOTE:**

For restore, you can use a different device than you used for backup. For information on restoring using another device, see the online Help index: “selecting, devices for restore”.

6. Click **Restore MS SQL Server** and then **Next** to select **Report level** and **Network load**.

Click **Finish** to start restore.

Restore options

Table 7 SQL Server restore options

Option	Description
Backup version	Specifies the backup session from which the selected objects will be restored.
Point-in-time restore	Specifies a point in time to which the database state will be restored (you also need to select Backup version and set Stop at). After recovery, the database is in the state it was at the specified date and time. Only transaction logs written before the specified date and time are applied to the database.

Option	Description
Stop at	<p>Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable the database recovery to a particular point in time, backup you restore from must include transaction log backups.</p> <p>You cannot use this option used with NORECOVERY or STANDBY. If you specify Stop at time that is after the end of RESTORE LOG operation, the database is left in a non-recovered state (as if RESTORE LOG is run with NORECOVERY).</p>
Restore only this backup	<p>If you restored a database version and left it in a non-operational or standby state, you can subsequently restore differential or transaction log backups one by one, leaving each version non-operational to restore additional backups.</p>
Full restore of the database	<p>All necessary versions are restored, including the latest full backup, the latest differential backup (if one exists), and all transaction log backups from the last differential up to the selected version.</p>
Force restore over the existing database	<p>Overwrites the existing database residing on the target server. If a database with the same name as the one you are restoring already exists and has a different internal structure, SQL Server does not let you rewrite the database unless you select this option.</p>
Recovery completion state	<p>Enables selecting the database state after recovery. You may select from:</p> <ul style="list-style-type: none"> • Leaving the database operational. Once the last transaction log is restored and the recovery completed, the database becomes operational. • Leaving the database non-operational after the last transaction log is restored. You may restore additional transaction logs one by one. • Leaving the database in read-only mode. You may restore additional transaction logs before the database is set to read-write mode.
Restore database with a new name	<p>Restores the database under a different name. Specify the database logical filename and the destination filename (suboptions of Restore files to new locations).</p>

Option	Description
Restore files to new locations	Restores files to a new location. Specify the database logical filename and a destination target filename for the specified logical filename. Use this option if you restore data to another server, instance, or make a database copy on the same server.

 **TIP:**

To allow different restore scenarios, you can combine general restore options, such as **Restore database to another Microsoft SQL Server** and **Restore using a different device**, with object-specific restore options, such as **Point-in-time restore**, **Recovery completion state**, **Force restore over existing database**.

Restoring to another SQL Server instance or/and another SQL Server

Prerequisites

- Both SQL Servers must have the same local settings (code page and sort order). This information is displayed in the session monitor for each backup.
- The target SQL Server must be configured and reside in the same Data Protector cell as the original SQL Server. For configuration procedure, see [“Creating backup specifications”](#) on page 38.

 **NOTE:**

This restore type is supported on SQL Server 2000 and higher.

1. Select the databases you want to restore and their versions.
2. Select the following:
 - To restore to another SQL Server, select **Restore to another client** and the target client from the drop-down list.
 - To restore to another SQL Server instance, select **Restore to another instance**. If there are no instances in the drop-down list, enter the instance name yourself.
 - To restore to another SQL Server and another SQL Server instance, enter the name of the instance that exists on the target client. Otherwise, restore fails.

3. Specify new database locations.
4. Start restore. See “Restore” on page 44.

Restoring using the Data Protector CLI

From the `Data_Protector_home\bin` directory, run:

```
omnir -MSSQL -barhost MSSQL_Server_Name [-destination  
Target_MSSQL_Server]-base dbname[-session Session  
_ID][-nochain][-replace][-recovery rec|norec][-standby File]  
[-instance instance name]
```

Provide the `Session_ID` of the backup session. For object copies, do not use the copy session ID, but the object backup ID (equals the object backup session ID).

Example

To restore the database RONA running on the SQL Server Alma to the same destination, run:

```
omnir -MSSQL -barhost Alma -base RONA
```

Disaster recovery

Disaster recovery is a complex process involving products from different vendors. Therefore, you need to check the instructions from the database/application vendor on how to prepare for disaster recovery.

As a first step, perform a general disaster recovery procedure described in the *HP Data Protector disaster recovery guide*. Next, restore SQL Server databases. See the below sections for instructions.



IMPORTANT:

If a disk failure occurred, recover the operating system prior to any other recovery tasks. Data Protector disaster recovery is used to bring the operating system back on the damaged system.

 **IMPORTANT:**

When reinstalling SQL Server, ensure that you use original local settings. Before restoring to another client, ensure that local settings on the target system match the original.

Recovering the master database

The master database holds the vital information about SQL Server. If it gets corrupted or lost, all other databases become unavailable. Recover the master database first to make SQL Server operational:

1. Rebuild the master database.

Create the basic master database:

- a.** Shut down SQL Server if it is running.
- b.** Start the Rebuild Master utility `SQL\bin\rebuiladm.exe`.
- c.** Select an appropriate character set and sort order to match the backed up data. You can check this in the latest backup session report.
- d.** Rebuild the database.

For more information, see the SQL Server documentation.

2. Set user rights or reconfigure the integration.

Set user rights using SQL Server Enterprise Manager:

- a.** From the server desktop, click **Start - Programs - Microsoft SQL Server 7/Microsoft SQL Server - Enterprise Manager**.
- b.** Right-click the required server and select **Register Server**. Configure SQL Server to use trusted connections.
- c.** Go to **Security - Logins** and select appropriate user rights.
- d.** Return to the server, right-click its name, and select **Register Server**.
Enter the account you selected in **Manage - Logins**.
Perform any additional administration tasks required to run SQL Server.

Reconfigure the SQL Server integration as described in [“Creating backup specifications”](#) on page 38.

3. Stop SQL Server services by going to **Start - Programs - Microsoft SQL Server 7**, starting SQL Service Manager and stopping the services.
4. Start SQL Server service in a single-user mode:
 - a. In the **Control Panel**, go to **Administrative Tools, Services**.
 - b. Select the MSSQL Server Service.
 - c. Enter `-m` as a start-up parameter and start the services.
5. Restore the master database using the Data Protector Manager.
6. Restart SQL Server services in normal mode.

After the master database is recovered, the SQL Server service is automatically shut down. Start SQL Server Service Manager and restart SQL services.



IMPORTANT:

To complete disaster recovery, restore *all* other databases as well (or reattach databases if they exist on disks to the newly-rebuilt master database). See [“Recovering user databases”](#) on page 53.

Recovering user databases

To restore user databases, proceed as described in [“Restore”](#) on page 44.

Note that restoring databases to a certain state often requires a multiphase restore. This means that multiple versions need to be restored to retrieve the data. The latest full backup, the latest differential backup and all transaction log backups after the last full or differential backup must be restored.

Example

Suppose you have the following backup sequence:

F D T T D T T T T T

and want to restore the version marked *T*, then all the backup versions in *italic* will be restored.

**TIP:**

You cannot restore versions one by one to have more control over the restore process. Use the options **Restore only this backup** and **Recovery completion state** to do this.

For more information on disaster recovery, see the *HP Data Protector disaster recovery guide* and the SQL Server documentation.

Performance tuning

Performance tuning means customizing your environment to improve backup and restore performance. Follow these guidelines:

1. Ensure that SQL Server database files are on separate disks.
2. Calculate the number of devices to be used in parallel. Select a number of devices matching the bandwidth of the incoming data stream and identify the bottleneck. This can either be the network, if devices are connected to remote systems, or SQL Server, if the devices are connected locally.

As the network bandwidths are most often ~10 MB/s (100 Mbit Ethernet), though the actual throughput is usually lower, you will not need more than one fast device (such as DLT 7000 for remote backups).

There are two possibilities for locally connected devices:

- a. Devices are dedicated to local SQL Server backups and backup/restore performance is important. Use fast direct mode, which enables Data Protector to read data directly from the SQL Server shared memory and can therefore increase the backup speed to local devices.
- b. Devices are shared within the Data Protector cell and backup/restore performance is not very important. Disable fast direct mode.

Determine the maximum backup speed by backing up to a few null file devices on a local server, and select the number of devices that fits best with the measured performance.

**TIP:**

Create separate backup specifications for local and remote devices. It is not recommended to use both in one backup specification.

3. Adjust block sizes for local backup devices.

- Enable/disable **Fast direct mode**.

Use this option only if the highest performance is required. Due to specific device settings, these device definitions should not be shared with conventional (filesystem) backups. Therefore, using this option in general is not recommended.

Disable **Fast direct mode** (as well as special local device settings) if backup performance is not very critical and/or other data is backed up to devices connected to SQL Server.

 **NOTE:**

Fast direct mode is ignored for remote devices.

- Set the block size (if **Fast direct mode** is enabled).

Adjusted block sizes are calculated as follows:

`block size (kB) = 64*N + 4 (N=1,...64)`

`block size (kB) = 68, 132,4100 kB`

All selected devices must have the same block size.

You can gain some performance improvement by specifying a block size larger than 68 KB (recommended). You can also increase the block size step by step and compare the performance achieved for each step.

You can adjust block size during the initial device definition for local devices by checking the attached check box and selecting the block size. See [Figure 15](#) on page 56.

You can modify block size later; however, you must first calculate it using the formula above and then insert the value as shown in [Figure 16](#) on page 57.

- Modify the registry.

To use block size larger than 56 KB, some SCSI interface cards require you to adjust related values in the registry of the system where the device is connected. See the online Help index: "changing block size" for instructions.

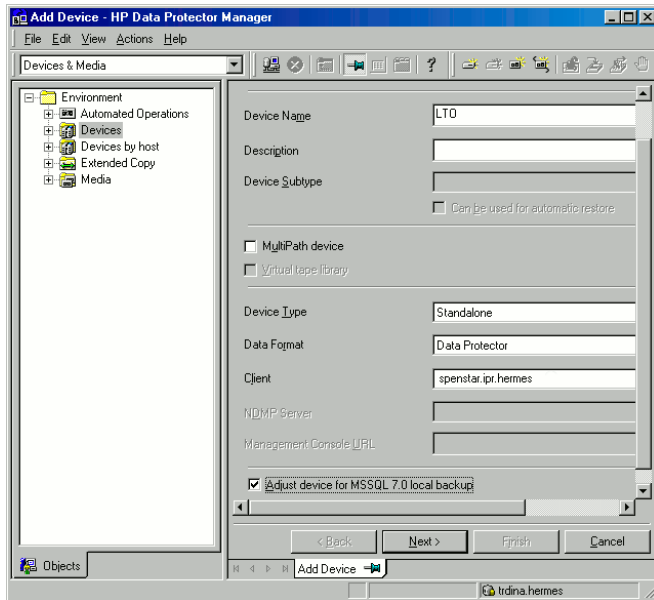


Figure 15 Adjusted local device

To modify block sizes of an existing device:

- a.** Switch to the **Devices & Media** context.

In the Scoping Pane, expand **Devices** and click the locally connected device you want to modify. In the Results Area, select **Settings**, and then click **Advanced**.

- b.** In the **Advanced Options** window, click **Sizes**.

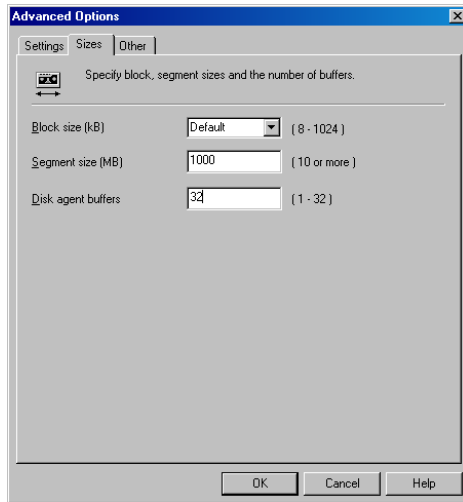


Figure 16 Advanced options

If **Fast direct mode** is activated and not all selected local devices in a backup specification are adjusted accordingly, you get the warning when saving the backup specification:




Figure 17 Block sSizes not adjusted

4. Scheduling.

Backup schedule depends on the number of transactions on the server. Generally, you should not let transaction log files grow over a certain limit, which depends on a specific production database and the size of its transaction log files. These are some general rules on how to schedule backups:

- Weekly full backup
- Differential backup daily
- Transaction log backups as needed

Schedule full and differential backups when the server is not heavily loaded (nights and weekends). Do transaction log backups several times a day.

The final decision on scheduling must be made according to the actual database configuration.

For more information, see the SQL Server documentation and the online Help.

Monitoring sessions

You can monitor currently running or view previous sessions in the Data Protector GUI. When you run an interactive session, the monitor window shows you the session progress. Closing the GUI does not affect the session.

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

For information on monitoring sessions, see the online Help index: “viewing currently running sessions” and “viewing finished sessions”.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector SQL Server integration. Start at “[Problems](#)” on page 60. If you cannot find a solution there, perform general checks and verifications.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Check that SQL Server services are running.
- Examine system errors reported in `Data_Protector_home\log\debug.log` on SQL Server functioning as a Data Protector client.
Additionally, check `errorlog` and `VDI.log` files in the `MSSQL\log` directory.
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Check that every SQL Server used with Data Protector has the `MS SQL Integration` component installed.
- Connect to SQL Server via SQL Server Enterprise Manager using the same login ID as you specified in the Data Protector **Configuration** dialog box.
- Perform a database backup using SQL Server Enterprise Manager. If the backup fails, fix any SQL Server problems, and then perform a backup using Data Protector.

Additionally, if your backup failed:

- Verify the configuration file to check if the Cell Manager is correctly set on SQL Server.
- If you do not see the SQL Server instance as the application database when creating a backup specification, enter the instance name yourself. When “not-named instance” is not displayed, insert the `DEFAULT` string.
- If Data Protector reports that the integration is properly configured, verify that the SQL Server user has appropriate rights to access the required databases.

During restore, if the following error occurred when executing an SQL statement:

```
Error message: "Microsoft SQL-DMO (ODBC SQLState: 01000)?15[152:5]
1646 [Microsoft][ODBC SQL Server Driver][SQL Server]The master database
has been successfully restored. Shutting down SQL Server.[Microsoft]
[ODBC SQL Server Driver] [SQL Server]SQL Server is terminating this
process."
```

note that this is the expected behavior when the master database is restored in a single user mode.

Problems

Problem

The integration is properly configured but the database backup fails after a timeout

- With an error similar to:

```
[Warning] From: OB2BAR@computer1.com "MSSQL70"
Time: 3/14/2000 8:19:22 PM
Error has occurred while executing SQL statement.
Error message: 'Microsoft SQL-DMO (ODBC SQLState: 42000) Error
number: bc5
[Microsoft][ODBC SQL Server Driver][SQL Server]Backup or restore
operation terminating abnormally.'
[Critical] From: OB2BAR@computer1.com "MSSQL70"
Time: 3/14/00 8:19:24 PM
Received ABORT request from SM => aborting
```

- SQL Server error log contains an entry similar to:

```
2000-03-14 20:19:21.62 kernel
BackupVirtualDeviceSet::Initialize: Open failure on backup
device 'Data_Protector_master'.
Operating system error -2147024891(Access is denied.).
```

- SQL Server **VDI.LOG** file contains an entry similar to:

```
2000/03/15 13:19:31 pid(2112)
Error at BuildSecurityAttributes: SetSecurityDescriptorDacl
Status Code: 1338, x53A Explanation: The security descriptor structure
is invalid.
```

SQL Server service and Data Protector Inet are running under different accounts. The integration cannot access SQL Server due to security problems.

Action

Restart the `Data Protector Inet` service under the same account as the SQL Server service is running.

Problem

Backup fails if concurrency is set to more than one and one of the devices fails or is not started at all

This can happen because of a medium error.

Action

Set the device concurrency to one or replace the invalid media.

Problem

Restore from an object copy fails

Action

Before restarting the restore:

- Increase the number of Disk Agent buffers for the device.
- If all backup objects are recorded in the IDB:
 1. In the **Internal Database** context, find the objects belonging to the same backup (identified by the same backup ID).
 2. Copy each object in a separate object copy session to a separate device, for example a file library. For each object, use a separate medium with the non-appendable media policy.
 3. Set the highest media location priority for the newly created copies.

Problem

Database is left in unrecovered state after “Invalid value specified for STOPAT parameter” is reported

The database remains in an unrecovered state as if the `RESTORE LOG` operation was run with `Leave the database non-operational`.

Action

Recover the database to the latest point in time using SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

Problem

Restore to another client in the Data Protector cell not configured for use with SQL Server fails

Action

Configure the SQL integration on this client (see “[Configuring the integration](#)” on page 30).

Problem

Database is left in unrecovered state after restore completed successfully

If you set the time for `stop` at beyond the end of the `RESTORE LOG` operation, the database remains in the unrecovered state as if the `RESTORE LOG` operation was run with `leave` the database non-operational.

Action

Recover the database to the latest point in time by using the SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

Problem

Restoring a Microsoft SQL Server 2005 instance to an alternate location when full-text indexing is enabled fails

When the Use full-text indexing option is enabled for a particular database in a Microsoft SQL Server 2005 instance, the restore session does not complete successfully, since restore of the full-text catalog of the SQL database fails. The session report contains warning messages about the full-text catalog file being used by the affected database.

Action

To solve the problem:

1. In the HP Data Protector Manager, switch to the **Restore** context.

2. In the Scoping Pane, expand **Restore Objects** and then **MS SQL Server**. Select name of the Microsoft SQL Server for which you want to perform restore.
3. In the Results Area, double-click the bar name corresponding to the particular Microsoft SQL Server instance. A list of backed up objects gets displayed.
4. Select the desired Microsoft SQL Server database, right-click it, and click **Properties**.
5. In the Properties window, click the **Advanced** tab.
6. Select the **Restore database with new name** option, and enter the new database name in the text box.
7. For all logical file names that are already present on the list, update contents of the Destination file name column accordingly.
8. Add the full-text catalog to the list.

In the Logical file name text box, enter the string `sysft_Full-Text_Catalog_Name`. In the Destination file name text box, enter the corresponding physical location.

 **NOTE:**

The full-text catalog is always restored to its original location, regardless of the specified physical location.

9. Click **Add/Set**.
10. In the Version and Options property pages, specify the appropriate options. For details, see [“Restoring using the Data Protector GUI ”](#) on page 45.
11. Click **OK** to close the Properties window.
12. In the Options, Devices, and Media property pages, specify the appropriate options. For details, see [“Restoring using the Data Protector GUI ”](#) on page 45.
13. Click **Restore** and then **Next** to select the Report level and Network load.
14. Click **Finish** to start the restore session.

2 Integrating Microsoft SharePoint Portal Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Microsoft SharePoint Portal Server integration (SPS integration). It describes concepts and methods you need to understand to back up and restore the following SharePoint Portal Server objects (SPS objects):

- Content databases (team databases)
- Site databases (*portal_name_SITE*, *portal_name_SERV*, *portal_name_PROF*)
- Index servers
- Single sign-on database
- Document library

Data Protector integrates with Microsoft SharePoint Portal Server (SPS Server) to back up SPS objects online. During backup, the SPS Server and Microsoft SQL Server instances are online and actively used.

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

Table 8 Backup types

Full	Backs up all selected objects.
Trans (MS SQL Server objects only)	Backs up only transaction logs of the selected SQL Server databases. Other selected SPS objects are backed up completely.

Differential (MS SQL Server objects only)	Backs up only changes made to the selected SQL Server databases since the last full backup. Other selected SPS objects are backed up completely. Before you run a differential backup, ensure that a full backup exists. Otherwise, a restore from such a differential backup session fails.
---	---

You can restore SQL Server databases to the original location or:

- To another SQL Server system
- To another SQL Server instance
- Under different names

You can restore SPS index servers to the original location or:

- To another client
- To another directory

This chapter provides information specific to the SPS integration. For general Data Protector procedures and options, see online Help.

Integration concepts

Data Protector integrates with SPS Server through the Data Protector integration agents (SPS and SQL agents), which channel communication between the Data Protector Session Manager and the clients in the SPS farm.

Whether your SPS environment consists of a single server system or multiple server systems (small, medium, or large farm), the architecture of the integration is basically the same.

[Figure 18](#) on page 67 shows the architecture of the Data Protector integration with a SPS medium server farm.

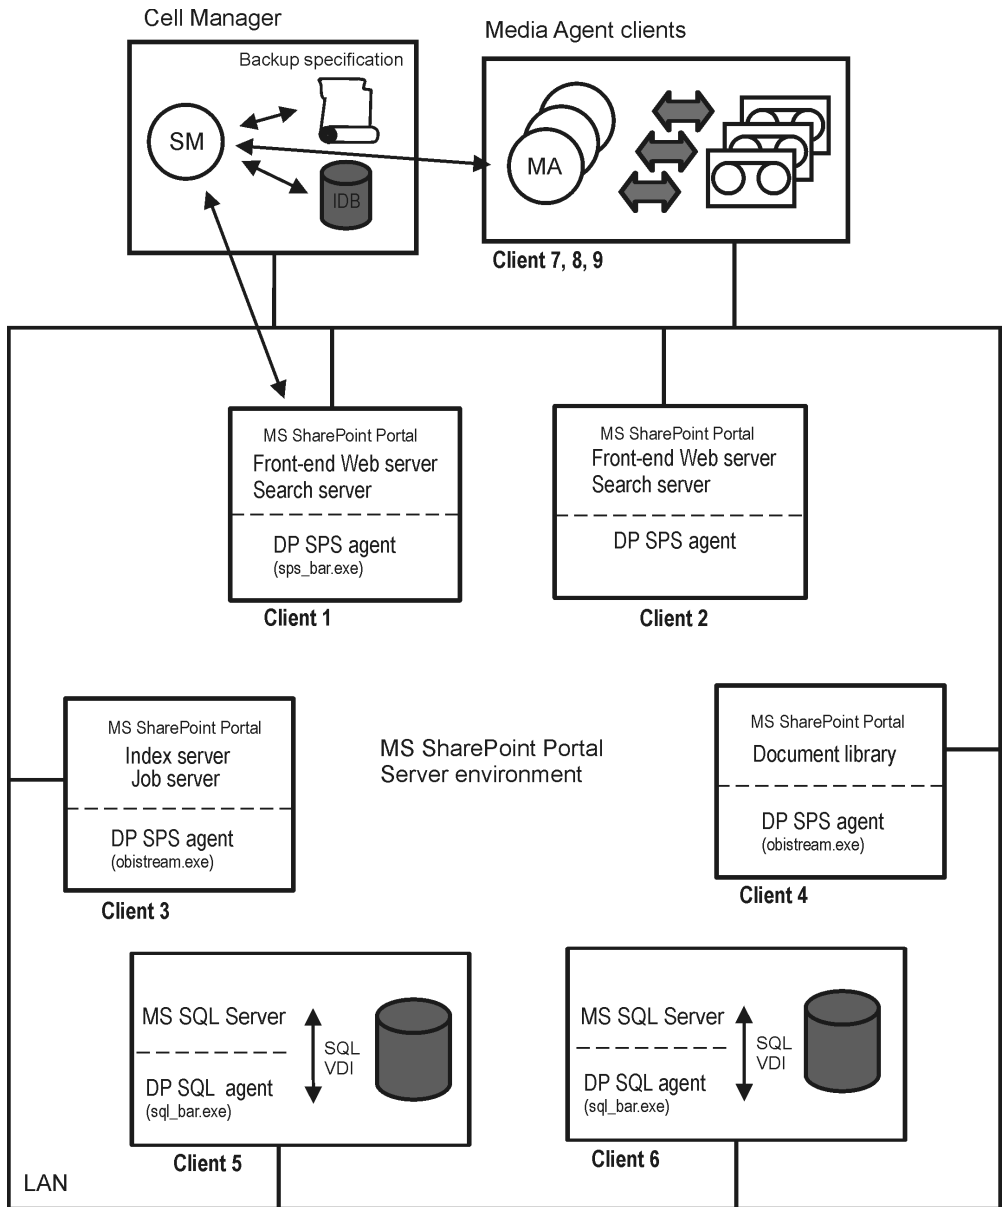


Figure 18 SPS integration architecture

Table 9 Legend

SM	Data Protector Session Manager: Backup Session Manager during
----	---

	backup and Restore Session Manager during restore
Backup Specification	A list of objects to be backed up, backup devices, and options to be used
IDB	The Data Protector Internal Database
DP SPS agent	A set of Data Protector executables that enable data transfer between SPS Server and Data Protector media
DP SQL agent	A set of Data Protector executables that enable data transfer between SQL Server and Data Protector media
SQL VDI	SQL Server virtual device interface, through which SQL Server and Data Protector exchange control and data
MA	Data Protector General Media Agent
LAN	Local Area Network

Table 10 briefly describes the SPS objects that you can back up and restore using the Data Protector SPS integration.

Table 10 SPS objects

SPS object	Description
Content database	Virtual server specific SQL Server database that stores Web site data.
Site databases	Portal specific SQL Server databases: <ul style="list-style-type: none"> • <i>portal_name_SITE</i> Contains user profiles and audiences. • <i>portal_name_SERV</i> Stores information for services, such as search and alerts, provided by the portal. • <i>portal_name_SITE</i> Contains site content information.
Index server	Data used by index management server, which builds and updates indexes, and crawls content.

SPS object	Description
Single sign-on database	SQL Server database that stores account credentials. The single sign-on functionality enables users to retrieve information from third-party applications without additional sign-on operations.
Document library	Folder that stores files. Each file is associated with user-defined information.

Backup and restore flow

The following procedure assumes that all SPS objects are selected for backup/restore.

- 1.** Data Protector Session Manager starts `sps_bar.exe` on the front-end Web server client, providing a list of objects to be backed up/restored.
- 2.** The `sps_bar.exe` agent browses the SharePoint Portal Server topology.
- 3.** For each portal:
 - a.** The `sps_bar.exe` agent starts `sql_bar.exe` on the SQL Server clients on which the portal site and content databases reside.

The `sql_bar.exe` agent establishes connection with the Session Manager and sends backup/restore requests to the SQL Server using the SQL Server VDI. The actual data transfer is then performed by the SQL Server. The databases are backed up/restored in parallel.
 - b.** If the portal provides index service, `sps_bar.exe` also starts `obistream.exe` on the client functioning as index server. `obistream.exe` establishes connection with the Session Manager to enable backup/restore data flow for the index server database.
- 4.** The `sps_bar.exe` agent starts `obistream.exe` on the document library client. `obistream.exe` establishes connection with the Session Manager to enable backup/restore data flow for the document library database.
- 5.** The `sps_bar.exe` agent starts `sql_bar.exe` on the SQL Server client on which the single sign-on database resides, which is then backed up/restored like all the other SQL Server databases.

Configuring the integration

You need to configure the SPS farm and SPS users.

Prerequisites

- Ensure that you have correctly installed and configured SPS Server.
 - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or <http://www.hp.com/support/manuals>.
 - For information on installing, configuring, and using SPS Server, see the Microsoft SharePoint Portal Server documentation.
- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HP Data Protector installation and licensing guide*.

The following Data Protector components must be installed:

- MS SharePoint Portal Server Integration - on SPS Server systems
- MS SQL Integration - on SQL Server systems

Before you begin

- Configure devices and media for use with Data Protector.
- On every client in the SPS farm, restart the Data Protector `Inet` service under a Windows domain user account that has Windows administrative rights on the client. For information, see the online Help index: "changing Data Protector Inet account".
- To test whether the SPS farm and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every client system in the farm.

Configuring SPS users

For every client in the SPS farm, add the Windows domain user account under which the Data Protector `Inet` service is running to the Data Protector `admin` or `operator` user group. For details on adding users to Data Protector groups, see the online Help index: "adding users".

Configuring SPS farms

You need to provide Data Protector with the following configuration parameters for the SPS farm:

- SPS administrator

- Password
- Domain

Data Protector then creates the SPS configuration file on the Cell Manager and verifies the connection to the farm.

Before You Begin

- Ensure that the SPS Server and SQL Server instances are online.

To configure the SPS Server, use the Data Protector Manager:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SharePoint Server**, and click **Add Backup**.
3. In the **Create New Backup dialog box**, click **OK**.

4. In **Client**, select the front-end Web server system of the SPS farm.

 **NOTE:**

In a farm with multiple front-end Web server systems, select one of them.

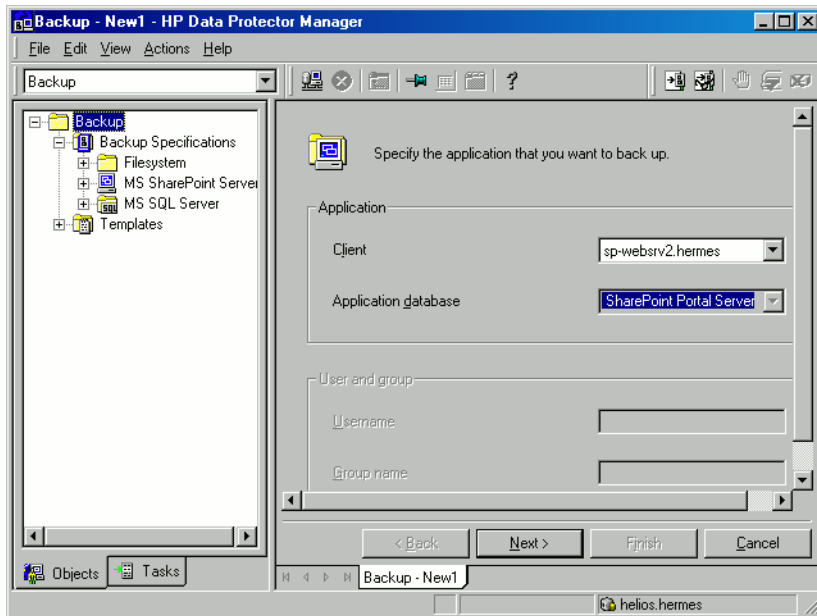


Figure 19 Selecting a front-end Web server system

Click **Next**.

5. In the **MS SharePoint Configuration** dialog box, type the SPS administrator, its password, and domain.

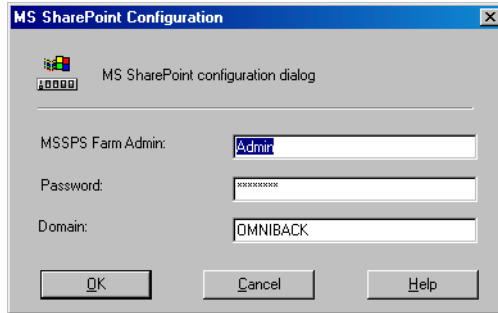


Figure 20 Configuring an SPS farm

Click **OK**.

6. The SPS Server farm is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#) on page 122.

Checking the configuration

You can verify the connection to the SPS farm after you have created at least one backup specification for the farm.

To verify the connection, use the Data Protector Manager:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS SharePoint Server**. Click the backup specification to display the SPS farm to be checked.
3. Right-click the front-end Web server system of the SPS farm and click **Check configuration**.

Backup

You can back up the following SPS objects:

- Content databases (team databases)
- Site databases (*portal_name_SITE*, *portal_name_SERV*, *portal_name_PROF*)
- Index servers
- Single sign-on database

- Document library

The integration provides online backups of the following types:

Table 11 Backup types

Full	Backs up all selected objects.
Trans (MS SQL Server objects only)	Backs up only transaction logs of the selected SQL Server databases ¹ . Other selected SPS objects are backed up completely.
Differential (MS SQL Server objects only)	Backs up only changes made to the selected SQL Server databases ¹ since the last full backup. Other selected SPS objects are backed up completely. Before you run a differential backup, ensure that a full backup exists. Otherwise, a restore from such a differential backup session fails.

¹SQL Server databases are content databases, site databases, and the single sign-on database.

For details on SQL Server transaction log backups and differential database backups, see the Microsoft SQL Server documentation.

 **NOTE:**

If the encryption key is used for the single sign-on service, back up the encryption key as described in Microsoft SharePoint Portal Server documentation.

Before you begin

- Ensure that the SPS services run under the SPS administrator account.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SharePoint Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, click **OK**.

4. In **Client**, select the front-end Web server system of the SPS farm.

 **NOTE:**

In a farm with multiple front-end Web server systems, select the one that was specified during the configuration.

Click **Next**.

5. If the SPS farm is not configured for use with Data Protector, the **MS SharePoint Configuration** dialog box is displayed. Configure it as described in “[Configuring SPS farms](#)” on page 70.
6. Select SPS objects to be backed up.

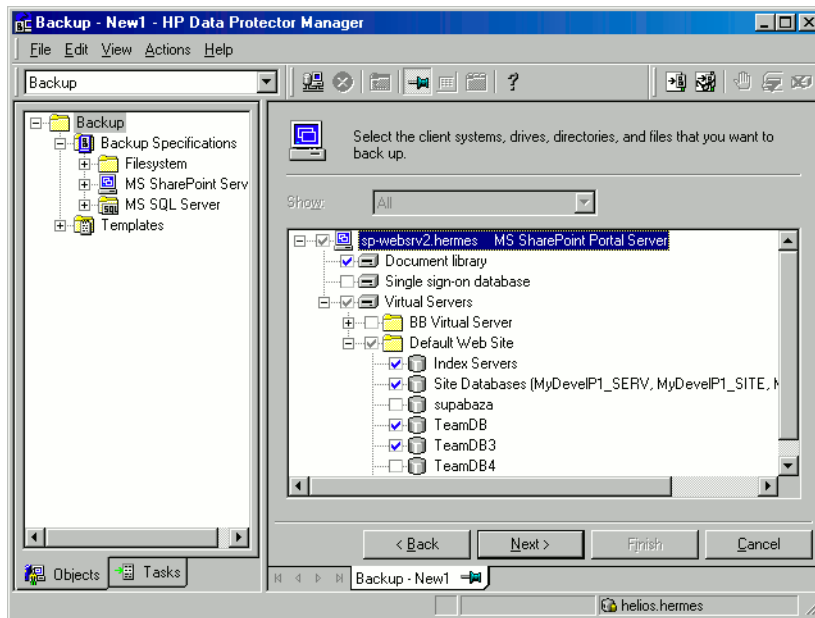


Figure 21 Selecting SPS objects

Click **Next**.

7. Select devices to use for the backup.

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

 **NOTE:**

Every site database, service database, profile database, content database, portal metadata, document library, and index server can be backed up in a separate stream.

Click **Next**.

8. Set backup options. For information on application specific backup options, see [Table 12](#) on page 77.

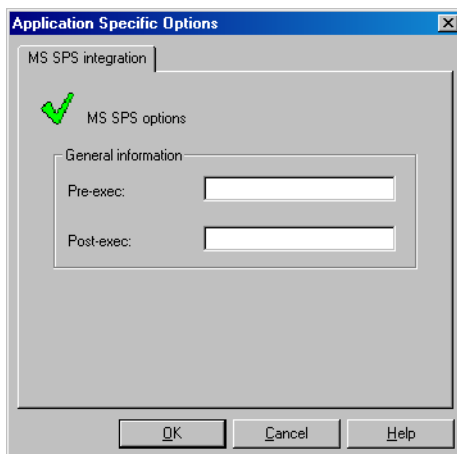


Figure 22 Application specific options

Click **Next**.

9. Optionally, schedule the backup. See [“Scheduling backup specifications”](#) on page 77.

Click **Next**.

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

 **TIP:**

Preview your backup specification before using it for real. See “Previewing backup sessions” on page 78.

Table 12 SPS backup options

Options	Description
Pre-exec, Post-exec	Specify a command to be run by <code>sps_bar.exe</code> on the front-end Web server system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes. Type only the name of the command and ensure that the command resides in the <code>Data_Protector_home\bin</code> directory on the front-end Web server system.

Modifying backup specifications

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

Scheduling backup specifications

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

Scheduling example

To schedule transaction log backups of selected SPS objects at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**. See [Figure 23](#) on page 78. Under **Session options**, select **Trans (for SQL objects only)** from the **Backup type** drop-down list.

Click **OK**.

3. Repeat [Step 1](#) on page 77 and [Step 2](#) on page 77 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

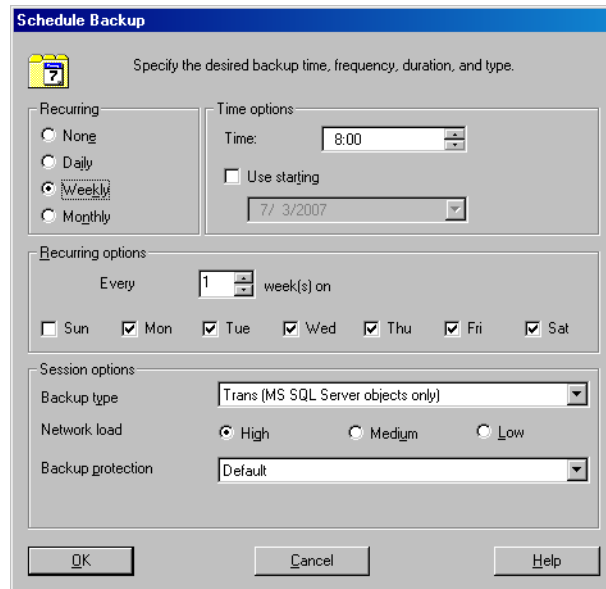


Figure 23 Scheduling a backup specification

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the directory `Data_Protector_home\bin`, run:

```
omnib -mssps_list backup_specification_name -test_bar
```

What happens during the preview?

The following is tested:

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

Starting backup sessions

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

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

Before you begin

- Ensure that the SPS Server and SQL Server instances are online.

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the directory `Data_Protector_home\bin`, run:

```
omnib -mssps_list backup_specification_name [-barmode
SPS_mode] [List_options]
```

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

```
full|trans|diff
```

For *List_options*, see the `omnib` command in the HP Data Protector command line interface reference.

Examples

To start a full backup using the backup specification `MyWebPortals`, run:

```
omnib -mssps_list MyWebPortals -barmode full
```

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

```
omnib -mssps_list MyWebPortals -barmode diff
```

Preparing for disaster recovery

To prepare for a disaster recovery, back up the following objects:

Table 13 What must be backed up

Objects	How to back up
IIS database (from all the front-end Web server clients)	Use Data Protector filesystem backup. For details, see online Help. IIS database of a particular client is located in the client CONFIGURATION.
Encryption key (if it is used for the single sign-on service)	See the Microsoft SharePoint Portal Server documentation.
Master database (of all the SQL Server instances)	Use the Data Protector SQL Server backup. For details, see Chapter 1 on page 27.
SPS objects	Use the Data Protector SPS backup.

Restore

You can restore SPS objects using the Data Protector GUI or CLI.

Before you begin

- Ensure that the SPS Server and the SQL Server instances are online and that the SPS services run under the SPS administrator account.
- If you plan to restore SQL Server databases to another location:
 - Ensure that the destination SQL Server system is part of the SharePoint Portal Server environment and has the `MS SQL Integration` component installed.
 - Ensure that the destination SQL Server instance exists, is configured for use with Data Protector, and is online.
- The following is applicable only if encryption key is used for the single sign-on service: before restoring the single sign-on database, ensure that you use the same encryption key as was used during backup.

Considerations

- If you have two or more SPS farms within the same Data Protector cell, you cannot restore/migrate SPS objects from one farm to another.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS SharePoint Server**, expand the front-end Web server system of the SPS farm to be restored, and then click **MS SharePoint Server [MS SharePoint Portal Server]**.

3. In the **Source** page, select SPS objects to be restored.

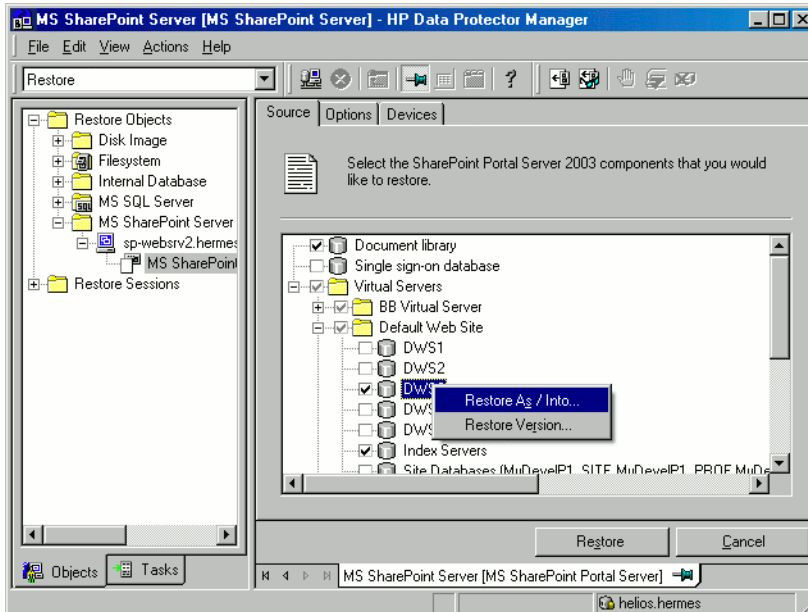


Figure 24 Selecting SPS objects for restore

You can specify a backup version to restore from for each SPS object separately: right-click an object and click **Restore Version**.

For an SQL Server database, Data Protector automatically restores the full chain. For example, if you select:

- | | |
|---------------------|---|
| Full backup | Data Protector restores the selected backup session. |
| Differential backup | Data Protector first restores the latest Full backup and then the selected Differential backup. |
| Trans backup | Data Protector first restores the latest Full backup, then the latest Differential backup (if one exists) and all the Trans backups from the latest Differential or Full backup up to the selected version. |

You can specify the restore destinations for SQL Server databases and index servers: right-click a database/index server and click **Restore As/Into**.

You can restore an SQL Server database to a different SQL Server system, SQL Server instance, or under a different name. See [Figure 25](#) on page 83.

 **NOTE:**

To restore to a not-named (default) SQL Server instance, type (DEFAULT) in the **Instance name** text box.

 **NOTE:**

Site databases (*portal_name_PROF*, *portal_name_SERV*, *portal_name_SITE*) cannot be restored under a different name.

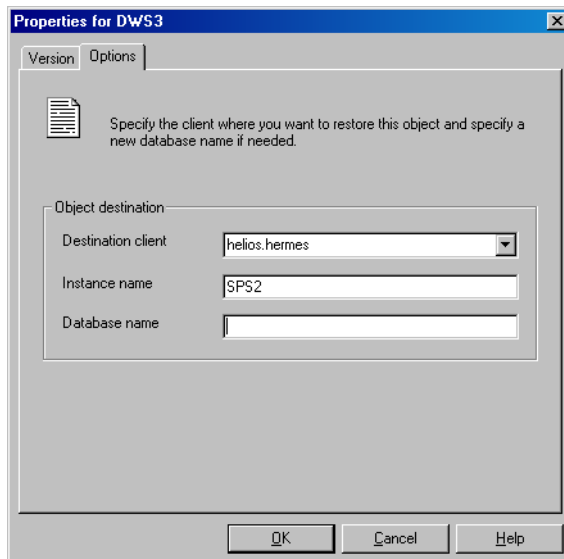


Figure 25 Specifying the restore destination for an SQL Server database

You can restore an index server to a different client or directory. See [Figure 26](#) on page 84.

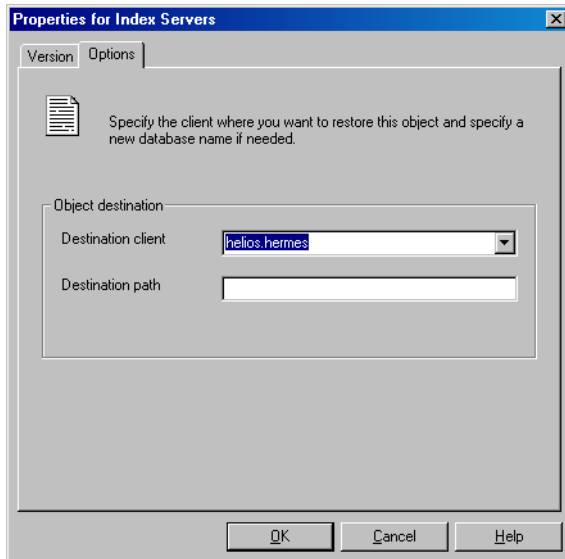


Figure 26 Specifying the restore destination for an index server

4. In the **Options** page, select the SPS specific restore options. These options are applicable only to farms that are centralized (having the master portal and child portals). Otherwise, the options are ignored.

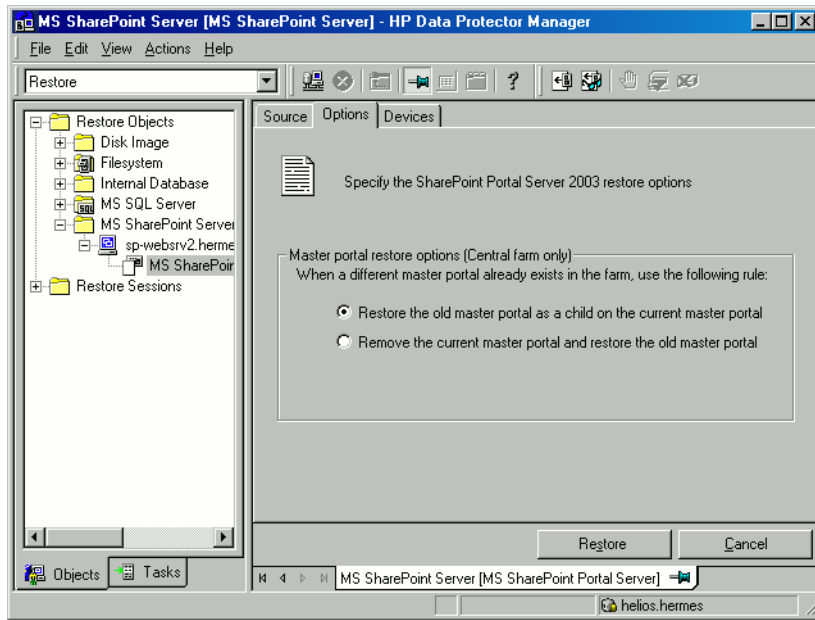


Figure 27 SPS restore options

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

The **Automatic device selection** option is selected by default, but it is recommended to select the **Original device selection** option.

 **IMPORTANT:**

If you decide to select the **Automatic device selection** option, ensure that the number of available devices is equal to or greater than the number of devices that were used for backup.

6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.

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

Click **Finish** to start the restore.

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

Table 14 SPS restore options

Option	Description
Restore the old master portal as a child on the current master portal	Applicable only if your SPS farm is centralized (having the master and child portals) and you are restoring the master portal. If this option is ON, the old master portal is restored as a child portal on the current master portal. Default: ON.
Remove the current master portal and restore the old master portal	Applicable only if your SPS Server farm is centralized (having the master and child portals) and you are restoring the master portal. If this option is ON, the current master portal becomes a child of the restored master portal. Default: OFF.

Restoring using the Data Protector CLI

From the directory `Data_Protector_home\bin`, run:

```
omnir -mssps -barhost front_end_server MSSPS_options
```

where `MSSPS_options` are:

```
[-portal virtual_server {[-teambd db_name SQL_options] [-index index_options] [-sitedbs sitedbs_options]]} [-ssodb SQL_options] [-doclib -session sessionID] [-changemaster]
```

`SQL_options` are:

```
-session sessionID [-tohost client] [-instance instance] [-as new_dbname]
```

`index_options` are:

```
-session sessionID [-tohost client] [-todir directory]
```

and `sitedbs_options` are:

```
-session sessionID [-tohost client] [-instance instance]
```

Parameter description

<i>front_end_server</i>	Front-end Web server system. In a farm with multiple front-end Web servers, specify the same front-end Web server system that was used for the backup.
<i>virtual_server</i>	SPS virtual server of the portal to be restored.
-teamdb	Specifies a portal's content database for restore.
-session	Specifies a backup session to restore from.
-tohost	Specifies a client system to restore to. By default, SPS objects are restored to the original clients.
-instance	Specifies an SQL Server instance to restore to. By default, SQL Server databases are restored to the original SQL Server instances.
-as	Specifies a new name for the database to be restored. By default, databases are restored with the original names.
-index	Specifies the portal's index server for restore.
-todir	Specifies a directory to which to restore the portal's index server. By default, index servers are restored to the original directories.
-sitedbs	Specifies the portal's site databases (<i>portal_name_PROF</i> , <i>portal_name_SERV</i> , <i>portal_name_SITE</i>) for restore.
-doclib	Specifies the document library for restore.
-ssodb	Specifies the single sign-on database for restore.
-changemaster	Applicable only if you are restoring the master portal. If this option is specified, the current master

portal becomes a child of the restored master portal.

Example

To restore the content database `TeamDB1` (from the backup `2006/10/9-34`) and Index Server (from the backup `2006/2/7-31`) of the portal created on the SPS virtual server `Virtual1` that belongs to the farm with the `Front1.company.com` front-end Web server system, run:

Example

```
omnir -mssps -barhost Front1.company.com -portal Virtual1  
-teamdb TeamDB1 -session 2006/10/9-34 -index -session  
2006/2/7-31
```

To restore the content database `TeamDB1` of the portal created on the SPS virtual server `Virtual2` that belongs to the farm with the `Front1.company.com` front-end Web Server system, under the name `TeamDB1_backup`, from the backup `2006/10/9-34`, run:

```
omnir -mssps -barhost Front1.company.com -portal Virtual2  
-teamdb TeamDB1 -session 2006/10/9-34 -as TeamDB1_backup
```

Restoring using another device

You can restore using a device other than that used for backup.

Using the Data Protector GUI

On how to specify another device for restore using the Data Protector GUI, see the online Help index: "restore, selecting devices for".

Using the Data Protector CLI

If you are restoring using the Data Protector CLI, specify the new device on the Cell Manager in the file:

```
Data_Protector_home\Config\Server\cell\restoredev
```


Disaster recovery

Disaster recovery is a very complex process, involving products from different vendors. Therefore, you need to check the instructions from the database/application vendor on how to prepare for disaster recovery.

To do a disaster recovery of a SharePoint Portal Server farm:

1. Reinstall the SharePoint Portal Server farm (including Windows Server 2003, IIS, SQL Server). For details, see the Microsoft SharePoint Portal Server documentation.
2. Reinstall the Data Protector integration components and import the SharePoint Portal Server clients in the Data Protector cell.
3. Recover the master databases of all the SQL Server instances. For details, see [“Recovering the master database”](#) on page 52.
4. Restore IIS databases to corresponding front-end Web server clients using the Data Protector filesystem restore.
5. Restore the encryption key (if it was used for the single sign-on service). For details, see the Microsoft SharePoint Portal Server documentation.
6. Restore other SPS objects (site databases, content databases, index servers, document library, single sign-on database) using the Data Protector SPS integration.

Monitoring sessions

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

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

On how to monitor a session, see the online Help index: “viewing currently running sessions”.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the SPS integration.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- On the client system, examine system errors reported in the `debug.log` located in `Data_Protector_home\log`.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see online Help.
- For each client in the SPS farm, ensure that the Data Protector `Inet` service is running under a Windows domain user account that is in the Windows Administrators group.

Additionally, if your configuration or backup failed:

- Ensure that the SPS Server and SQL Server instances are online.

Additionally, if your backup failed:

- Check the configuration of the SPS farm as described in “[Checking the configuration](#)” on page 73.

Problems

Problem

Restore of a portal fails

If the Site databases (`portal_name_SITE`, `portal_name_SERV`, `portal_name_PROF`) of a portal are deleted or corrupted, SPS Server may lock the configuration database, which prevents you from restoring the portal (and the Site databases).

Action

1. Restart the services of the SPS Server and the SQL Server instance to unlock the configuration database.
2. Restart the restore.

Problem

Restore fails with “Object reference not set to an instance of an object”

When you start a restore of individual content databases (not the whole portal) to an unextended virtual server, an error similar to the following is displayed:

```
[Critical] From: OB2BAR_SPS_BAR@siska.hermes.com "MSSPS" Time:
10.1.2007 16:12:32 SPS_FarmRestore failed with error
System.NullReferenceException: Object reference not set to
an instance of an object.
```

Action

1. Extend the virtual server.
2. Restart the restore.

You do not need to extend the virtual server when restoring the whole portal.

Problem

Restore from a differential backup fails

When you start a restore from a differential backup, an error similar to the following is displayed:

```
[Critical] From: OB2BAR_Main@siska.hermes.com "MSSPS" Time:
21.12.2006 11:55:57

There are no objects in the Data Protector Internal Database
for object 'TESTNOOK1_SERV'.
```

You cannot restore from a differential backup if you do not have a full backup.

Action

Ensure that a full backup exists before you run a differential backup.

Problem

Backup of index servers fails

If the SPS services run under a user account different than the SPS administrator account and you start a backup of index servers, the following error is displayed:

```
Exception occurred during backup of search server! Access is denied.
```

Action

1. Start the SPS services under the SPS administrator account.
2. Restart the backup.

Problem

Restore to another SQL Server instance fails

If you start a restore of SQL Server databases to an SQL Server instance that does not exist or is not online, an error similar to the following is displayed:

```
[Critical] From: OB2BAR_SPS_BAR@sp-websrv2.hermes.com "MSSPS"  
Time: 1/4/2007 9:21:56 AM
```

```
RestorePortal failed with error  
System.Data.SqlClient.SqlException: SQL Server does not exist  
or access denied.
```

Action

1. Ensure that the SQL Server instance that you plan to restore to exists, configure it for use with Data Protector, and ensure that it is online.
2. Restart the restore.



NOTE:

The error message may also contain a part starting with:

```
Unhandled Exception: System.Data.SqlClient.SqlException:
```

This is a Microsoft issue and can be fixed by installing the hotfix KB 904422. After installing the hotfix, this part of the error is not displayed. For more information, see:

<http://support.microsoft.com/?id=904422>

3 Integrating Microsoft Exchange Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server integration. It describes the concepts and methods you need to understand to back up and restore Microsoft Exchange Server (**Exchange Server**) database objects.

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

Table 15 Exchange Server online backup types

Full database backup	Includes all data (database and all log files) regardless of the changes made after the last backup.
Incremental backup	Includes log files only. Refers to the previous full or incremental backup, whichever was performed last. After the backup, log files are deleted. Before you run an incremental backup, ensure that a full database backup exists. Otherwise, a restore from such an incremental backup session fails.

Using the Exchange Server integration, you can back up and restore the whole server or particular databases listed below:

- Microsoft Exchange Server (Microsoft Information Store)
- Microsoft Exchange Server (Microsoft Key Management Service)
- Microsoft Exchange Server (Microsoft Site Replication Service)
- Single mailboxes. See [Chapter 4](#) on page 115.

This chapter provides information specific to this integration. For general Data Protector procedures and options, see the online Help.

Integration concepts

Data Protector integrates with Exchange Server through the Data Protector `ese_bar.exe` executable, installed on Exchange Server. It controls the activities between Exchange Server and Data Protector backup and restore processes.

You can perform interactive and scheduled full and incremental database backups. The last full backup combined with incremental prevents data loss if a disk failure occurs. Transaction logs are backed up to perform rollforward recovery.

Exchange Server databases are grouped into **storage groups**. Exchange Server 2000/2003 supports up to 4 storage groups and up to 5 databases per storage group. Exchange Server 2007 supports up to 50 storage groups and up to 50 databases, where each storage group is limited to a maximum of 5 databases. The databases within a storage group are backed up sequentially. Storage groups are backed up in parallel. The maximum number of devices used in a session equals the number of storage groups you want to back up.

Using the Data Protector User Interface, you define which objects and object versions to restore. Data Protector then passes the information about objects and backup versions to the backup API. General Media Agents are started and the data flows from the media to the target Exchange Server. See [Figure 28](#).

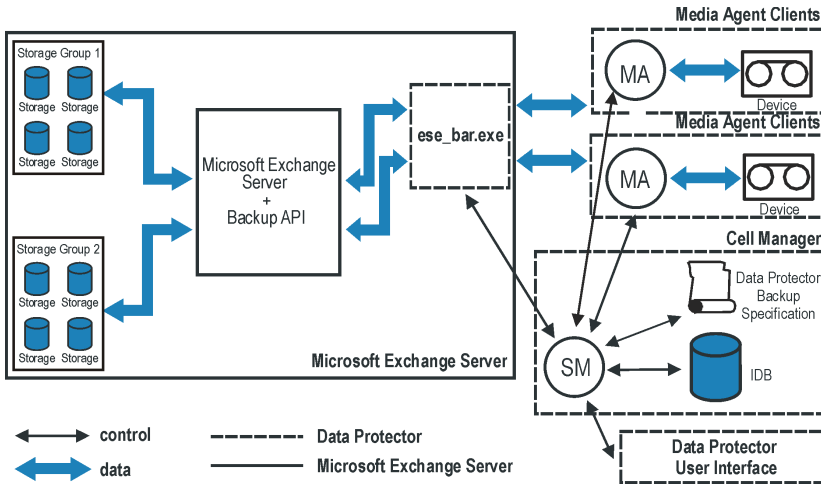


Figure 28 Data Protector Exchange Server integration architecture

Table 16 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
Backup API	The Microsoft defined interface that enables data transfer between Data Protector and Exchange Server.
MA	Data Protector General Media Agent.
Storage Group	A collection of mailbox stores and public folder stores that share a set of transaction log files.

Configuring the integration

Prerequisites

- You need a license to use the Exchange Server integration. See the *HP Data Protector installation and licensing guide* for information.
- Ensure that you correctly installed and configured Exchange Server.
 - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or <http://www.hp.com/support/manuals>.

- For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing the Data Protector Exchange Server integration, see the *HP Data Protector installation and licensing guide*. Every Exchange Server system to be used with Data Protector must have the MS Exchange Integration component installed.

Limitations

- Due to incompatibility between Microsoft Exchange Server 2007 and earlier Exchange Server versions, Exchange Server 2007 backup objects cannot be restored to Data Protector clients on which an earlier Exchange Server version is installed, and vice versa.

Before you begin

- Configure devices and media for use with Data Protector. See the online Help index: “configuring devices” and “creating media pools” for instructions.
- To test whether Exchange Server and Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore. See the online Help for instructions.
- Before performing incremental backups, disable **circular logging** for all storage groups.

If the application is cluster-aware, disable circular logging on all cluster nodes.

- Add the `Exchange_home\bin` directory to the Windows **Path** environment variable:
 1. In the Windows Explorer, right-click **My Computer** and click **Properties**.
 2. In the **Properties** dialog box, click **Advanced** and then **Environment Variables**.
 3. Select **Path** in the **System Variables** list and click **Edit**.
 4. Add `Exchange_home\bin` in the **Variable Value** text box and click **OK**.

See [Figure 29](#).

If the integration is cluster-aware, perform this procedure on all cluster nodes.




Figure 29 Path system variable

Backup

To run an online backup of an existing Exchange Server backup specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

For information on starting interactive backups using the CLI, see the `omnib` man page.

Limitations

- Backup preview is not supported.

Considerations

- You can perform incremental backups only if circular logging is disabled for the involved Exchange Server.

Circular logging is a Microsoft Exchange mode, where transaction logs are automatically overwritten when the data they contain is committed to the database. If enabled, this option reduces disk storage space requirements, but does not allow you to perform incremental backups.

- Do not use double quotes (" ") in object-specific pre and post-exec commands.

Configuring Exchange Server Backup

To configure a backup:

1. Configure devices and media for backup.
2. Create a Data Protector Exchange Server backup specification.

Creating backup specifications

Create a backup specification using the Data Protector Manager:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange Server**, and click **Add Backup**.

3. In the **Create New Backup** dialog box, select the **Blank Microsoft Exchange Server Backup** template, and click **OK**.

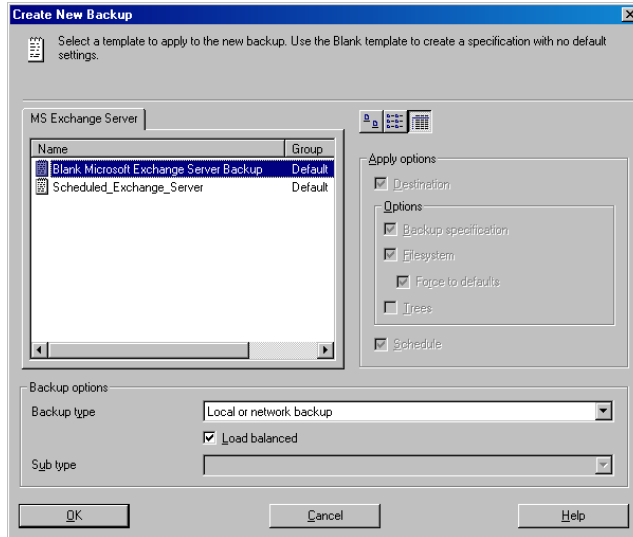


Figure 30 Selecting a blank template

4. In **Client**, select Exchange Server. For cluster environments, select the virtual server of the Exchange Server resource group.

In **Application database**, select one of the following:

- **Microsoft Exchange Server (Microsoft Information Store)**
- **Microsoft Exchange Server (Microsoft Key Management Service)** (if installed)
- **Microsoft Exchange Server (Microsoft Site Replication Service)** (if installed)

Click **Next**.

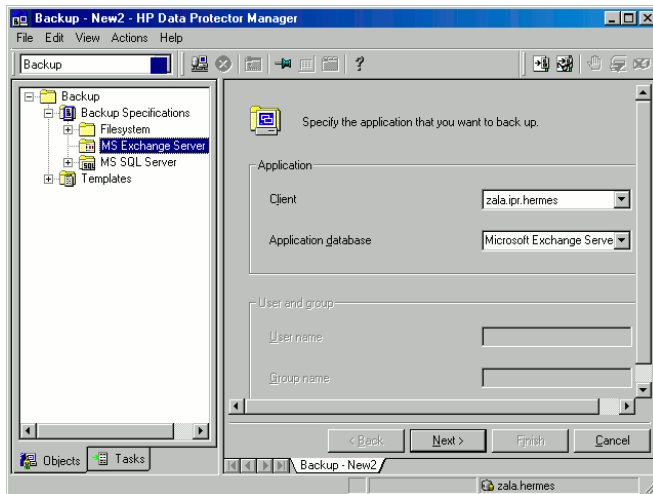


Figure 31 Client name and application database

5. Select Exchange Server databases you want to back up.

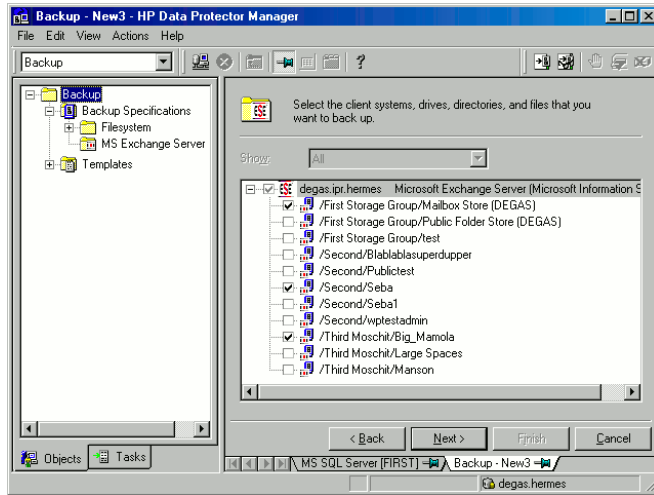


Figure 32 Backup objects

Click **Next**.

6. Select the device(s). Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on options, click **Help**.

To create additional backup copies (mirrors), click **Add mirror/Remove mirror**. Select separate devices for each mirror. The minimum number of devices for mirroring equals the number of devices used for backup.

For information on object mirroring, see the online Help index: “object mirroring”.

 **NOTE:**

The recommended maximum device concurrency is two for devices connected directly to the server, and one for those connected remotely.

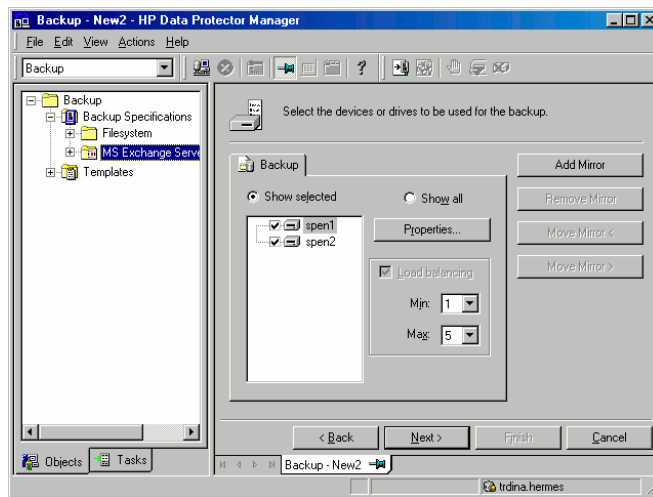


Figure 33 Backup devices

Click **Next** to proceed.

7. Select the backup options.

For information on **Backup Specification Options** and **Common Application Options**, see the online Help.

For information on **Application Specific Option**, see [SQL Server Specific Backup Options](#) or the online Help.

Click **Next**.

8. Optionally, schedule the backup. For information on scheduler, press **F1**.

9. Save the backup specification.

Once saved, the backup specification can be started by clicking **Start Backup**.

Exchange Server specific backup options

You access these options from the **Options** property page by clicking the **Advanced** button next to **Application Specific Options**.

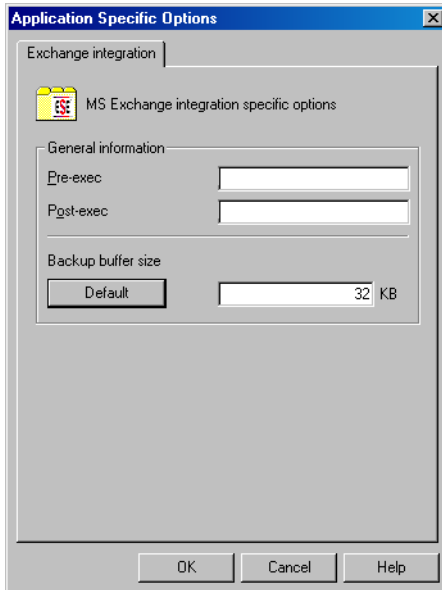


Figure 34 Application specific options

Table 17 Application specific options

Pre-exec	Specifies a command with arguments or a script started on Exchange client before backup. Only the filename must be provided in the backup specification.
Post-exec	Specifies a command with arguments or a script started on Exchange client after backup. Only the filename must be provided in the backup specification.
Backup buffer size	The size of the Exchange buffer, used for transferring data to Data Protector.

 **NOTE:**

Pre- and post-exec scripts must reside in the `Data_Protector_home\bin` directory on the Exchange Server.

Scheduling backups

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

Scheduling example

To schedule a database backup at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon, Tue, Wed, Thu, and Fri**.

Click **OK**.

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

Incremental backup backs up transaction log files that record changes to the database. Exchange Server automatically deletes `transaction log` files after they are backed up.

Starting backup sessions

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

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. Right-click the backup specification you want to start and select **Start Backup**.

3. Select **Backup type** and **Network load**. For information on these options, click **Help**. Click **OK**.

Click **OK**.

Restore

You can restore Exchange Server databases using the Data Protector GUI or CLI.

IMPORTANT:

The database (store) must be unmounted before a restore.

To unmount a database, use the Exchange Administration GUI:

1. In the **Exchange System Manager** window, right-click the backed up object (**Mailbox Store** or **Public Folder Store**), and select **Dismount Store** from the pop-up menu.

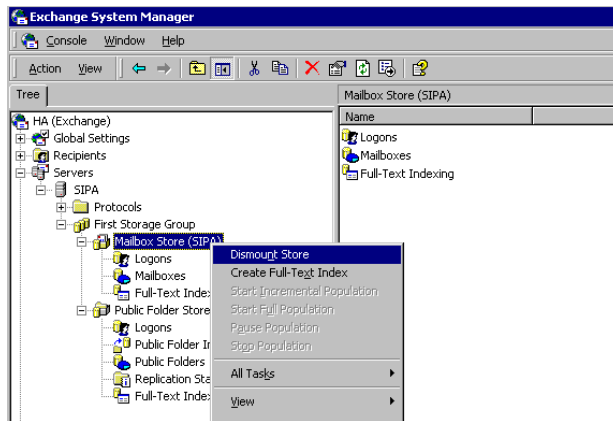


Figure 35 Unmounting a database

2. A warning appears. Click **Yes** to continue unmounting.

When unmounting completes, you may start restore.

After hard recovery, databases can be mounted automatically. See [Table 18](#) for details.



NOTE:

Log files for storage groups are saved in the subdirectory of the specified log directory.

Restoring using the GUI

Proceed as follows using the Data Protector Manager:

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects, MS Exchange Server**, and then select the client from which you want to restore. A list of backed up objects is displayed in the Results Area.

3. Select the restore objects.

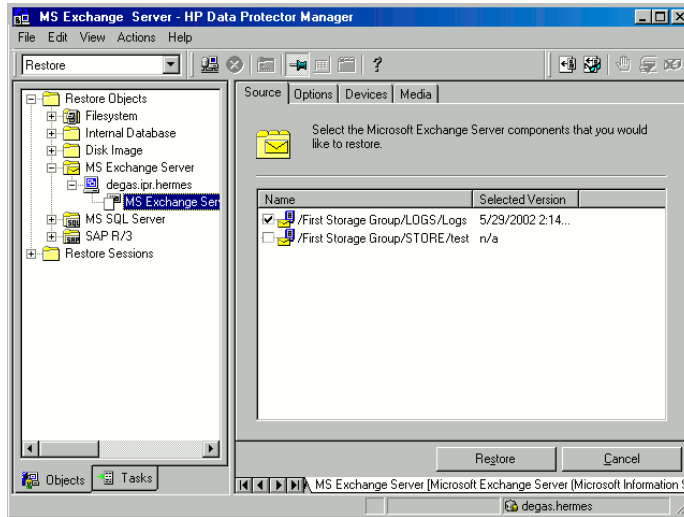


Figure 36 Restore objects

To select a backup version, right-click the object and select **Properties**.

 **IMPORTANT:**

When restoring several databases from the same storage group, ensure that their backup versions are the same. Otherwise, you need to restore them in separate sessions.

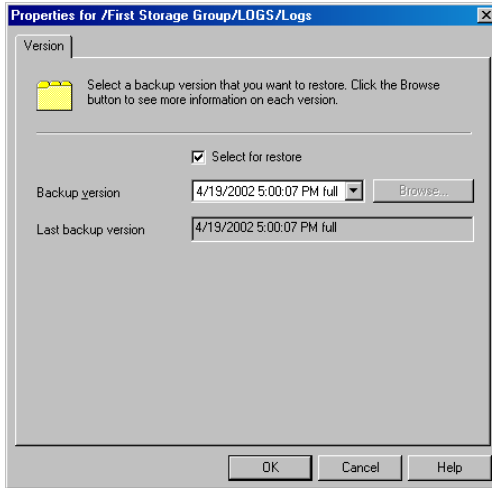


Figure 37 Selecting a backup version



NOTE:

Restoring databases to a certain state often requires a multiphase restore (multiple versions must be restored to retrieve data). During incremental backups, only transaction logs of storage groups are backed up (without information on physical location of the storage groups); therefore, you must restore last full backup first and then all transaction log backups made after the last full backup.



IMPORTANT:

When restoring from a full database backup, make sure you selected database files and transaction logs from the same version.

Example

Suppose you have the following backup sequence:

F T T *F T T T*

and you want to restore the version marked T, then restore all versions in *italic*: first full and transaction log backup, second transaction log backup, and the last transaction log backup (**Last restore set (start recovery)** selected).

4. In the **Options** property page, select restore options. See [Table 18](#) for details.
5. Click **Devices** and **Media** to select devices, verify device information, and set media priorities.

Note that you can use a different device for restore than the one used for backup. See the online Help index: “selecting, devices for restore”.

 **IMPORTANT:**

If the devices for restore are not those used for backup, select the same number of devices in the **Devices** property page as you used during backup.

6. Click **Restore**. Review your selection and click **Finish** to start restore.

If **Mount databases after recovery** is not specified, mount dismounted Information Stores after restore using the Exchange System Manager.

Table 18 Exchange Server restore options

Restore to another client	By default, the target client is the Exchange Server from which the application data was backed up. Nevertheless, you can restore databases to a different Exchange Server. The new target server must be a part of the Data Protector cell and have the MS Exchange Integration component installed.
Directory for temporary log files	Sets the temporary directory for log files restore. Using this directory, Exchange Server recovers the database - this is called hard recovery .
Last restore set (start recovery)	Performs a hard recovery after restore. Use to restore the last set of files. If you do not set this option, start the recovery manually by running <code>eseutil /cc /t</code> from the appropriate subdirectory of the directory for temporary log files.
Mount databases after recovery	Automatically mounts restored databases after hard recovery.
Last consistent state	Restores the database to its last consistent state. The latest log files, created after backup, are applied to the restored database during recovery.

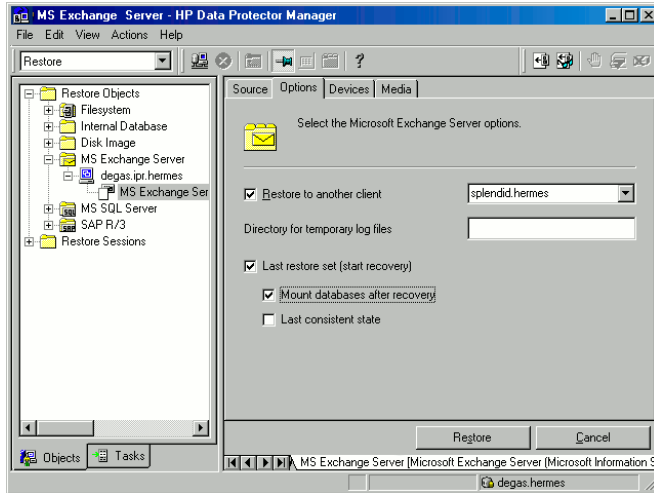


Figure 38 Restore options

Restoring to another client

1. Install the same version of Exchange Server on a separate system, and install the same Exchange Server Service Pack version(s).



NOTE:

New system name can be different.

2. On the new Exchange Server, create *all* storage groups that existed on the backed up Exchange Server. For every storage group, use the *same* name, location, and parameters.
3. For every newly created storage group, create *all* stores (databases) that existed in this particular storage group on the backed up Exchange Server. When creating a store, use the *same* name, location, and parameters.
4. Install the Data Protector Exchange integration on this system.

5. Restore the last full backup of the Exchange Server database. Follow the normal restore procedure using the Data Protector GUI and set the following options in the **Options** property page:
 - Restore to another client and specify the target client name.
 - The directory for temporary log files on the target client, for example `c:\EsseRestore`.
 - Last restore set (start recovery) to restore the last set of files (if you have no incremental backups of the last full backup).

See [Table 18](#) for details.

6. Restore all subsequent incremental backups and specify the same directory for temporary log files on the target client as for the restore of the last full backup.

When restoring the last incremental backup, select `Last restore set (start recovery)` to initiate automatic hard recovery of the Exchange Server database. If this option is not set, start recovery manually by running `eseutil /cc /t` from the temporary log files directory.

If hard recovery is initiated after restore of the last set of files (`Last restore set (start recovery)` selected), temporary log files are deleted after recovery.

Restoring using the CLI

From the `Data_Protector_home\bin` directory, run:

```
omnir -msese
-barhost ClientName [-destination ClientName]
-appname full_application_name {-base DBName -session SessionID}...
-logpath Path [-last [-mount] [-consistent]]
```

Provide the `Session_ID` of the backup session. For object copies, do not use the copy session ID, but the object backup ID (equals the object backup session ID).

For the description of options, see the `omnir` man page.

Example

To restore Information Store with the `/First Storage Group/STORE/Public Folder Store store` and `/First Storage Group/LOGS/Logs logs` to `computer.company.com` (where it was backed up), using the session ID `2003/07/07-13`, plus restore log files to `c:\temp`, perform hard recovery after restore and mount the database after hard recovery, run:

```
omnir -msese -barhost computer.company.com -appname "Microsoft Exchange Server (Microsoft Information Store)" -base "/First Storage Group/LOGS/Logs" -session "2003/07/07-13" -base "/First Storage Group/STORE/Public Folder Store" -session "2003/07/07-13" -logpath c:\temp -last -mount
```

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Exchange Server integration. Start at [Problems](#). If you cannot find a solution there, perform general checks and verifications.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Check that Exchange Server services (Microsoft Exchange System Attendant and Microsoft Exchange Information Store) are running.
- Using Exchange System Manager, check that all stores to be backed up are mounted and all stores to be restored are dismounted.
- Perform a backup of the Exchange Information Store using Windows Backup. If the backup fails, fix Exchange Server problems first, and then perform a backup using Data Protector.
- Ensure that the Cell Manager is correctly set on Exchange Server by checking the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBack II\Site
```


Its name and value must be `CellServer` and "`Cell Manager hostname`", respectively.

- Examine system errors reported in `Data_Protector_home\log\debug.log` on Exchange Server functioning as a Data Protector client. Additionally, examine the errors reported in the Windows Event log.
- Check if the following directories exist on the Data Protector Cell Manager:
`Data_Protector_home\config\server\barlists\msese`
`Data_Protector_home\config\server\barschedules\msese`
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Create a backup specification to back up to a null or file device and run the backup. If the backup succeeds, the problem may be related to backup devices. See the *HP Data Protector troubleshooting guide* for instructions on troubleshooting devices.
- Try to restart the Microsoft Exchange Server and start the backup again.
- Check that the `Exchange_home\bin` directory is added to the Windows `Path` environment variable. For details, see "[Configuring the integration](#)" on page 95.
- When performing incremental backups, ensure that circular logging is disabled by starting Exchange System Manager and selecting **Properties** from the storage group you are backing up.
- If you cannot mount the storage after a successful restore, check that LOGS storage on the same storage group is also restored.
- Define a directory for temporary log files in the **Restore** context. Check if the specified directory exists. If it does not, create it or specify another existing directory.
- To restore to another system, make sure Exchange Server is installed on that system and has the same organization and site names as the restored server.

Problems

Problem

Restore fails

The following message is reported:

```
[Critical]
Target Instance, specified for restore, is not found or log files
do not match the backup set logs.
```

The problem occurs when there is a gap in the sequence of the restored and the current log files.

Action

At the command prompt, run `eseutil` from the directory with temporary log files of the corresponding storage group:

- If the storage group name consists only of the ASCII characters A-Z, a-z, 0-9, and space, run the following command from *Storage_group_name*:

```
eseutil /cc /t
```

- If the storage group name consists of unicode characters, proceed as follows:
 1. One of the subdirectories in the temporary log file directory contains an empty file whose filename equals the name of the storage group you are restoring. Identify the subdirectory where the file is located. The subdirectory name conforms to the following template:

```
Storage Group Number
```

2. Run:

```
Drive_letter:
```

```
cd "\Temporary_log_files_directory_path\Storage Group  
Number"
```

```
eseutil /cc /t
```

4 Integrating Microsoft Exchange Single Mailbox and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Single Mailbox integration (**Exchange Single Mailbox integration**). It describes concepts and methods you need to understand to back up and restore mailboxes and Public Folders from or to a Microsoft Exchange Server system.

You can back up the entire content of a mailbox or Public Folders, including e-mail messages, task assignments, calendar schedules, contacts, and so on (**Exchange items**). Or you can back up only individual Exchange items from different mailboxes and Public Folders.

Data Protector integrates with Microsoft Exchange Server (**Exchange Server**) to back up and restore Exchange items online, enabling the Exchange Server to be actively used during the session.

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

Table 19 Exchange Single Mailbox backup types

Full	Backs up all selected Exchange items.
Incr1	Backs up changes made to selected Exchange items since the last full backup.
Incr	Backs up changes made to selected Exchange items since the last backup of any type.

You can restore Exchange items:

- To the original location.

- To a new folder, created in the root of the mailbox or All Public Folders.
- To another mailbox.
- To another Exchange Server system.

This chapter provides information specific to the Data Protector Exchange Single Mailbox integration. For general Data Protector procedures and options, see online Help.

Integration concepts

The main component of the Data Protector Exchange Single Mailbox integration is `mbx_bar.exe`, installed on the Exchange Server system, which channels communication between the Data Protector Session Manager, and, via the **MAPI interface**, the Exchange Server. [Figure 39](#) on page 116 shows the architecture of the Data Protector Exchange Single Mailbox integration.




Figure 39 Exchange Single Mailbox integration architecture

Legend:

MAPI	The Messaging Application Programming Interface, enabling applications and messaging clients to interact with messaging and information systems.
SM	The Data Protector Session Manager, which controls the session.

mbx_bar.exe	The Data Protector component started by SM that logs in through the MAPI profile to the Exchange Server administrator's mailbox, establishing an MAPI session. Having access to all other mailboxes, <code>mbx_bar.exe</code> logs in to each mailbox selected for backup or restore and initiates data transfer between Exchange Server and Data Protector media.
MA	The Data Protector General Media Agent.
IDB	The Data Protector internal database.

While the Exchange Server is responsible for read/write operations to disk, Data Protector reads from and writes to devices, and manages media.

Configuring the integration

Configure every Exchange Server you intend to back up from or restore to and the corresponding Exchange Server users.

Prerequisites

- Ensure that you have correctly installed and configured Exchange Server.
 - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or <http://www.hp.com/support/manuals>.
 - For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- On Microsoft Exchange Server 2007 systems, ensure that the package Microsoft Exchange Server MAPI Client and Collaboration Data Objects is installed. The package can be obtained free of charge from the Microsoft web site <http://www.microsoft.com/downloads/Search.aspx?displaylang=en>.
- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HP Data Protector installation and licensing guide*.
Every Exchange Server system you intend to back up from or restore to must have the Data Protector `MS Exchange Integration` component installed.

Limitations

- The Data Protector Exchange Single Mailbox integration is supported only on Exchange Server systems. You cannot back up and restore Exchange items from or to other clients.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the Exchange Server system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the Exchange Server system.

Cluster-aware clients

Configure the integration on all cluster nodes.

Configuring Exchange Server users

Add the Exchange Server administrator to the Data Protector `admin` or `operator` user group. For information, see the online Help index: “adding users” and “user groups”.

See the Exchange Server documentation for further information on different types of connections, roles and permissions of Exchange Server administrators, and security issues.

Configuring Exchange servers

Provide Data Protector with the name, password, and domain of the Exchange Server administrator. Data Protector then creates the Exchange Server configuration file on the Cell Manager and verifies the connection to the Exchange Server.



IMPORTANT:

Reconfigure the Exchange Server every time the Exchange Server administrator’s password changes.

Prerequisites

- Ensure that the Exchange Server is online.

Configure the Exchange Server using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange Single Mailboxes**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, click **OK**.
4. In **Client**, select the Exchange Server system. In a cluster environment, select the virtual server of the Exchange Server resource group.

Click **Next**.

5. In the **Configure Single Mailbox** dialog box, provide the username, password, and domain of the Exchange Server administrator.

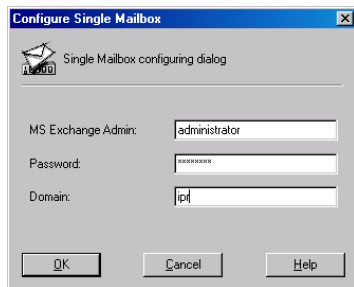


Figure 40 Configuring the Exchange Server

Click **OK**.

6. The Exchange Server is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#) on page 122.

Checking the configuration

You can check the configuration of the Exchange Server after you have created at least one backup specification for the Exchange Server.

Check the Exchange Server configuration using the Data Protector Manager.

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS Exchange Single Mailboxes**. Click the backup specification to display the Exchange Server to be checked.

3. Right-click the Exchange Server and click **Check configuration**.

Backup

The integration provides online backups of the following types:

Table 20 Exchange Single Mailbox backup types

Full	Backs up all selected Exchange items.
Incr1	Backs up changes made to selected Exchange items since the last full backup.
Incr	Backs up changes made to selected Exchange items since the last backup of any type.

Limitations

- Backup sessions that back up the same mailbox cannot run simultaneously.
- The Data Protector Exchange Single Mailbox backup is slower and requires more media space than the Data Protector Exchange Server backup. In the latter case, a message that has been sent to several recipients is saved only once and linked to all recipients, whereas in the first case, the entire message is saved for each recipient separately.

 **IMPORTANT:**

Do not use Data Protector Exchange Single Mailbox backups as a replacement for Data Protector Exchange Server backups. The latter are still needed to successfully recover a crashed system. For information, see “[Backup](#)” on page 97.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange Single Mailboxes**, and click **Add Backup**.

3. In the **Create New Backup** dialog box, select the template you want to use.

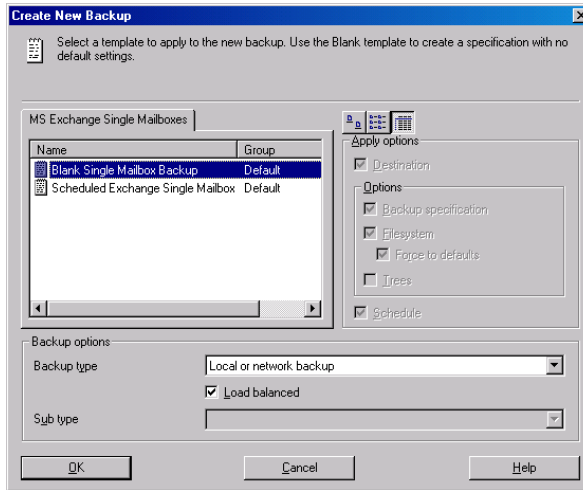


Figure 41 Selecting a template

4. In **Client**, select the Exchange Server system. In a cluster environment, select the virtual server.

Click **Next**.

5. If the Exchange Server is not configured for use with Data Protector, the **Configure Single Mailbox** dialog box is displayed. Configure it as described in [“Configuring Exchange servers”](#) on page 118.

6. Select the Exchange items you want to back up.

Mailboxes are organized alphabetically. For example, mailboxes starting with the letter S are collected under the S folder.

To back up all mailboxes and Public Folders, select the Exchange Server system at the top. Or you can browse for and select individual mailboxes and Public Folders or individual Exchange items from different mailboxes and Public Folders.

 **NOTE:**

Empty folders will not be backed up.

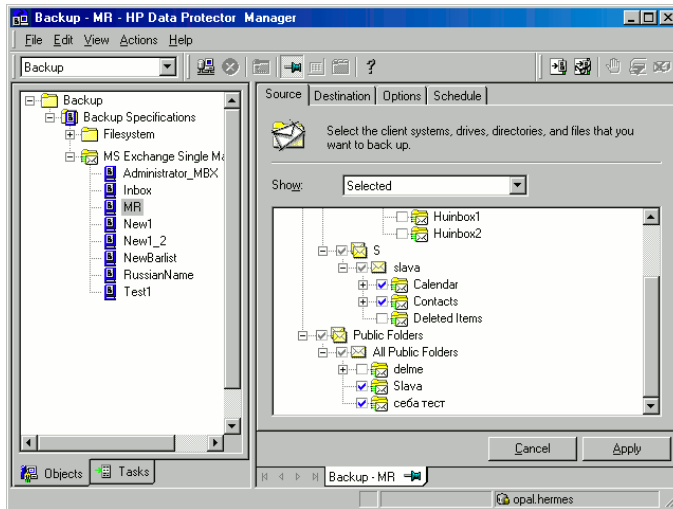


Figure 42 Selecting Exchange items for backup

Click **Next**.

7. Select devices to use for the backup.

To specify device options (for example, the device concurrency and the media pool to be used), right-click the device and click **Properties**.

Click **Next**.

8. Set backup options. For information on application specific backup options (Figure 43 on page 123), see Figure 42 on page 122.

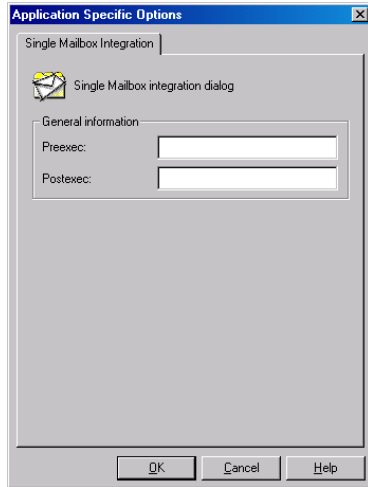


Figure 43 Exchange Single Mailbox specific backup options

Click **Next**.

9. Optionally, schedule the backup. See “Scheduling backup specifications” on page 77.

Click **Next**.

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

 **TIP:**

Preview your backup specification before using it for real. See “Previewing backup sessions” on page 78.

Table 21 Exchange Single Mailbox specific backup options

Option	Description
Pre-exec, Post-exec	Specify a command to be run by <code>mbx_bar.exe</code> on the Exchange Server system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes.

Option	Description
	Type only the name of the command and ensure that the command resides in the <i>Data_Protector_home\bin</i> directory on the Exchange Server system.

Modifying backup specifications

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

Scheduling backup specifications

You can run unattended backups at specific times or periodically. For details on scheduling, see the online Help index: "scheduled backups".

Scheduling example

To perform Incr1 backups of selected Exchange items at 14:45, 18:00, and 20:00 on Sundays:

1. In the **Schedule** page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **14:45**. Under **Recurring Options**, select **Sun**. Under **Session Options**, select the **Incr1** backup type. See [Figure 44](#) on page 125.
Click **OK**.
3. Repeat [Step 1](#) on page 124 and [Step 2](#) on page 124 to schedule backups at 18:00 and 20:00.
4. Click **Apply** to save the changes.

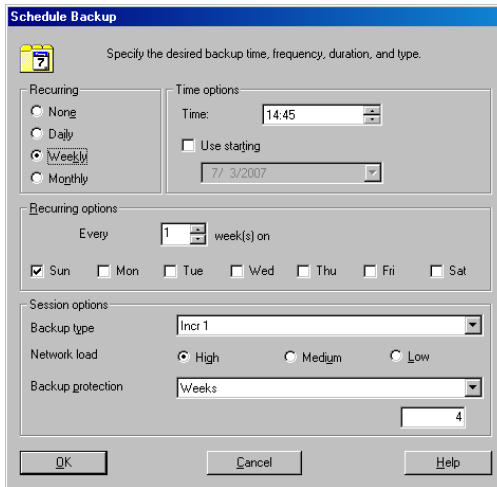


Figure 44 Scheduling a backup specification

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

Run:

```
omnib -mbx_list backup_specification_name -test_bar
```

What happens during the preview?

The following are tested:

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

After that, the Exchange Server part of the preview starts, which checks if the selected Exchange items are in an appropriate state for backup.

Starting backup sessions

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

Use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS Exchange Single Mailboxes**. Right-click the backup specification you want to start and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

On the Exchange Server system, run:

```
omnib -mbx_list backup_specification_name [-barmode mailbox_mode] [list_options]
```

where *mailbox_mode* is one of the following:

```
{full|incr|incr1}
```

For *list_options*, see the `omnib` man page.

Example

To start an incremental backup using the backup specification `FIRST` and to set data protection to 5 days, run:

```
omnib -mbx_list FIRST -barmode incr -protect 5
```

Restore

Restore Exchange items using the Data Protector GUI or CLI.

Before you begin

- If you intend to restore Exchange items to another mailbox, ensure that the destination mailbox exists on the destination Exchange Server.
- If you intend to restore Exchange items to another Exchange Server system, ensure that the destination Exchange Server system has the `MS Exchange 2000/2003 Integration` component installed and that the Exchange Server is configured for use with Data Protector.
- If you intend to restore Exchange items from Data Protector A.05.50 backups, ensure that the `MBX_RESTORE_55 omnirc` variable is set to 1.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Exchange Single Mailboxes**, the client from which the data to be restored was backed up, and then click **MS Exchange Single Mailboxes**.

3. In the **Source** page, browse for and select Exchange items to restore.

To restore all mailboxes and Public Folders, select **Mailboxes** and **Public Folders**. Or you can browse for and select individual mailboxes and Public Folders or individual Exchange items from different mailboxes and Public Folders.

Mailboxes are organized alphabetically. For example, mailboxes starting with the letter **S** are collected under the **S** folder.

See [Figure 45](#) on page 128.

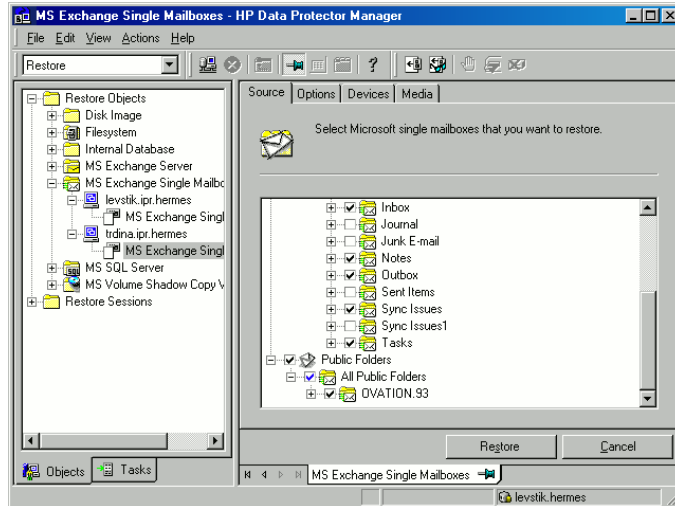


Figure 45 Selecting Exchange items for restore

You can specify the backup version, the chain of backups to be used, and the restore destination for each mailbox or Public Folders separately.

By default, the last backup session is used for restore. To restore from another session, right-click the relevant mailbox or **Public Folders**, and click **Properties**. See [Figure 46](#) on page 129.

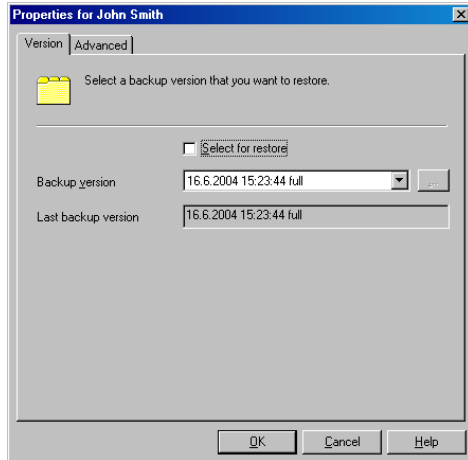


Figure 46 Version properties

To specify the restore destination and the chain of backup sessions to be used, click the **Advanced** tab. See [Figure 47](#) on page 129.

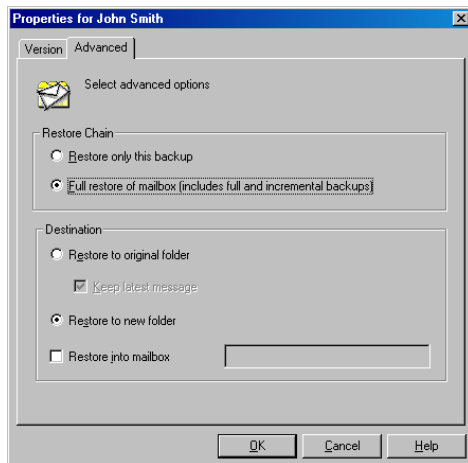


Figure 47 Advanced properties

For details on these options, see [Table 22](#) on page 131.

 **NOTE:**

Which Exchange items are displayed in the Results Area depends on the selected backup session and the **Restore Chain** options.

For example, if **Restore only this backup** is selected, only the Exchange items backed up in the selected session are displayed, whereas if **Full restore of mailbox** is selected, all Exchange items backed up in the restore chain of backup sessions are displayed.

The **Full restore of mailbox** and **Restore to new folder** options are selected by default.

4. In the **Options** page, specify the destination Exchange Server system. By default, the original Exchange Server system is selected.

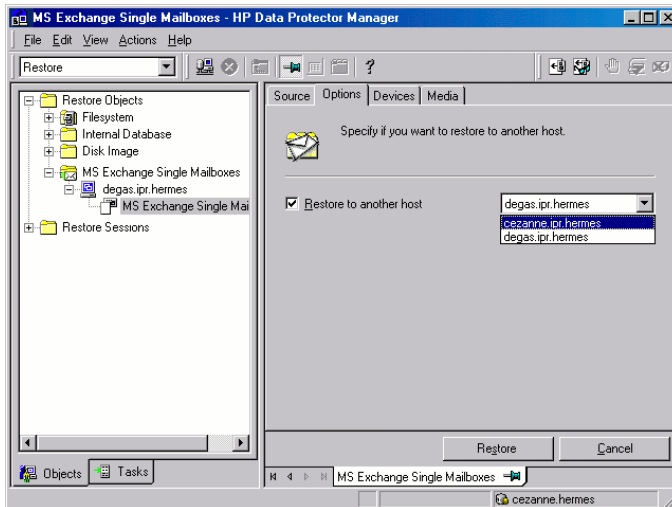


Figure 48 Selecting the destination Exchange Server system

5. In the **Devices** page, select devices to use for the restore.
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.

8. Specify Report level and Network load.

Click **Finish** to start the restore.

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

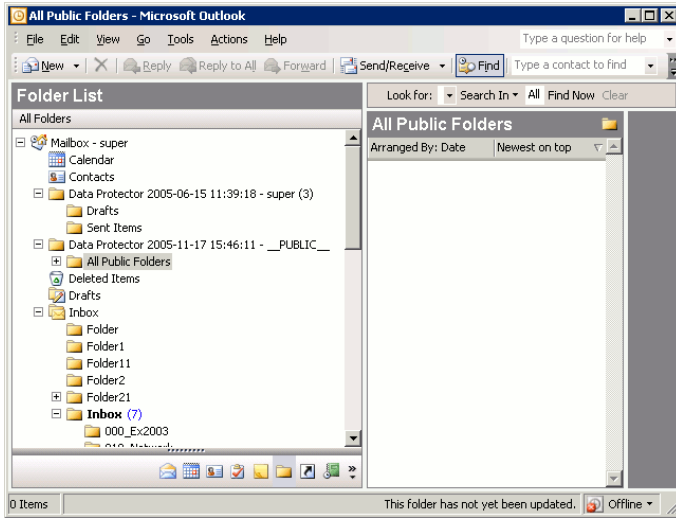


Figure 49 Restored mailbox and public folders content with the restore to new folder option selected.

To transfer restored data to .pst files:

1. On the client system, create a .pst file.
2. Connect to the Exchange Server system.
3. Move the restored data from the *Data Protector backup date backup time* folder or the *Data Protector backup date backup time - public folder* folder to the previously created .pst file.

Table 22 Exchange Single Mailbox restore options

Option	Description
Restore only this backup	Select this option to restore data only from the selected backup session.

Option	Description
Full restore of mailbox	<p>Selected by default. Data is restored, not only from the selected backup session, but also from the latest full, the latest incremental1 (if it exists), and any incremental backups from the last incremental1 up to the selected backup version.</p> <p>Note that any Exchange item that was backed up in any of these sessions is displayed and can be selected for restore.</p>
Restore to original folder	<p>Data Protector restores Exchange items to the same location from which they were backed up.</p> <p>If <code>Keep latest message</code> is selected, existing messages in the destination mailbox or Public Folders are not restored even if they differ from their backed up versions.</p> <p>If <code>Keep latest message</code> is not selected, all messages are restored, replacing their current versions (if they exist). If different versions of the same message exist in the mailbox or Public Folders (for example, if you have a copy of the message), only one is replaced with the backed up version and all other versions remain intact.</p> <p>The messages in the mailbox that were not backed up in the specified backup session (or the restore chain of backup sessions) always remain intact.</p> <p>By default, this option is not selected.</p>
Restore to new folder	<p>Selected by default. Data Protector creates a new folder in the root of the mailbox (or in the root of All Public Folders) and restores Exchange items into it. See Figure 49 on page 131.</p> <p>When restoring a mailbox, the folder is named <code>Data Protector backup_date backup_time</code>. When restoring Public Folders it is named <code>Data Protector backup_date backup_time - public folder</code>.</p> <p>If you restore from the same backup several times, a number is appended to the folder name. For example, in the second restore session of a mailbox, the folder <code>Data Protector backup_date backup_time (1)</code> is created.</p>
Restore into mailbox	<p>By default, Exchange items from a mailbox are restored to the original mailbox. Select this option to specify a different destination mailbox. Note that you can restore Exchange items from different mailboxes to the same mailbox.</p> <p>For privacy protection, you cannot restore Exchange items from mailboxes to Public Folders.</p>

Option	Description
Restore to another host	By default, Exchange items are restored to the original Exchange Server system. Select this option, to specify a different destination Exchange Server system.

Restoring using the Data Protector CLI

On the Exchange Server system, run:

```
omnir -mbx -barhost client_name [-destination
dest_client_name] -mailbox mailbox_name -session session_ID
[mailbox_options] -public -session session_ID
[Public_Folders_options]
```

Select among the following *mailbox_options*:

```
-destMailbox dest_mailbox_name -folder folder -exclude
ex_folder -originalfolder {-keep_msg | -overwrite_msg} -chain
```

Select among the following *Public_Folders_options*:

```
-folder folder -exclude ex_folder -originalfolder {-keep_msg
| -overwrite_msg} -chain
```

NOTE:

Specify

```
-mailbox mailbox_name -session session_ID [mailbox_options] as
many times as the number of mailboxes you want to restore.
```

Specify `-folder folder` and `-exclude ex_folder` as many times as the number of folders you want to restore or exclude from restore.

Parameter description

client_name

Original Exchange Server system, from which Exchange items to be restored were backed up.

dest_client_name

Destination Exchange Server system, to which the Exchange items will be restored (needed only if you are not restoring to the original Exchange Server system).

<i>session_ID</i>	Backup version ID. For object copies, use the object's backup ID (which equals the object's backup session ID). Do not use the object's copy session ID.
<i>mailbox_name</i>	Original mailbox, from which Exchange items to be restored were backed up. If the name contains a space, put the name in quotes. For example, "John Smith".
<i>dest_mailbox_name</i>	Destination mailbox, to which the Exchange items from the mailbox will be restored (needed only if you are not restoring to the original mailbox).
<i>folder</i>	Folder to be restored. Specify its pathname, starting from the root directory in the mailbox or Public Folders. If the pathname contains a space, put the pathname in quotes. For example, "Inbox\My folder".
<i>ex_folder</i>	Subfolder to be excluded from restore of the mailbox or Public Folders.

Option description

<code>-originalfolder</code>	This option is equivalent to the Data Protector GUI option <code>Restore to original folder</code> . If not specified, the same results occur as if the Data Protector GUI option <code>Restore to new folder</code> was selected.
<code>-chain</code>	This option is equivalent to the Data Protector GUI option <code>Full restore of mailbox</code> . If not specified, the same results occur as if the Data Protector GUI option <code>Restore only this backup</code> was selected.

Limitations

- If any of the mailbox names or folder names specified in the `omnir` command contains a slash (/), backslash (\), or double quote (") character, the restore fails.

Restore examples

Example 1

To restore the mailbox `FIRST`, backed up in the session `2005/01/10-1` from the Exchange Server system `infinity.ipr.hermes`, to a new folder in the mailbox `TEMP` on the same Exchange Server system, run:

```
omnir -mbx -barhost infinity.ipr.hermes -mailbox FIRST  
-session 2005/01/10-1 -destMailbox TEMP
```

Example 2

To restore the folder `Inbox` from the mailbox `User 1`, backed up in the session `2005/03/10-18` from the Exchange Server system `exchange.hp.com`, to the original folder without overwriting the messages in the original folder, run:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 1" -session  
2005/03/10-18 -folder Inbox -originalfolder -keep_msg
```

Example 3

To restore the mailbox `User 2`, backed up in the session `2005/03/10-19` from the Exchange Server system `exchange.hp.com`, to a new folder in the original mailbox, without restoring the messages from the folder `Deleted Items`, run:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 2" -session  
2005/03/10-19 -exclude "Deleted Items"
```

Example 4

To restore two public folders, `Administration` and `Addresses`, which are subfolders of `All Public Folders`, and the mailbox `My Mailbox`, backed up in the session `2005/06/10-19` from the Exchange Server system `exchange.hp.com`, to a new folder in `Public Folders` and to the original folders in the mailbox respectively, run:

```
omnir -mbx -barhost exchange.hp.com -public -session  
2005/06/10-19 -folder "All Public Folders\Administration"
```

```
-folder "All Public Folders\Addresses" -mailbox "My Mailbox"  
-originalfolder -keep_msg
```

Monitoring sessions

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

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

On how to monitor a session, see the online Help index: “viewing currently running sessions”.

Performance tuning

Performance tuning means customizing Exchange Server and Data Protector to achieve better backup and restore results.

Data Protector creates a separate backup object out of selected Exchange items from a single mailbox or Public Folders. This object is then backed up as a separate data stream. `mbx_bar.exe` spends a significant amount of time creating Data Protector backup objects and logging mailboxes on/off. Meanwhile, the Data Protector devices are in an idle state, waiting for the actual data transfer to start.

Backup performance can be enhanced by streaming two or more backup objects to the same device concurrently. While one stream is preparing the backup object and logging the mailbox on/off, data from the other backup object is being transferred to the tape, keeping the device busy.

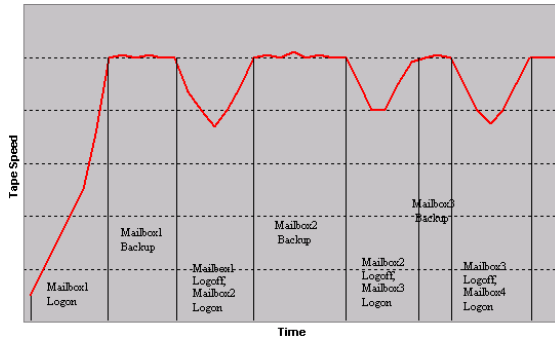


Figure 50 Example of backup with concurrency 1

Tests have shown that best performance is achieved when backing up mailboxes and Public Folders using two concurrent data streams, either by specifying one device with concurrency 2 or two devices with concurrency 1.

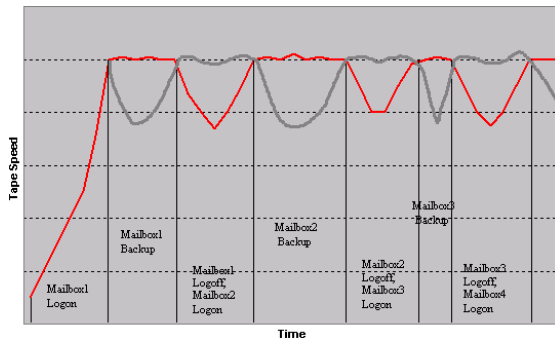


Figure 51 Example of backup with concurrency 2



NOTE:

Data Protector cannot create more than one backup object out of Exchange items from a single mailbox or Public Folders.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector Exchange Single Mailbox integration. Start at “Problems” on page 139. If you cannot find a solution there, perform general checks and verifications.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Ensure that the following directories exist on the Data Protector Cell Manager:
`Data_Protector_home\config\server\barlists\Mailbox`
`Data_Protector_home\config\server\barschedules\Mailbox`
- Examine errors reported in
`Data_Protector_home\log\debug.log` on the Exchange Server system.

Additionally, if your backup or restore failed:

- Ensure that the Cell Manager is correctly specified on the Exchange Server system: ensure that a value entry with the name `CellManager` and the value "`Cell Manager`" exists under the key:
`HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBack II\Site`
- Examine errors logged in the Windows Event log.

Additionally, if your backup failed:

- Preview the Data Protector Exchange Single Mailbox backup.
If the Exchange Server part of the preview fails, ensure that the Exchange Server is online.
If the Data Protector part of the preview fails:
 - Ensure that the Exchange Server is configured for use with Data Protector. See "[Configuring Exchange servers](#)" on page 118.

- Create an Exchange Single Mailbox backup specification to back up to a null or file device.

If the backup succeeds, the problem is probably related to devices. For information on troubleshooting devices, see online Help.

Problems

Problem

You do not have permissions to log in to the system

Data_Protector_home\log\debug.log on the Exchange Server contains one of the following messages:

Error = 596

Logon failure: the user has not been granted the requested logon type to this computer.

or:

[MBX_ImpersonateUser] A required privilege is not held by the client.

Action

Check if the Domain Controller system has domain-level policy settings defined. Go to:

Start > Settings > Control Panel > Administrative Tools > Domain Security Policy > Local Policies > User Rights Assignment

and check if the Act as part of the operating system and Log on as a service user rights are set to Defined.

If domain-level policy settings are defined:

1. On the Domain Controller system:
 - a. Go to:
Start > Settings > Control Panel > Administrative Tools > Domain Security Policy > Local Policies > User Rights Assignment.
 - b. Set Act as part of the operating system and Log on as a service user rights for the Exchange Server administrator.
 - c. Run:
`secedit /refreshpolicy machine_policy /enforce`
2. On the Exchange Server system:
 - a. Log off from the system and log in again under the same user account.
 - b. Go to:
Start > Settings > Control Panel > Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment.
 - c. Ensure that Act as part of the operating system and Log on as a service user rights are set for the Exchange Server administrator in both Local Setting and Effective Setting columns.
 - d. Restart the Data Protector Inet service.

If domain-level policy settings are not defined:

1. Log in to the Exchange Server system.
2. Go to:
Start > Settings > Control Panel > Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment.
3. Set Act as part of the operating system and Log on as a service user rights for the Exchange Server administrator.
4. Log off from the system and log in again under the same user account.
5. Restart the Data Protector Inet service.

Problem

Configuration of the Exchange Server fails

`Data_Protector_home\log\debug.log` on the Exchange Server system contains the following message:

An error has occurred while creating a profile administration object.

Action

1. Log in to the Exchange Server system.
2. Delete the incorrect administrator's profile:
`mbx_bar.exe delete`
3. Manually create a new profile:
`mbx_bar.exe create`
4. In the **Choose Profile** page, click **New**.
5. Follow the setup wizard. Type `$$$Data Protector` for the profile name. Specify the Exchange Server system and the name of the Exchange Server administrator's mailbox. See [Figure 52](#) on page 141.

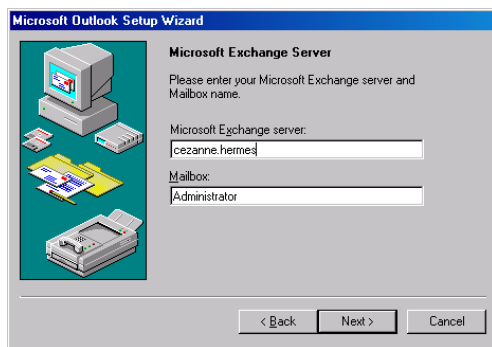


Figure 52 Specifying the Exchange Server administrator's mailbox

Problem

Restore to another client fails

Action

Ensure that Exchange Server and Data Protector **MS Exchange Single Mailbox integration** component are installed and configured on the destination system to which you restore.

Problem

Restore to another mailbox fails

Action

Ensure that the destination mailbox exists on the destination Exchange Server.

5 Integrating Microsoft Volume Shadow Copy Service with Data Protector

Introduction

A traditional backup process is based on the direct communication between the backup application and the application whose data is backed up. This backup method requires from the backup application an individual interface for each application it backs up.

The number of applications on the market is constantly increasing. The necessity of handling application specific features can cause difficulties in backup, restore, and storage activities. An effective solution to this problem is introducing a coordinator among the actors of the backup and restore process.

Volume Shadow Copy Service

Volume Shadow Copy Service (VSS) is a software service introduced by Microsoft on Windows operating systems. This service collaborates with the backup application, applications to be backed up, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.

Data Protector supports the integration with VSS.

The Data Protector Volume Shadow Copy integration provides a unified communication interface that can coordinate backup and restore of an application regardless of their specific features. With this approach, a backup application does not need to handle each application to be backed up specifically. However, the production application as well as the backup application must conform to the VSS specification.

[Figure 53](#) on page 145 and [Figure 54](#) on page 145 show the differences between the traditional backup model and the model with the Data Protector Microsoft Volume Shadow Copy integration.

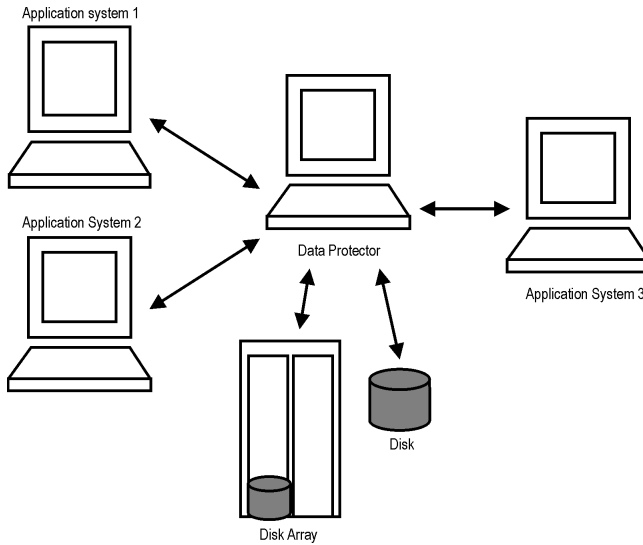


Figure 53 Actors of the traditional backup model

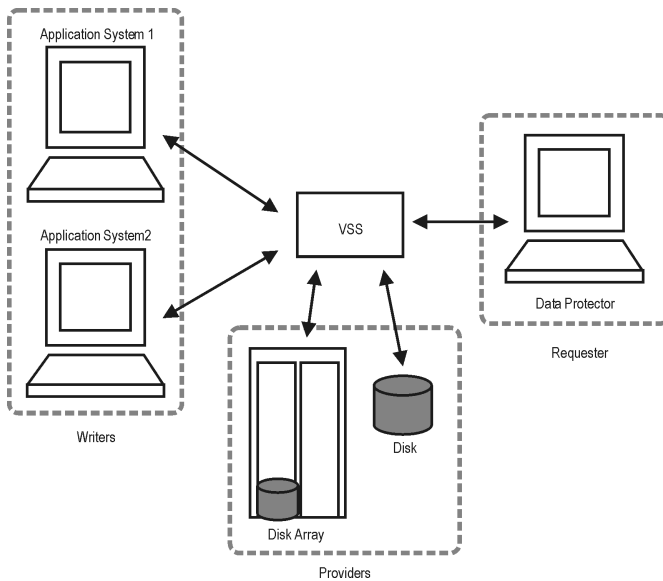


Figure 54 Actors of the Data Protector VSS integration backup model

Without using the Volume Shadow Copy Service, Data Protector has to communicate with each application to be backed up individually. The Data Protector VSS integration introduces a unified backup and restore interface and provides the coordination among the participants of the backup and restore process.

Advantages

The advantages of using the Data Protector Microsoft Volume Shadow Copy Service integration are the following:

- **Central Management for all backup operations**
The administrator can manage backup operations from a central point.
- **Media Management**
Data Protector has an advanced media management system that allows users to monitor media usage and set protection for stored data, as well as to organize and manage devices in media pools.
- **Backup Management**
Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.
- **Scheduling**
Data Protector has a built-in scheduler that allows the administrator to automate backups to run periodically. Using the Data Protector Scheduler, the backups you configure run unattended at specified times, as long as the devices and media are properly set.
- **Device Support**
Data Protector supports a wide range of devices: files, standalone drives, very large multiple drive libraries, etc.
- **Reporting**
Data Protector has reporting capabilities that allow you to get information on your backup environment. You can schedule reports to be issued at a specific time or to be attached to a predefined set of events, such as the end of a backup session or a mount request.
- **Monitoring**
Data Protector has a feature that allows the administrator to monitor currently running sessions and view finished sessions from any system that has the Data Protector GUI installed.
All backup sessions are logged in the IDB, which provides the administrator with a history of activities that can be queried later.

Integration concepts

The Data Protector integration with the Microsoft Volume Shadow Copy Service provides full support for certified VSS writers. This includes automatic detection of the VSS writers and backup and restore functionality.

For a complete list of supported VSS writers and providers refer to the latest support matrices at <http://www.hp.com/support/manuals>.

Benefits of using the integration

Advantages of using the Data Protector VSS integration are the following:

- Unified backup interface is provided for all applications that provide a writer.
- Data integrity is provided on application level, because it is provided by the writers. No interference is needed from the backup application.

VSSBAR agent

The central part of the integration is the **VSSBAR agent**, which links Data Protector with the Microsoft Volume Shadow Copy Service. Data Protector Microsoft Volume Shadow Copy integration uses the VSSBAR agent for automatic browsing of VSS-aware writers, coordinating backup and restore. VSSBAR agent is responsible for the following actions:

- detecting VSS writers
- examining and analyzing Writer Metadata Document (WMD)

Writer Metadata Document (WMD) is metadata provided by each writer. Writers identify themselves by the metadata and instruct the backup application what to back up and how to restore the data. Thus, Data Protector follows the requirements provided by the writer when selecting the volumes to be backed up and the restore method.

- requesting shadow copy creation
- backing up writers' data to media
- coordinating restore session start
- restoring the Writer Metadata Document
- restoring writer's data from media

Backup

During the Data Protector VSS integration backup, Data Protector does not interact directly with each writer, but through the VSS interface. It uses the VSSBAR agent to coordinate the backup process. The consistency of data is a responsibility of the VSS writer and not dependent on Data Protector functionality. The backup process of the VSS writers consists of the following phases:

1. When you select writers and components you want to back up and start a VSS integration backup, Data Protector communicates with the Volume Shadow Copy Service (backup coordinator) to notify that the backup is about to start.
2. The coordinator identifies all writers that support the VSS feature and passes the list of available writers and their characteristics (Writer Metadata Document) back to Data Protector.
3. Data Protector examines Writer Metadata and identifies the volumes that contain the data to be backed up. Then the VSS informs available writers about selected components.
4. Data Protector prepares a list of volumes (shadow copy set) that must be put into consistent state, and passes the list back to the coordinator for preparing a shadow copy.
5. The VSSBAR agent notifies the writers about the shadow copy creation. The VSS mechanism ensures that there are no writes on the volume while the shadow copy is being created.



NOTE:

When the VSSBAR agent creates a shadow copy of the volume, this volume is marked in order to avoid attempts to simultaneously create another shadow copy of the same volume. In order to prevent any deadlocks arising from volume locking, only a single VSSBAR agent at a time is allowed to define a shadow copy set.

6. When the writers are fully prepared for the consistent shadow copy backup, the VSSBAR agent passes shadow copy creation requests to VSS.
7. After a shadow copy is created, the VSS service returns the related information to the Data Protector.

8. Data Protector backs up the data from the shadow copy to media and then notifies the VSS service that the shadow copy can be released. VSS issues a command to the shadow copy provider to destroy the shadow copy that has been already backed up. Figure 55 on page 149 shows the relations between the actors of a local or network VSS backup.

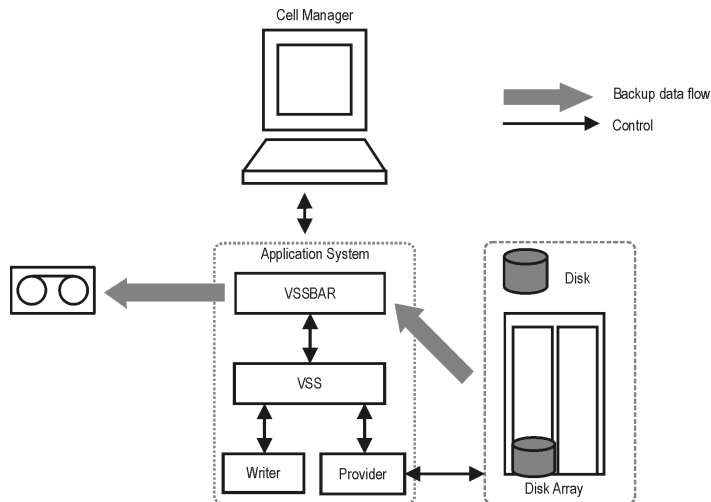


Figure 55 Local or network VSS backup

Data consistency

The **filesystem backup** does not guarantee application data consistency, only filesystem consistency. Application data consistency can be achieved only by using supported application writers, such as Microsoft Exchange Server writer.

Restore

Data Protector offers two restore modes:

- **component restore** using the VSS service
- **file restore** using the DMA instead of VSS.

By default, Data Protector restores writer components using the VSS service.

Instant recovery is also available within the Data Protector VSS integration. This functionality requires VDS hardware providers.

Restoring components

During the restore procedure, the Data Protector VSS integration coordinates communication between Data Protector and the writers. In general, the restore flow consists of the following phases: preparing for restore, restoring components, and notifying the application writers that a restore has been completed. The restore procedure of the VSS-aware writers consists of the following phases.

1. Data Protector first restores the metadata, which was collected during the backup. Then it examines the metadata to identify the backup components and determine the restore method. It also checks if restore to specific volumes is possible.
2. Data Protector connects to the coordinator (VSS service) to notify that the restore is about to start, which in turn communicates with the writer. Data Protector restores the data from the backup media to the locations specified in the backup metadata. During the restore, Data Protector follows the writers' instructions regarding any additional checking or processing specified in the WMD.
3. After the data are successfully restored from the backup media, Data Protector informs the coordinator that the restore is completed and the writers can now access the newly-restored data and start the internal processing, for example recovery.

Restoring files

For a successful restore of a VSS component, all files comprising this component must be restored. If a restore of a single file fails, the restore of the a whole component fails. Data Protector offers an additional restore mode for restoring single files that does not use the Microsoft Volume Shadow Copy Service, thus solving this problem. This mode can also be used for restoring to systems that do not support VSS or do not have a VSS writer installed.

When restoring files or a group of files, DMA is started and the files are restored using the standard Data Protector filesystem restore procedure.

IMPORTANT:

As the file restore mode does not utilize VSS services, additional tasks that are performed after a component restore – such as database recovery – are not performed and your application data may be left in an inconsistent state, requiring additional manual procedures before the application is recovered.

Microsoft Exchange Server 2007 writer concepts

Backup

With Microsoft Exchange Server 2007 writer, you can use two models of replication to protect your data:

- Local Continuous Replication (LCR)
With LCR, you can create and maintain an exact copy (LCR copy) of databases in a storage group. An LCR copy is used in case of data corruption 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 load on a production database is minimal.
- Cluster Continuous Replication (CCR)
The same characteristics as with LCR apply to CCR. The difference is only that in the CCR environment, databases and logs are replicated to separate servers. Therefore, a CCR copy can be used for disaster recovery. You can perform backups using VSS on your passive Exchange Server node where the CCR copy is located and thus reducing the load on the active node.

The replicated storage groups are represented as a new Exchange writer instance, Exchange Replication Service. They are backed up like original or production storage groups. Backups of the replicated storage groups (LCR and CCR copies) are possible only using VSS hardware provider.

Restore

With Microsoft Exchange Server 2007 writer, you can restore your data not only to the original location (from where the backup was performed) but also to a different location. You can restore:

- A whole storage group (including LCR or CCR copies)
- A single store (including LCR or CCR copies)

You can restore them to:

- Original storage group
- Different storage group
- Non-Exchange location. After restore is completed, Recovery Storage Group (RSG) can be created automatically.
- Recovery server (different client and different storage group)

If you perform restore to a different storage group, you can access single mailboxes or individual mails on an alternate location without the need to compromise the original storage group with a potentially unsuccessful restore operation. Also, if the whole server is destroyed, a restore to an alternate Exchange server (recovery server) will minimize downtime of your mailboxes.

Prerequisites and limitations

This is a list of prerequisites and limitations for the Data Protector Microsoft Volume Shadow Copy Service integration. Additional limitations and recommendations that are not directly connected to the integration (such as operating system and GUI limitations) and disk array limitations are listed in the *HP Data Protector product announcements, software notes, and references*.

Prerequisites

Common prerequisites

- Before you begin, ensure that you have correctly installed and configured Data Protector, writers and shadow copy providers. Refer to the:
 - *HP Data Protector product announcements, software notes, and references* for an up-to-date list of supported versions, platforms, devices, disk arrays, limitations, and other information.
 - *HP Data Protector installation and licensing guide* on how to install Data Protector on various architectures and how to install the Data Protector Microsoft Volume Shadow Copy Service integration.
 - Writers and shadow copy providers documentation for instructions on how to install and configure writers and providers on your system.

Limitations

For a list of general Data Protector limitations, see the *HP Data Protector product announcements, software notes, and references*.

Common limitations

- Maximum 64 shadow copies per volume are allowed. The number of shadow copies per volume is limited by system resources.
- To run a VSS integration backup, the writer's data must be on an NTFS filesystem.

- The VSS integration backup of writers which store their data on network shared volumes is not supported.
- The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore method. These writers are by default not presented by Data Protector.

If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. Refer to the writers documentation for additional information on the restore methods.

- Preview is only possible for VSS filesystem backup sessions.

Configuring the integration

The Data Protector Microsoft Volume Shadow Copy integration by itself does not require any configuration steps neither on the Data Protector nor on the application side, unless you are configuring a cluster-aware Data Protector VSS integration.

VSS writers are either a part of Windows operating system or delivered with applications. Data Protector automatically detects writers when the VSS backup specification is created and registers them.

You may check which writers and providers are installed and registered on your system using the following Windows operating system command:

- For a list of writers: `VSSadmin list writers`
- For a list of VSS providers: `VSSadmin list providers`
- VDS hardware providers should be present in the list of installed software. Check **Control Panel > Add/Remove Programs**.

Configuring the Data Protector VSS cluster-aware integration

The configuration of the Data Protector VSS cluster-aware integration consists of:

1. Configuring an VSS cluster-aware client. Refer to [“Configuring a VSS cluster-aware client”](#) on page 154.
2. Configuring a cluster-aware VSS integration backup. Refer to [“Configuring backups for a cluster-aware VSS client”](#) on page 154

Find below an overview of global configuration tasks with cluster-specific steps.

Configuring a VSS cluster-aware client

The client configuration must be performed on one cluster node per one VSS client, since the Data Protector VSS configuration file resides on the Cell Manager.

Configuring backups for a cluster-aware VSS client

To configure backups for a cluster-aware VSS client, create a Data Protector VSS backup specification, as explained in [“Creating backup specifications using GUI”](#) on page 160 taking into account the VSS infrastructure specifics described below.

The Microsoft VSS infrastructure does not identify writers that run as cluster resources (for example, cluster-aware writers). Therefore, the Microsoft VSS integration agent cannot distinguish between the cluster-aware and non-cluster-aware writers when creating a backup specification. This means you need to configure different backup specifications for cluster-aware and non-cluster-aware writers.

When backing up cluster-aware writers (such as SQL Server via the MSDE Writer), specify the name of the VSS client system as the virtual server name given in the particular writer resource group.

When backing up writers that are not cluster-aware (such as System Writer or Event Log Writer), specify the name of the VSS client system as the physical node.

Example 1. VSS cluster specifics

The example below shows why it is necessary to create different backup specifications for cluster-aware and non-cluster-aware writers.

You have `node_A` and `node_B`, and MS Exchange Server running on a virtual host `exchsvr`. When creating a backup specification, you can select, among others, MS Exchange Writer and Event Log Writer. Suppose, at the time of a backup, Exchange is running on `node_A`. If you create just one backup specification for both writers, the following problems will occur:

- If you select `node_A` as your source host, you have Event Log Writer and Exchange Server associated with `node_A`. While it is true, that Event Log Writer is a property of the physical node, it is wrong to associate Exchange Server with it, as it is a property of the virtual server.

Suppose that after a failover, MS Exchange Server is running on `node_B`. When you try to restore the data to `node_A`, the restore will fail because Exchange disks are now owned by `node_B` and you cannot write to them. However, the restore of Event Log Writer will succeed.

- If you select exchsvr as your source host, you have Event Log Writer and Exchange Server associated with a virtual server exchsvr. While it is true, that Exchange Server is a property of the virtual server, it is wrong to associate Event Log Writer with it, as it is a property of the physical node.

Suppose that after a failover, MS Exchange Server is running on `node_B`. When you try to restore Event Log Writer data to exchsvr, it will overwrite (or try to overwrite) the data in the Event Log of `node_B` with the data from the Event Log of `node_A`. The restore of Event Log Writer will fail, while Exchange Server will be restored successfully.

Microsoft Exchange Server writer specific configuration

A backup of the Microsoft Exchange database can be considered successful only if the consistency check of the replicated datafiles succeeds.

Microsoft Exchange Server 2007 configuration in LCR environment

When configuring an LCR environment, ensure that the original database and the database copy have the same directory structure for their files. Otherwise, after restore of an LCR copy to original database, mounting of the database will not be correct.

Writers specifics

This section describes specific information about VSS writers, that you need to take into account before backing up or restoring the writers.

VSS writers either come with the Windows operating system or with applications. For a complete list of supported VSS writers and providers refer to the latest support matrices at <http://www.hp.com/support/manuals>.

The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore. If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. Refer to the writers documentation for additional information on the restore methods.

 **NOTE:**

Writers requiring custom restore methods are by default not shown by Data Protector. The `omni.rc` variable `OB2VSS_SHOWALLWRITERS` must be set to 1 for all writers to be displayed.

Table 23 on page 156 provides a description of VSS writers.

Table 23 Writer description

Writer name	Description	Restore method
Certificate Authority Writer	This is a system writer, used to back up and restore Certificate Authority (CA) Service database. This service issues, revokes, and manages certificates employed in public key-based cryptography technologies.	Files are restored after a reboot.
Cluster Service Writer	This VSS writer using a custom API, is used to back up and restore Cluster Service on Microsoft Cluster Server (MSCS). The Cluster Service is a component on Windows servers used to control server cluster activities on cluster nodes. It is fundamental to the operation of the cluster.	Custom restore method
COM+ REGDB Writer	This VSS writer using a custom API, is used to back up and restore COM+ Database Service. This service provides automatic distribution of events to subscribing COM+ components.	Custom restore method
DHCP Jet Writer	This is a system writer, used to back up and restore DHCP Service database. DHCP Service provides dynamic IP address assignment and network configuration for Dynamic Host Configuration Protocol (DHCP) clients.	Files are restored after a reboot.
Event Log Writer	This is a system writer, used to back up and restore Event Logs. Event Logs are files where the Windows operating system saves information about events, such as starting and stopping services or the logging on and logging off of a user.	Files are restored after a reboot.

Writer name	Description	Restore method
FRS Writer	This VSS writer using a custom API, is used to back up and restore File Replication Service data. File Replication Service is a multithreaded replication engine that replicates system policies and logon scripts stored in System Volume (SYSVOL). FRS can also replicate data for Distributed File System (Dfs), copy and maintain shared files and folders on multiple servers simultaneously.	Custom restore method
IIS Metabase Writer	This is a system writer, used to back up and restore Microsoft Internet Information Server (IIS). IIS 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).	Files are restored after a reboot.
MSDE Writer	This is a writer used to back up and restore Microsoft SQL Server 2000. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language.	Refer to “MSDE writer restore specifics” on page 177.
Microsoft SQL 2005 Writer	This is a writer used to back up and restore Microsoft SQL Server 2005. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language.	Standard VSS restore and instant recovery.
Microsoft Data Protection Manager 2006 Writer	This is a writer used to back up and restore Microsoft Data Protection Manager 2006. Microsoft Data Protection Manager is a server that creates and stores replicas of clients and uses them for recovering the data on clients	Refer to “Microsoft Data Protection Manager 2006 writer restore specifics” on page 184.
Microsoft Exchange Server Writer	This is a writer used to back up and restore Microsoft Exchange Server. Microsoft Exchange Server is a mail and groupware server.	Standard VSS restore and instant recovery.
Microsoft Virtual Server 2005 Writer	This is a writer used to back up and restore Microsoft Virtual Server 2005. Microsoft Virtual Server 2005 is a virtualization platform for	Standard VSS restore and instant recovery.

Writer name	Description	Restore method
	<p>Microsoft Windows servers. Data Protector supports live back up of individual virtual machines and the Virtual Server configuration, ensuring data consistency of the backup and restore. Hardware providers are not supported if a Virtual Server Machine is in online mode; use a software provider or put the Virtual Server Machine in offline mode. For details about Virtual Server online and offline mode, see the Microsoft Virtual Server documentation.</p> <p>Cluster configurations are not supported, only individual nodes can be backed up.</p>	
NTDS Writer	<p>This is a system writer used to back up and restore Microsoft Active Directory on Windows servers. Active Directory Service is a Windows server directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides methods for storing directory data and making this data available to network users and administrators.</p>	<p>To restore Active Directory, boot into Directory restore mode. Files will be restored if they can be overwritten.</p>
Registry Writer	<p>This VSS writer using a custom API, is used to back up and restore Windows Registry. Windows Registry is a database repository of information containing the Windows system configuration.</p>	<p>Custom restore method</p>
Remote Storage Writer	<p>This is a system writer used to back up and restore Remote Storage Service (RSS). RSS is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened.</p>	<p>Files are restored after a reboot.</p>
Removable Storage Manager Writer	<p>This is a system writer used to back up and restore Removable Storage Manager Service. This service manages removable media, drives, and libraries.</p>	<p>Files are restored after a reboot.</p>
System Writer	<p>This is a system writer that backs up a specific set of Windows dynamic link libraries (DLL).</p>	<p>Files are restored after a reboot.</p>

Writer name	Description	Restore method
TermServlicensing Writer	This is a system writer that backs up Windows Terminal Services. These services provide a multi-session environment that allows client systems to access a virtual Windows desktop session and Windows-based programs running on the server.	Files are restored after a reboot.
WINS Jet Writer	This is a system writer, used to back up and restore Windows Internet Name Service (WINS). WINS is a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on a TCP/IP network.	Files are restored after a reboot.
WMI Writer	This is a system writer, used to back up and restore Windows Management Instrumentation (WMI). WMI is a unified management infrastructure in Windows for monitoring system resources.	Files are restored after a reboot.

Backing up writers data

Overview

To run backups and restores of the VSS writers and filesystems, you need to configure the Data Protector Microsoft Volume Shadow Copy Service integration backup specifications.

To configure the backup using the VSS integration, perform the following steps:

Configuration steps

1. Configure devices, media, and media pools needed for the backup. See the online Help for instructions.
2. Create a Data Protector VSS backup specification specifying the VSS components to back up, the media and devices to which you want your data to be backed up, as well as the Data Protector backup options that define the behavior of your backup or restore session.

Overview

Replica creation and reuse

A new replica is created and added to the replica set when the specified **Number of replicas rotated** is not reached.

The oldest replica in the set is deleted and the new one is created when the specified **Number of replicas rotated** is reached.

Creating backup specifications using GUI

The procedure below shows how to back up Microsoft VSS objects using the Data Protector GUI. Some writers have specific limitations. For writers specific limitations, see:

- For Microsoft Exchange Server specifics, see [“Microsoft Exchange Server writer backup specifics”](#) on page 163.
- For Microsoft Data Protection Manager 2006 specifics, see [“Microsoft Data Protection Manager 2006 writer specifics”](#) on page 167.

You need to configure different backup specifications for cluster-aware and non-cluster-aware writers. Refer to [“Configuring the Data Protector VSS cluster-aware integration”](#) on page 153.

To create a new backup specification for the VSS integration, proceed as follows:

- 1.** In the **HP Data Protector Manager**, switch to the **Backup** context.
- 2.** In the Scoping Pane, expand **Backup Specifications**.
- 3.** Right-click **MS Volume Shadow Copy Writers** and click **Add Backup**.
- 4.** In the **Create New Backup** dialog box, select **Local or network backup** for the backup type.
- 5.** In Application system, specify the name of the client that has the VSSBAR agent installed.

When backing up cluster-aware writers (such as SQL Server via the MSDE Writer, or Exchange Server in the CCR environmen), specify the virtual server name given in the particular writer resource group.

Click **Next**.

- 6.** This page only displays the selection of your VSS client. Click **Next**.

7. Select the backup objects you want to back up.

You can specify a **full client backup** by selecting the top-level item (the name of the client), a single writer or a writer's component backup by selecting a lower-level item.

If full client is selected, Data Protector checks which writers exist on the client and backs up all of them at backup time.

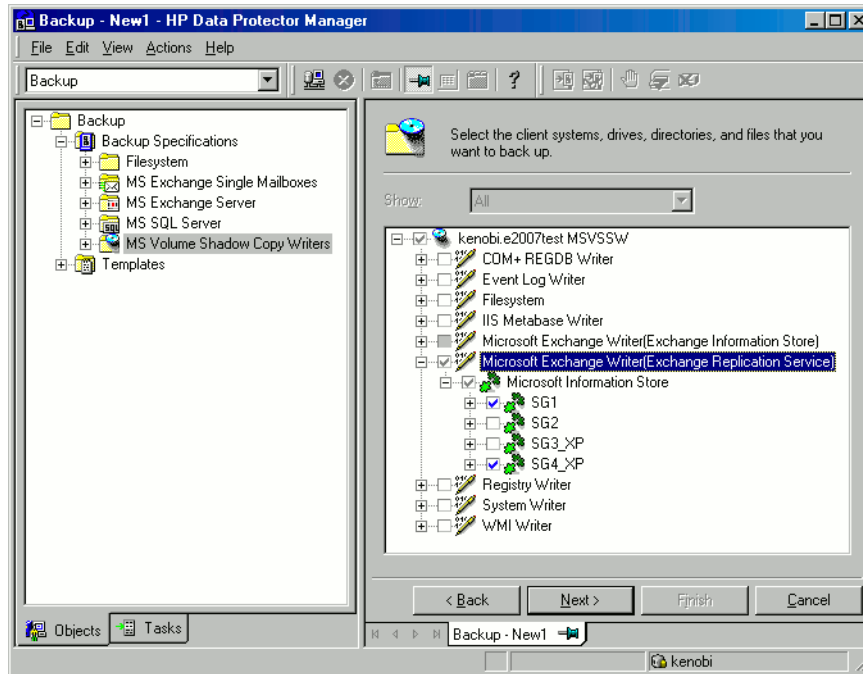


Figure 56 Selecting backup objects (Microsoft Exchange Server 2007 CCR)

If a writer requires all of its components to be backed up, lower-level items are automatically selected. If you select such a writer for backup, all its components will be backed up.

If a writer has no components to be backed up, it is not displayed in the list of writers, and is not backed up when the full client is selected.

The **Filesystem** item displays all mounted disks. If another disk is mounted under a disk, the parent disk is displayed twice. The first item is the parent disk (for example `c:`) and the second one is the container for the mountpoint (for example `c:\mnt\1`). To select a mounted disk, select the mountpoint. A hardware provider is required for this backup object.

Microsoft Exchange Writer: Optionally, to specify options for the consistency check of Microsoft Exchange Server Writer, right-click **Microsoft Exchange Writer** and click **Additional options**.

Microsoft Exchange Server 2007 Writer in a CCR environment: Optionally, to specify cluster options, right-click **Microsoft Exchange Writer** and click **Additional options**.

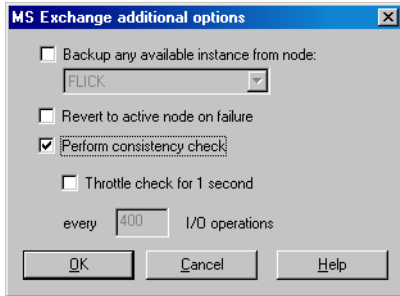


Figure 57 Additional options for Microsoft Exchange Server 2007 in CCR environment

8. Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**. If you do not select a device, *only backup to disk* will be available.

You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror.

For detailed information on the object mirror functionality, see the online Help.

9. Following the wizard, select the backup options and schedule your backup.

 **TIP:**

If you are not sure about selecting the backup options, keep the default values.

See the online Help for details about the options common to all Data Protector backup specifications.

10. Once you have defined all backup options and the schedule, you need to name and save the newly-created backup specification. You have now completed the creation of a Microsoft Volume Shadow Copy Writers backup specification.
11. You can review the newly-created and saved backup specification in the **Backup** context, under the specified group of backup specifications.
12. You can run backup using one of the following methods:
 - Schedule the backup of an existing Microsoft Volume Shadow Copy Writers backup specification using the Data Protector Scheduler.
 - Start an interactive backup of an existing Microsoft Volume Shadow Copy Writers backup specification.

VSS specific backup options

Table 24 VSS specific backup options

Option	Description
Pre-exec	Specify a command that will be started by <code>vssbar.exe</code> on the application system before the backup. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in <code>Data_Protector_home\bin</code> .
Post-exec	Specify a command that will be started by <code>vssbar.exe</code> on the application system after the backup. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in <code>Data_Protector_home\bin</code> .

Microsoft Exchange Server writer backup specifics

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

Backup types

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

- *Full* - backs up databases, transaction logs, and checkpoint files. The transaction logs are truncated.
- *Incremental* - backs up the transaction logs to record changes since the last full or incremental backup. The transaction logs are truncated. Requires Service Pack 1.

- *Differential* - similar as incremental backup, but the transaction logs are not truncated. Requires Service Pack 1.
- *Copy* - a Full backup, but the logs are not truncated. This type of backup is not intended for use in recovering failed systems.

Limitations

- A combination of VSS snapshot backups and incremental stream backups is not possible.
- You can back up only the whole server or full storage groups. Single stores cannot be backed up.
- Incremental and differential backups cannot be mixed in one restore chain.
- Circular logging must be disabled; otherwise, only full backup recovery is possible.
- With Exchange Server 2003, only one VSS backup session can be running on the same application system at once. With Exchange Server 2007, only one VSS backup session backing up a specific storage group can be running on the same application system at once. Consequently, if you start a backup while another one is in progress, the latter waits until the first one finishes.

Recommendations

Rollforward recovery

Transaction logs must be backed up to be able to perform the rollforward operation.

Consistency check

A backup of the Microsoft Exchange Server database is considered as successful only if the consistency check of the replicated datafiles succeeds. To enable consistency check, click on a created backup specification, right-click **Microsoft Exchange Writer** in the Source tab, and then click **Additional options**.

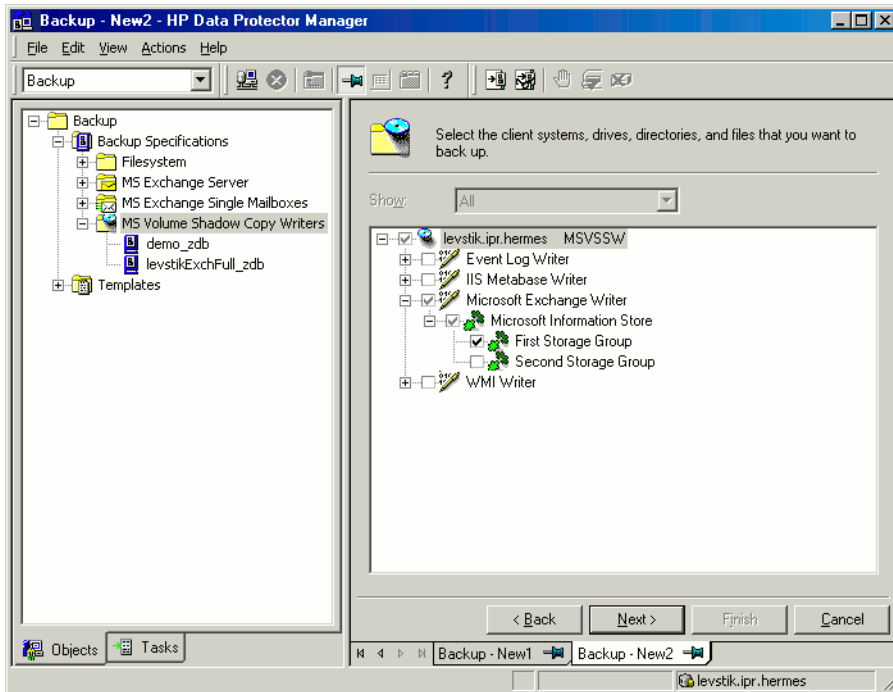


Figure 58 Selecting Microsoft Exchange Server 2003 storage groups

Microsoft Exchange Server 2007 writer backup specifics

In an LCR or CCR environment, the replicated storage groups are represented as a new instance of Exchange writer, **Exchange Replication Service**. The replicated storage groups are backed up in the same way as original production storage groups.

You can select any of these storage groups (original or replicated) for a backup. However, you cannot select both in the same backup specification. See [Figure 59](#).

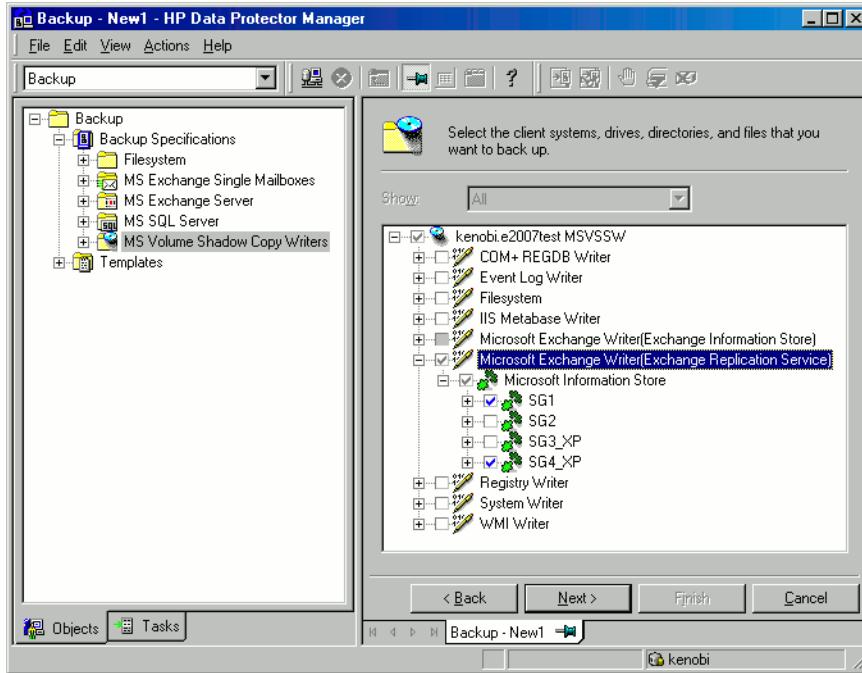


Figure 59 Selecting a replicated Microsoft Exchange Server 2007 storage group

In a CCR environment, you can select a cluster node from where you want your backup to be performed, regardless of which instance (Information Store or Replication Service) resides on this node. If you select the cluster node, Data Protector backs up any available instance on this node ignoring the selection of the instance in the GUI even if you, for example, select the replicated storage group (Exchange Replication Service) as backup object.

To specify the cluster node from where you want to perform backup of any instance residing on this node, right-click an Exchange writer and select the node in the additional options dialog box under **Back up any available instance from node**. See [Figure 57](#).

When backing up Replication Service instance, it may happen that it cannot be backed up due to the following reasons:

- The selected node is not available.
- The status of the storage group to be backed up is not "Healthy".
- Data Protector is not running on the selected node.
- `vssbar.exe` could not be started on the selected node.

To avoid the session failure, you can select the option **Revert to active node on failure** in the same dialog box. The backup will be restarted on the original server (active cluster node) and the original storage group will be backed up. If Information Store instance is backed up, this option is ignored.

Limitations

- In a CCR environment, if a storage group or a replicated storage group (Replication Service instance) is selected for backup, no other writer (for example, filesystem writer or SQL Server writer) can be selected to be backed up in the same backup session. This is because the two writers are located on different servers and an Data Protector limitation is that two different servers cannot be backed up in the same session.
- In an LCR or CCR environment, a storage group can control only one store.
- The LCR and CCR clusters can have only two nodes.

Considerations

- After any Microsoft Exchange Server backups, the transaction log files are truncated. The LCR and CCR Exchange backups guarantee 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.

Microsoft Data Protection Manager 2006 writer specifics

Microsoft Data Protection Manager 2006 (DPM) is a server application that creates replicas of the clients, synchronizes them through LAN, and stores these replicas as snapshots.

The Data Protection Manager writer is used to back up:

- the Data Protection Manager database and the Data Protection Manager Report database
- the *latest version* of the DPM replicas.

 **IMPORTANT:**

To ensure data consistency, schedule a DPM replica synchronization before starting a backup.

The DPM uses DPM snapshots for restore. These snapshots are *not* backed up. To be able to recreate DPM snapshots you must manually schedule a backup of the replica each time after the DPM creates a new replica.

Two backup types are supported:

- *Full* (for the DPM databases and replicas)
- *Incremental* (replicas only).

If you select unsupported backup types (*Copy* or *Differential*) when scheduling the backup, Data Protector will abort the backup and display an error message.

Prerequisite

The MSDE writer (used for backing up the DPM databases) must be installed.

Limitations

- Hardware providers are not supported with DPM.
- If you start a backup while an incremental synchronization of the DPM replica is still in progress, the backup gets corrupted, although Data Protector does not report any errors. In case of synchronization with consistency check, Data Protector automatically aborts the backup session.

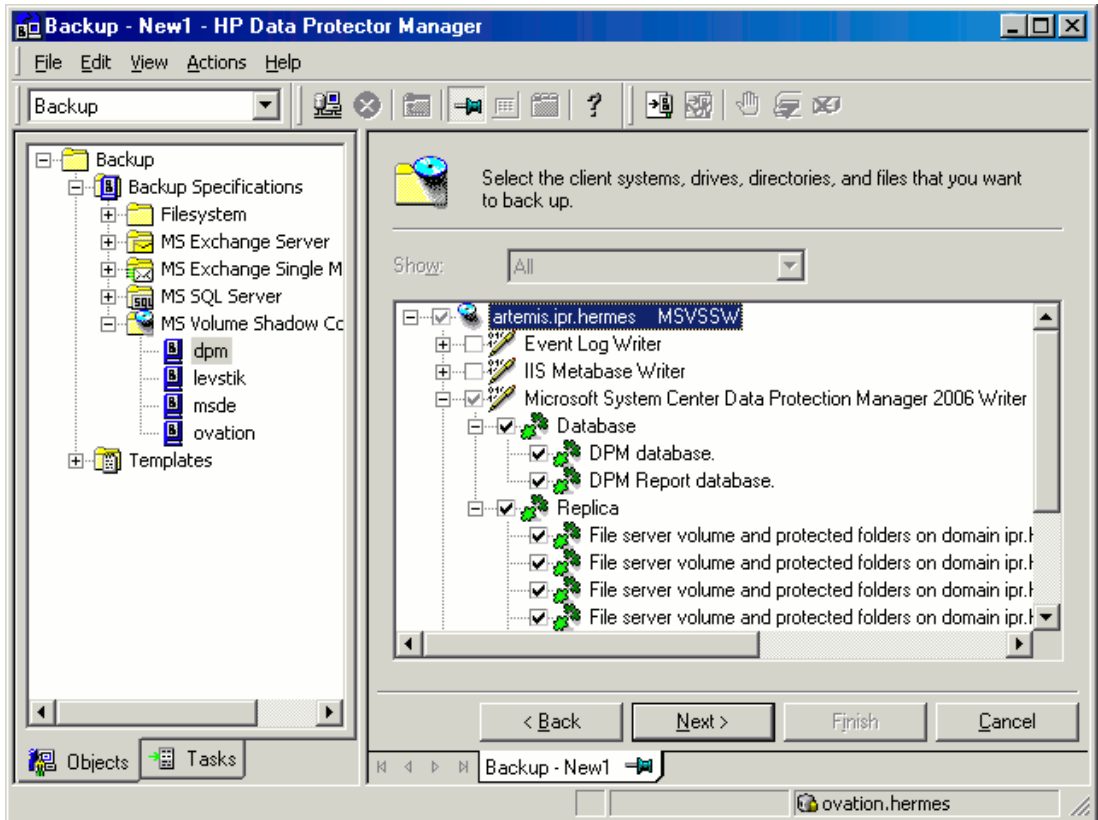


Figure 60 Selecting Microsoft Data Protection Manager database and replicas

Scheduling the backup

For more detailed information on scheduling, refer to the online Help index: “scheduled backups”.

To schedule a Microsoft Volume Shadow Copy Writers backup specification, perform the following steps in the Data Protector GUI:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup**, then **Backup Specifications**. Click **MS Volume Shadow Copy Writers**.

A list of available backup specifications is displayed in the Results Area.

3. Double-click the backup specification you want to schedule and click the **Schedule** tab to open the **Schedule** property page.
4. In the **Schedule** property page, select a date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
5. Specify **Recurring**, **Time options**, **Recurring options**, and **Session options**.

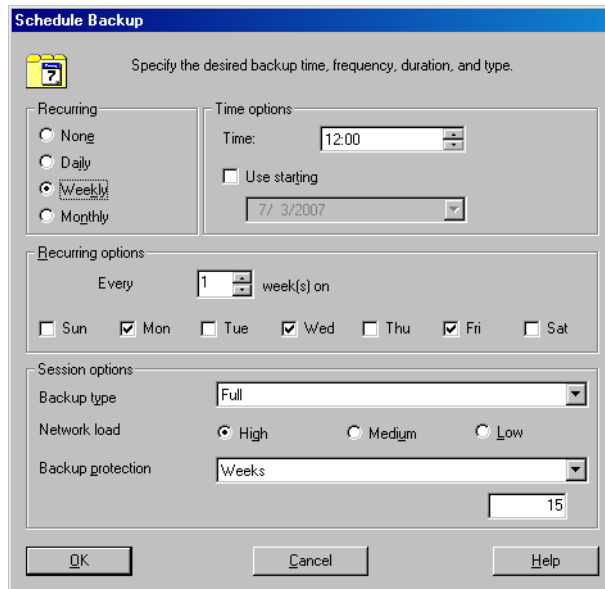


Figure 61 Scheduling a backup

6. Click **OK** to return to the **Schedule** property page.
7. Click **Apply** to save the changes.

Running an interactive backup

An interactive backup can be started using the Data Protector GUI by following these steps:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup**; then expand the **Backup Specifications** and the **MS Volume Shadow Copy Writers** items.

3. Right-click the backup specification you want to run, and then select **Start Backup** from the pop-up menu.

The **Start Backup** dialog box appears.

Select the backup type and the network load (**High**, **Medium**, or **Low**).

Refer to online Help for a description of network load.

4. Click **OK**. Upon successful completion of the backup session, a `Session Completed Successfully` message appears.

Restoring writers data

You can restore the Data Protector Microsoft Volume Shadow Copy Service integration backup objects using the Data Protector GUI.

NOTE:

Data Protector first restores the Writer Metadata collected during the backup time. This metadata contains the information about the backup components and the restore method. Data Protector performs restore according to the restore method specified by the writers.

Limitations for custom restore

- Data Protector Microsoft VSS integration does not automatically provide any restore method for writers requesting custom restore. If a writer specifies custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector file restore functionality. You can use the `Restore Into` option to specify an alternate restore path for these plain files. You can then perform the custom restore from these plain files manually. For information on writer's custom restore, refer to the writers documentation.

NOTE:

Writers requiring custom restore methods are by default not shown by Data Protector. The `omnirc` variable `OB2VSS_SHOWALLWRITERS` must be set to 1 for all writers to be displayed.

Restore procedure

The procedure below shows how to restore Microsoft VSS components using the Data Protector GUI. Some writers require custom restore procedures and/or have specific limitations. See also the appropriate sections:

- For Microsoft Exchange Server Writer specifics, see [“Microsoft Exchange Server writer restore specifics”](#) on page 179.
- For MSDE Writer specifics, see [“MSDE writer restore specifics”](#) on page 177.
- For Microsoft Data Protection Manager 2006 Writer specifics, see [“Microsoft Data Protection Manager 2006 writer restore specifics”](#) on page 184.

To restore Microsoft VSS objects using the Data Protector GUI, proceed as follows:

- 1.** In the **HP Data Protector**, switch to the **Restore** context.
- 2.** Expand **Restore Objects**, expand **MS Volume Shadow Copy Writers**, expand the client from which you want to restore the data, and then click **MS Volume Shadow Copy Writers**. In the Results Area, a list of writers, which were backed up on this client, is displayed.
- 3.** Select the Restore mode:
 - **Restore components**
With this option selected, whole components are restored using the Volume Shadow Copy Service. Individual files cannot be selected for restore.
 - **Restore files to temporary location**
With this option, you can select individual files or a group of files that were backed up using the selected writer. The files are restored using the Data Mover Agent and not the Volume Shadow Copy service.

4. In the Results Area, select the writers or writers' components (for component restore) or files or a group of files (for file restore mode).




Figure 62 Restore objects

You can select the top-level item (full writer restore) or only specific components. If you select a full writer restore, but some components of this writer were not backed up in the same session, the unavailable components are shaded and you cannot select them. To select the version (the date of a backup), right-click the object name and click **Properties**. The last backup version is selected by default, however, you can select a different version from the drop-down list.

Exchange Server Writer: Optionally, to specify options for the consistency check of a Microsoft Exchange writer, right-click the writer and click **Additional options**.

You can perform *point-in-time* or *rollforward* restore of Microsoft Exchange Server writer. For more information, see [Microsoft Exchange Server writer restore specifics](#).

Exchange Server 2007 Writer:

If you restore an LCR or CCR copy to the original location, note that restore will be performed to the original database (Exchange Information Store) and not to the database copy (Exchange Replication Service).

To perform *restore to a different location*:

- a. Right-click a storage group, a store, or logs and click **Restore as**.
- b. In the MS Exchange Additional options dialog box, select the target location for the components you want to restore: target server, target storage group, and target stores. The following options are possible:

- **Restore to a different store**

This option is selected by default. It enables you to select the target stores for each store to be restored (original stores). Select the target server first and then from the Original and Target drop-down list, select the stores pairs. For example, to restore only one store, select only one target store for the original store you want to restore. Note that only stores cannot be restored, so logs are automatically selected in the restore session to different location.

- **Restore to a non-Exchange location**

Select this option to restore your data to a non-Exchange location. This means that the restored data will not be managed by Exchange Server and Recovery Storage Group (RSG) will not be created. You can create RSG when data is restored. Select the target server first and then from the **Original** drop-down list, select the stores to be restored.

- **Restore to a non-Exchange location and create RSG**

Select this option to restore your data to a non-Exchange location. After restore, Data Protector will create a Recovery Storage Group called TBD on the target server. The stores that you have selected and logs will be restored to this recovery group. Select the target server first and then from the **Original** drop-down list, select the stores to be restored.

Optionally, select the new location for the storage group (**Restore into location**) if you want to be different than default, which is `C:\Omni`.

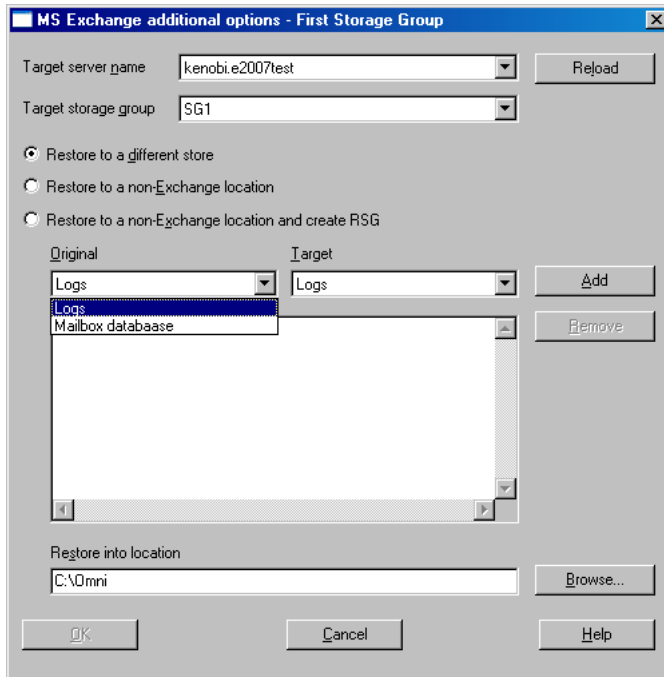


Figure 63 Restore to different location options (Exchange Server 2007 Writer)

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. Refer to “[Restore options](#)” on page 177.
6. In the **Devices** and **Media** property pages, the devices and media for restore are automatically selected.

Note that you can change the device used for the restore. Therefore, you have the possibility of using a different device for a restore than the one that was used for the backup. See the online Help index: “selecting devices for restore”.

7. Click **Restore**. Review your selection, and then click **Finish** to start a restore session.

The restore session messages are displayed in the Results Area.

8. If you are restoring a VSS writer that requires a custom restore, continue manually, using the writers specific methods, if it is provided by a writer. Refer to the writers’ documentation.

Restore options

The following restore options are specific to the Data Protector MS Volume Shadow Copy integration.

Restore to another client

By default, the components or files are restored to the client from which the application data was backed up. However, you may restore the data to another VSS client if you specify the **Restore to another client** option. The new target Microsoft VSS client must be a part of the Data Protector cell. For component restore, it must also run on the same platform and have the MS Volume Shadow Copy Integration software component installed. For file restore, the MS Volume Shadow Copy Integration software component is not required.

Restore into the following directory

By default, you restore the data to the same directory from which it was backed up (it can be on the original client or on some other client which you selected).

However, if you specify the **Restore into the following directory** option, your data will be restored to another directory. When defining the restore location, you can specify the path to the directory where you want to restore your data.

MSDE writer restore specifics

MSDE writer is used to back up and restore Microsoft SQL database.

 **IMPORTANT:**

Before restoring the SQL system databases (*master*, *model*, *msdb* and *pub*), you have to stop the SQL service.

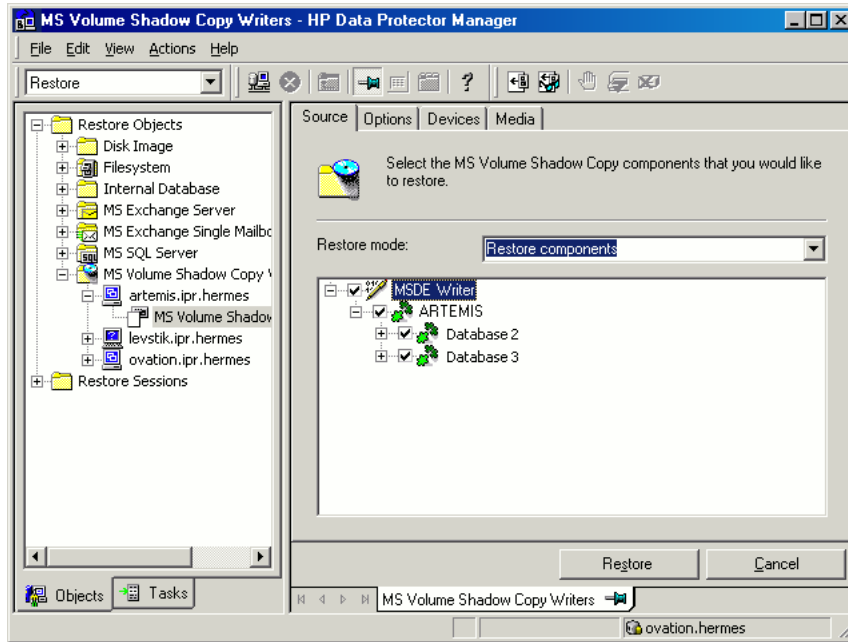


Figure 64 MSDE writer

When you expand the MSDE Writer item in the Results Area, all Microsoft SQL Server instances are displayed. Each instance contains all databases it includes. System databases (*master*, *model*, *msdb* and *pub*) are always listed there.

 **IMPORTANT:**

If system databases are restored, the whole internal database structure will be changed.

 **NOTE:**

Only point-in-time restore is possible. Rollforward restore is not supported.

User databases will be restored only if it is possible to overwrite the files. MSDE writer will take the user databases offline before the restore, while SQL service will have to be stopped manually in order to restore the system databases.

Microsoft Exchange Server writer restore specifics

Microsoft Exchange Server Writer is used to restore Microsoft Exchange Server database files.

When restoring from a Microsoft Exchange backup, the following two scenarios are possible:

- One or more databases are corrupted, but the log files are not damaged. In this case the database is restored and transaction logs are applied. See [“Rollforward recovery from the loss of one or more databases”](#) on page 179.
- The log files are corrupted or missing. In this case all databases and log files need to be restored. A rollforward recovery of the database is not possible. See [“Point-in-time restore after loss of a log file”](#) on page 180.

Microsoft Exchange Server 2003 limitations

- Shadow copies cannot be restored to alternate locations on the backup system.
- You cannot restore the shadow copy to the Recovery Storage Group.
- Rollforward recovery cannot be performed after a point-in-time restore.
- Restore to Microsoft Exchange Server 2007 from backup made with Microsoft Exchange Server 2003 is not supported.

Rollforward recovery from the loss of one or more databases

For a rollforward recovery:

1. Dismount all stores from the storage group in which the target store resides using Microsoft Exchange System Manager.

2. In the Data Protector GUI, switch to the **Restore** context. Expand **Restore Objects** and **MS Volume Shadow Copy Writers** and select the client from which you want to restore the data.

In the Results Area, expand Microsoft Exchange Writer and select the stores you want to recover. The **Logs** component is shaded and cannot be selected. You cannot select versions of individual stores, because a rollforward recovery is performed only to the current storage group state.

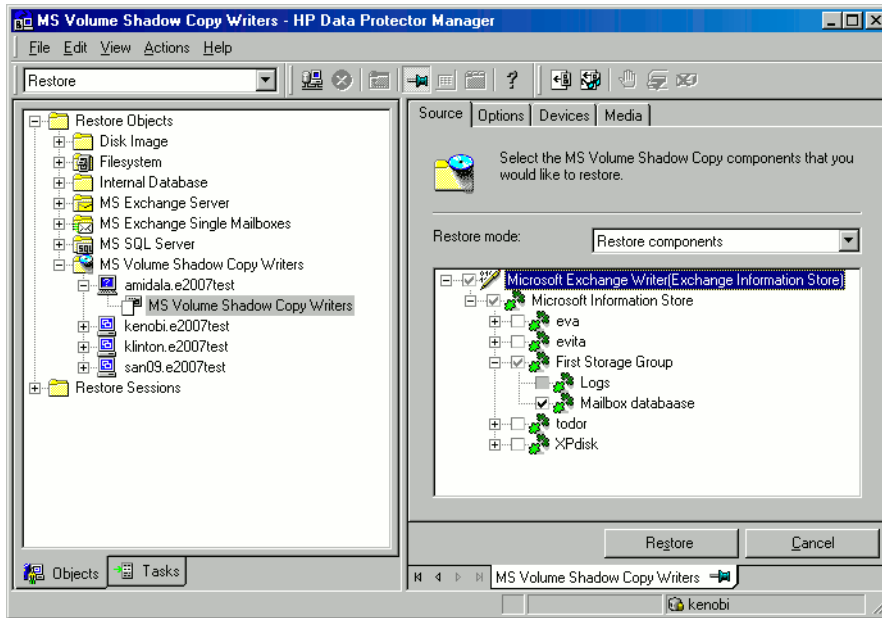


Figure 65 Selecting Microsoft Exchange Server 2007 stores for rollforward recovery

3. Proceed as with general VSS writer restore. See “[Restore procedure](#)” on page 172 for the general VSS writer restore procedure.
4. Mount all stores from the storage group in which they reside using Exchange System Manager. Selected stores are recovered.

Point-in-time restore after loss of a log file

To perform a point-in-time restore:

1. Start Exchange System Manager and check if the storage group is already unmounted. If not, unmount the whole group.

2. Switch to the **Restore** context. Expand **Restore Objects** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

In the Results Area, expand Microsoft Exchange Writer and select the whole storage group. Do not select individual stores.

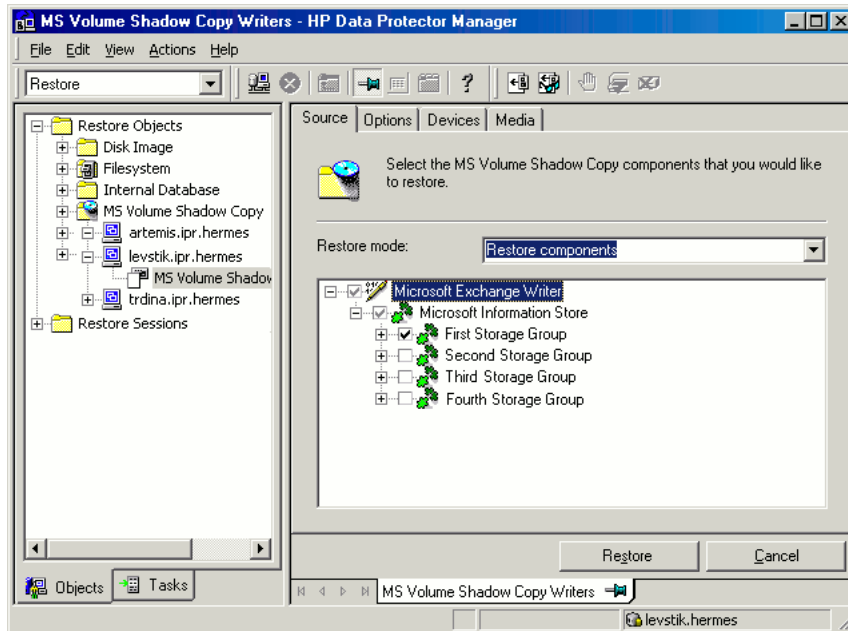


Figure 66 Selecting Microsoft Exchange Server 2003 stores for point-in-time restore

3. Proceed as with general VSS writer restore. See [“Restore procedure”](#) on page 172 for the general VSS writer restore procedure.
4. Mount the stores from the storage group in which the target stores reside using Exchange System Manager. All stores are mounted and put in the state as they were at the last selected full, incremental, or differential backup.

Microsoft Exchange Server 2007 writer restore specifics

Exchange Server 2007 offers additional restore option:

- Restore to a different location than the location from where the backup was made (original location).

The following restore scenarios are supported:

- Restore of a whole storage group or individual stores to the original location.
- Restore of LCR and CCR copies (their storage groups or individual stores) to the original database.
- Restore of a whole storage group (including LCR or CCR copies) or individual stores to a different location:

- *Restore to a different storage group*

Data can be restored to a **different storage group**. During the restore, data is restored into the target component destination path and all needed information is stored there. Exchange Server can then replay the restored logs from the original storage group on the target storage group to bring the restored database up to date. If you perform restore to a different storage group, you can access single mailboxes or individual mails on an alternate location without the need to compromise the original storage group with a potentially unsuccessful restore operation.

- *Restore to a non-Exchange location*

Data can be restored to a **non-Exchange location**. On session level, this option is already available if you use **Restore into the following directory** option. But on a storage group level, one storage group can be restored to a target directory, another to a different storage group in the same session if the target restore client system is the same. Also single store can be restored to a non-Exchange location. During the restore, data is restored into target component destination path.

Once the data is restored, it is not managed by Exchange Server.

- *Restore to a non-Exchange location and creating Recovery Storage Group (RSG)*

Data can be restored to a non-Exchange location and after that a **Recovery Storage Group** can be created for the databases. Exchange Server recognizes the Recovery Storage Group created by Data Protector, which is called `TBD`.

- *Restore to Recovery Server*

If a whole Exchange Server system is unusable, then the whole storage group can be restored to some other Exchange Server system (**Recovery Server**) and some other storage group. If the target system is identical to original (the same host name, storage groups names, stores names, and the same directory structure), restore can be started as normal restore to original location. If there are differences, then restore to a different location must be used.

To restore to a different location on a storage group or store level, right-click the storage group or store or logs to be restored and click **Restore as**. Select the target location for all the components you want to restore.

You can restore two storage groups to both original and a different location in the same restore session if the target restore client system is the same.

During restore to original location, all stores in the storage group to be restored are dismounted even if only one store from the storage group will be restored. Stores are left dismounted after the restore.

Prerequisites

- For restore to a different location, `Inet` must run under domain account that has the permissions assigned to the Exchange Server Administrators group. The account must also be a member of the local Administrators group on that computer.
- If you restore to different storage group and if there are any dismounted stores you want to use for the restore, mount them before the restore. Otherwise, Data Protector cannot list them for the target selection in the GUI.
If this is done while the Data Protector GUI is opened, you can refresh the current GUI page by clicking **Reload** in the **MS Exchange additional options** dialog box.

IMPORTANT:

If you restore to a non-Exchange location and create RSG, check that Recovery Storage Group does not already exist on this location. If Recovery Storage Group already exists, it will be deleted by Data Protector. All data stored in the existing Recovery Storage Group will be lost.

Limitations

- Restore to an LCR or CCR storage group copy (Exchange Replication Service instance) is not supported due to the Microsoft Exchange Server limitations. Restore of an LCR or CCR copy to original location will be performed to the original database (Exchange Information Store) and not to the database copy (Exchange Replication Service).
- When restoring to a different location, you cannot restore a specific database store without the transaction logs. When a store is selected for restore, logs are automatically included for the restore.
- Only one storage group can be restored with the option **Restore to a non-Exchange location and create RSG**. This is due to the Microsoft Exchange Server limitation that only one RSG per server is supported.
- You cannot mount public folders to RSG due to the Microsoft Exchange Server limitations.

- All stores in a RSG must originate from the same storage group due to the Microsoft Exchange Server limitations.

Microsoft Data Protection Manager 2006 writer restore specifics

When restoring the DPM writer, you can:

- Restore the DPM server first and then use the DPM to restore clients.
In case of a disaster, when the entire DPM server is lost, perform a standard disaster recovery procedure first and continue with restoring the DPM server. See [“Restore the DPM server first”](#) on page 184
- Restore individual DPM *clients* directly, without using the DPM server (for example, if you cannot restore the DPM server or if you want to avoid the additional step of recreating the DPM snapshot). When restoring the DPM clients directly you can select between component restore and file restore modes. See [“Restore the DPM clients directly”](#) on page 187.



NOTE:

Although the Data Protection Manager databases could also be restored using the MSDE writer, this method is not recommended, because DPM is *not* shut down automatically as with the DPM writer. If you really need to use this writer, shut down the DPM server manually.

Limitations

- Restore to another server is not supported by the Data Protection Manager writer.
- Parallel restore to different clients is not supported.

Restore the DPM server first

1. Start the DPM administrator console and add disks to the storage pool so that you have enough free space to restore the replicas.
Ensure that the DPM writer (service) is started.
2. Switch to the Data Protector **Restore** context. Expand **Restore** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

3. In the Results Area, expand the DPM writer and select *only* the Data Protection Manager databases.

Proceed as with general VSS writer restore. See “[Restore procedure](#)” on page 172 for the general VSS writer restore procedure.

4. Run the DPM command `DpmSync -Sync` to reallocate replicas.

5. Switch back to the Data Protector **Restore** context, and select and restore the necessary *replicas*.

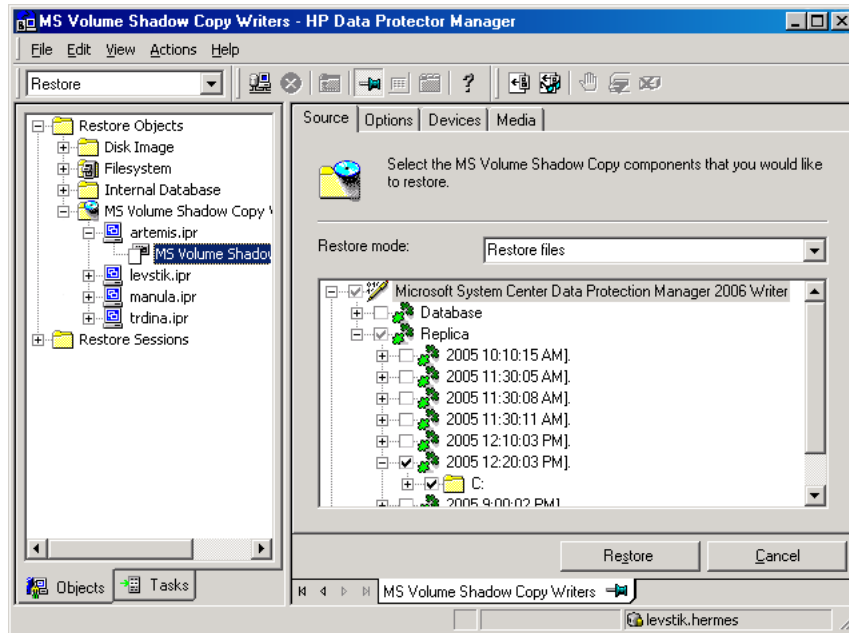


Figure 67 Restoring the Microsoft Data Protection Manager 2006 client

6. Use the DPM to restore individual clients.

 **IMPORTANT:**

The DPM console does not automatically check for new or restored snapshots. Before you can start the restore of clients, you must use the Data Protection Manager to recreate a DPM snapshot.

- a. In the DPM console, open the **Recovery** context. Under the **Browse** tab, select the server, right click on the restored replica, and select **Create shadow copy now**.
- b. Select and restore the new snapshot to the client.

Restore the DPM clients directly

1. Switch to the **Restore** context. Expand **Restore** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.
2. Select the restore modes:
 - **Restore Components**
Use this mode *only* if the client to which you want to restore supports VSS, for example if you restore to Windows 2003 clients.
You can restore only entire replicas.
 - **Restore Files**
The client does not need to support VSS and you can restore individual folders or files.
3. When selecting the DPM writer for restore, select *only* the **Replica** components. Do not select the DPM database.
4. Click the **Options** tab, and under **Restore to another client** enter the name of the target client. Click **Next**.
5. Proceed as with general VSS writer restore. See “[Restore procedure](#)” on page 172 for the general VSS writer restore procedure.

Monitoring a VSS backup and restore

The Data Protector GUI enables you to monitor current or view previous backup and restore sessions.

Monitoring is automatically activated when you start a restore or a backup interactively.

Monitoring current sessions

To monitor a currently running session using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Monitor**.
In the Results Area, all currently running sessions are listed. See [Figure 68](#) on page 188.
2. Double-click the session you want to monitor.




Figure 68 Monitoring a current session

Clearing sessions

To remove all completed or aborted sessions from the Results Area of the **Monitor** context, proceed as follows:

1. In the Scoping Pane, click **Current Sessions**.
2. In the **Actions** menu, select **Clear Sessions**. Or click the **Clear Sessions** icon on the toolbar.

To remove a particular completed or aborted session from the current sessions list, right-click the session and select **Remove From List**.

NOTE:

All completed or aborted sessions are automatically removed from the Results Area of the **Monitor** context if you restart the Data Protector GUI.

For detailed information on a completed or aborted session, see “[Viewing previous sessions](#)” on page 188.

Viewing previous sessions

To view a previous session using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand **Sessions** to display all the sessions stored in the IDB.

The sessions are sorted by date. Each session is identified by a session ID consisting of a date in the YY/MM/DD format and a unique number.

3. Right-click the session and select **Properties** to view details on the session.

4. Click the **General**, **Messages** or **Media** tab to display general information on the session, session messages, or information on the media used for this session, respectively. See [Figure 69](#) on page 189.

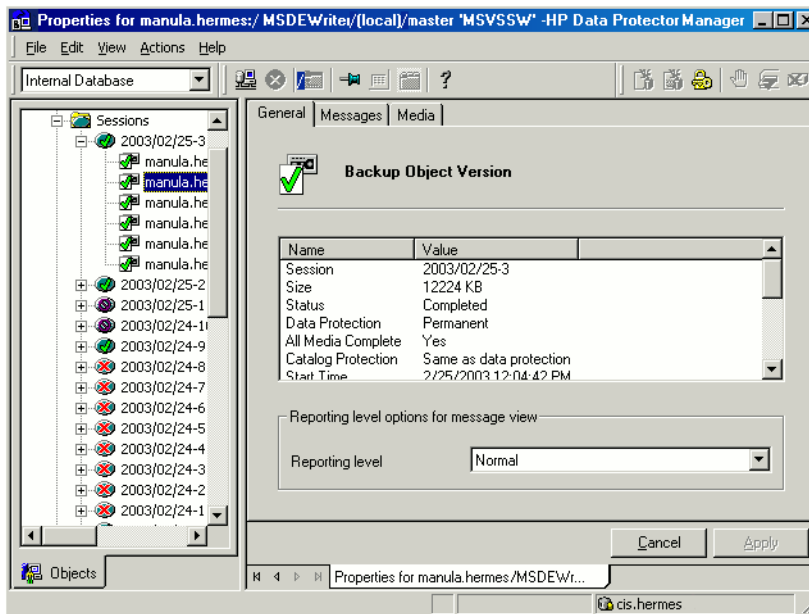


Figure 69 Viewing a previous session

Troubleshooting

This section lists problems you might encounter when using the Data Protector Microsoft Volume Shadow Copy integration.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.

- See <http://www.hp.com/support/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

- On the application and backup systems, examine system errors reported in:
HP-UX, Solaris: `/var/opt/omni/log/debug.log`
Windows: `Data_Protector_home\log\debug.log`

Backup problems

Problem

Microsoft Exchange Server 2003 writer backup failed.

When you start a backup of a Microsoft Exchange Server 2003 Writer, the following error is displayed:

```
[MAJOR]: Writer 'Microsoft Exchange Writer' failed to prepare files  
for backup.
```

The reason may be the failure of the previous backup and, due to a Microsoft Exchange Server 2003 issue, the Exchange Server Writer could not perform a proper cleanup.

Action

Restart the Information Store.

For more information, see the web page <http://support.microsoft.com/kb/945424/en-us>.

Problem

Microsoft Exchange Server 2003 aborts the backup if the shadow copy creation takes over 20 seconds.

If an Exchange Server 2003 Writer is being backed up, the session can fail with VSSBAR reporting:

```
Snapshot could not be created.
```

In the application event log on the application system, the following event is recorded:

Event Type: Error
Event Source: ESE Event
Category: (16)
Event ID: 2004
Information Store (4916) Shadow copy 3 time-out (20000 ms).

Action

The following can help to solve the problem:

- Limit the number of users that are accessing the management system.
- Reduce the number of volumes in a snapshot set. Create a backup specification dedicated to each storage group instead of one specification for the whole server. See "[Microsoft Exchange Server writer backup specifics](#)" on page 163.

TBD: this whole problem should be checked if applicable for DP ZDB

Problem

Backup of a Microsoft Exchange Server 2007 CCR database copy fails.

A backup session of a database copy in a CCR environment reports a major error that the backup session has failed.

A backup session may fail in a CCR environment when you want to back up a database copy, which has "Healthy" status in the Exchange Management Console, but in reality the status of the database is "Failed". This is the Microsoft Exchange Server 2007 issue.

Action

To repair the Exchange Management Console to show the real status of the database copy, perform one of the following two actions:

- 1.** Install the Microsoft Exchange Server 2007 Service Pack 1.
Or:
- 2.** Perform reseeding procedure as described below:
 - a.** On the passive node, where the copy of the database exists, suspend the replication by using `Suspend-StorageGroupCopy` cmdlet.
 - b.** Delete all log files from the `Logs` directory of the database copy.
 - c.** Seed the database copy or re-synchronize the original and copy databases by using the `Update-StorageGroupCopy` cmdlet.
 - d.** Resume the database copy by using the `Resume-StorageGroupCopy` cmdlet.

- e. Check the status of the Exchange Replication Service by using `vssadmin list writers` command on the passive node. If the status is not stable, restart the Exchange Replication Service.

Restore problems

Problem

After the restore of system writers was aborted, the Windows operating system is corrupted when you restart it.

If the restore of some system writers (for example, System Writer) is aborted for any reason (hardware or software failure, manually aborted, etc.), the Windows operating system may be corrupted after the restart (for example, the GUI or some system services cannot be started, etc.).

Action

Depending on the nature of the corruption, repair or re-install the operating system from the Windows installation CD-ROM.

Problem

Some components are not restored during the restore session.

If a component cannot be restored to the location specified in Writer Metadata Document (for example, if this location is locked or it is not possible to perform regular restore), this component will be skipped during the restore procedure.

Action

Specify a location, where skipped files will be redirected in case of failure, by setting the `OB2VSS_DUMPTO` environmental variable in the `Data_Protector_home\omnirc` file. Restart the Data Protector services to apply the changes in the `omnirc` file.

Example

If you want the files that are skipped during the restore to be copied to the `F:\Restore` directory, set `OB2VSS_DUMPTO=F:\Restore` in the `omnirc` file. In case the SQL component `Company` was skipped during the restore, it will be copied to the specified directory as follows:

```
F:\Restore\2002-12-09-23\G\SQL\Log\Company.ldf
F:\Restore\2002-12-09-23\G\SQL\Log\Company.mdf
```


The pathname includes the backup session ID and the pathname to the original location.

Glossary

access rights	See user rights .
ACSL	<i>(StorageTek specific term)</i> The Automated Cartridge System Library Server (ACSL) software that manages the Automated Cartridge System (ACS).
Active Directory	<i>(Windows specific term)</i> The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.
AML	<i>(EMASS/GRAU specific term)</i> Automated Mixed-Media library.
application agent	A component needed on a client to back up or restore online database integrations. See also Disk Agent .
application system	<i>(ZDB specific term)</i> A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume .
archived redo log	<i>(Oracle specific term)</i> Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: <ul style="list-style-type: none">• 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](#).

archive logging	<i>(Lotus Domino Server specific term)</i> Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.
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 <code>Data_Protector_home\Config\Server\dr\asr</code> on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/asr/</code> on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes 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	<i>(Oracle specific term)</i> Automatic Storage Management is an Oracle 10g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.
automigration	<i>(VLS specific term)</i> 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](#).

BACKINT	<p>(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.</p>
backup API	<p>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.</p>
backup chain	<p>See restore chain.</p>
backup device	<p>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.</p>
backup generation	<p>One backup generation includes one full backup and all incremental backups until the next full backup.</p>
backup ID	<p>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.</p>
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). 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 (full or incremental). 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 , incremental backup , and full backup .
backup set	A complete set of integration objects associated with a backup.
backup set	<i>(Oracle specific term)</i> A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.
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	<i>(ZDB specific term)</i> A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica. See also application system , target volume , and replica .
backup types	See incremental backup , differential backup , transaction backup , full backup , and delta backup .

backup to IAP	A Data Protector based backup to the HP Integrated Archiving Platform (IAP) appliance. It takes advantage of the IAP capability to eliminate redundancies in the stored data at a block (or chunk) level, by creating a unique content address for each data chunk. Only changed chunks are transmitted over the network and added to the store.
backup view	Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.
BC	<i>(EMC Symmetrix specific term)</i> Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices. See also BCV .
BC	<i>(HP StorageWorks Disk Array XP specific term)</i> The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets should be connected to the backup system. See also HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system .
BC EVA	<i>(HP StorageWorks EVA specific term)</i> Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware. See also replica , source volume , snapshot , and CA+BC EVA .
BC Process	<i>(EMC Symmetrix specific term)</i> A protected storage environment solution that has defined specially configured EMC Symmetrix

devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also [BCV](#).

BC VA *(HP StorageWorks Virtual Array specific term)* Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system. See also [HP StorageWorks Virtual Array LUN](#), [application system](#), and [backup system](#).

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected. See also [BC](#) and [BC Process](#).

Boolean operators The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. See also [BRBACKUP](#), and [BRRESTORE](#).

BRBACKUP *(SAP R/3 specific term)* An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data

files, or of all tablespaces and, if necessary, of the online redo log files.

See also [BRARCHIVE](#), and [BRRESTORE](#).

BRRESTORE

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

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

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.

See also [BRBACKUP](#), and [BRARCHIVE](#).

BSM

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

CA

(HP StorageWorks Disk Array XP specific term) Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.

See also [BC](#) *(HP StorageWorks Disk Array XP specific term)*, [Main Control Unit](#) and [HP StorageWorks Disk Array XP LDEV](#).

CA+BC EVA

(HP StorageWorks EVA specific term) The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.

See also [BC EVA](#), [replica](#), and [source volume](#).

CAP

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

catalog protection	Defines how long information about backed up data (such as file names and file versions) is kept in the IDB. See also data protection .
CDB	The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions,, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell. See also MMDB .
CDF file	<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. Central control is available to administer the backup and restore policies and tasks.
Cell Manager	The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.
centralized licensing	Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs. See also MoM .
Centralized Media Management Database (CMMDB)	See CMMDB .

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	<p><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:</p> <ul style="list-style-type: none"> • type 'disk' • type 'sbt_tape' <p>If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.</p>
chunking	<p><i>(IAP specific term)</i> The process of dividing data into blocks (chunks), where each chunk gets a unique content address. This address is then used to determine whether a particular chunk is already backed up to the IAP appliance. If the duplicate data is identified (two addresses are identical, that is the address is the same as for another data chunk already stored into IAP), it is not backed up. This way, the data redundancy is eliminated and the optimal data storage is achieved.</p> <p>See also backup to IAP.</p>
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 filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.
client backup with disk discovery	A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration

and improves backup coverage of systems that often mount or dismount disks.

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

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).

Cluster Continuous Replication (*Microsoft Exchange Server specific term*) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also [Exchange Replication Service](#) and [Local Continuous Replication](#).

CMD Script for Informix Server (*Informix Server specific term*) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also [MoM](#).

COM+ Registration Database	<i>(Windows specific term)</i> The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.
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.
Command View (CV) EVA	<i>(HP StorageWorks EVA specific term)</i> The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView Storage Management Appliance, and is accessed by a Web browser. See also HP StorageWorks EVA SMI-S Agent and HP StorageWorks SMI-S EVA provider .
Command View VLS	<i>(VLS specific term)</i> A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN. See also Virtual Library System (VLS) .
concurrency	See Disk Agent concurrency .
control file	<i>(Oracle and SAP R/3 specific term)</i> An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.
copy set	<i>(HP StorageWorks EVA specific term)</i> A pair that consists of the source volumes on a local EVA and their replica on a remote EVA. See also source volume , replica , and CA+BC EVA
CRS	The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account <code>root</code> .

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.
data file	<i>(Oracle and SAP R/3 specific term)</i> A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.
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. <i>See also catalog protection.</i>
data stream	Sequence of data transferred over the communication channel.
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.
Data Replication (DR) group	<i>(HP StorageWorks EVA specific term)</i> A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common CA EVA log. <i>See also copy set.</i>
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	<i>(Informix Server specific term)</i> An Informix Server physical database object. It can be a blob space, db space, 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 <code>dcbf</code> directory and is located in the <code>Data_Protector_home\db40</code> directory on a Windows Cell Manager and in the <code>/var/opt/omni/server/db40</code> directory on a UNIX Cell

Manager. You can create more DC directories and use a custom location. Up to 10 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 file system 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	<i>(EMC Symmetrix specific term)</i> A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.
device streaming	A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.
DHCP server	A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup	An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. See also incremental backup .
differential backup	<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. See also backup types .
differential database backup	A differential database backup records only those data changes made to the database after the last full database backup.
direct backup	A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems. See also XCOPY engine .
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.
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.
Disk Agent concurrency	The number of Disk Agents that are allowed to send data to one Media Agent concurrently.
disk discovery	The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have

been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

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. See also virtual full backup .
Distributed File System (DFS)	A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.
DMZ	The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server	In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.
domain controller	A server in a network that is responsible for user security and verifying passwords within a group of other servers.
DR image	Data required for temporary disaster recovery operating system (DR OS) installation and configuration.
DR OS	A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also 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.
dynamic client	See client backup with disk discovery .
EMC Symmetrix Agent (SYMA) (EMC Symmetrix specific term)	See Symmetrix Agent (SYMA) .
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.

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 Event Log is accessible only to Data Protector users in the <code>Admin</code> group and to Data Protector users who are granted the <code>Reporting and notifications</code> user rights. You can view or delete all events in the Event Log.
Event Logs	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. <i>See also Cluster Continuous Replication and Local Continuous Replication.</i>
exchanger	Also referred to as SCSI Exchanger. <i>See also library.</i>
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. <i>See also importing media.</i>

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.
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 StorageWorks EVA specific term)</i> An operation that reverses the roles of source and destination in CA+BC EVA configurations. See also CA+BC 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 StorageWorks Disk Array XP specific term)</i> HP StorageWorks Disk Array XP allows up to three mirror copies of a primary volume and each of these copies can have additional two copies. The three mirror copies are called first-level mirrors. See also primary volume and MU number .
flash recovery area	<i>(Oracle specific term)</i> Flash recovery area is an Oracle 10g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and 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 .
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 in the <code>/etc/opt/omni/server/options</code> directory on HP-UX and Solaris systems and in the <code>Data_Protector_home\Config\Server\Options</code> directory on Windows 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 cross-platform (HP-UX, Solaris, Linux, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI, Data Protector also provides a Java-based graphical user interface with the same look and feel. As Java can run on numerous platforms, the Data Protector Java GUI is supported on a larger number of platforms than the original Data Protector GUI.
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: <code>/etc/opt/omni/server/Holidays</code> on the UNIX Cell Manager and <code>Data_Protector_home\Config\Server\holidays</code> on the Windows Cell Manager.
host backup	See client backup with disk discovery .
hosting system	A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.
HP ITO	See OM .
HP OpC	See OM .
HP OpenView SMART Plug-In (SPI)	A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager software, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager software (OM).
HP OM	See OM
HP StorageWorks Disk Array XP LDEV	A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities. See also BC , CA (<i>HP StorageWorks Disk Array XP specific term</i>), and replica .
HP StorageWorks EVA SMI-S Agent	A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.

See also [Command View \(CV\) EVA](#) and [HP StorageWorks SMI-S EVA provider](#).

- HP StorageWorks SMI-S EVA provider** An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.
See also [HP StorageWorks EVA SMI-S Agent](#) and [Command View \(CV\) EVA](#).
- HP StorageWorks Virtual Array LUN** A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also [BC VA](#) and [replica](#).
- HP VPO** See [OM](#).
- 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 is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.
- 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. <i>See also exporting media.</i>
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. <i>See also backup types.</i>
incremental backup	<i>(Microsoft Exchange Server specific term)</i> A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up. <i>See also backup types.</i>
incremental mailbox backup	An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.
incremental1 mailbox backup	An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.
incremental (re)-establish	<i>(EMC Symmetrix specific term)</i> A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
incremental restore	<i>(EMC Symmetrix specific term)</i> A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and

the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

- 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](#).
- Inet** A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.
- Information Store** (*Microsoft Exchange Server specific term*) The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also [Key Management Service](#) and [Site Replication Service](#).
- Informix Server** (*Informix Server specific term*) Refers to Informix Dynamic Server.
- initializing** See [formatting](#).
- Installation Server** A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.
- instant recovery** (*ZDB specific term*) A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also [replica](#), [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), and [ZDB to disk+tape](#).

integration object	A backup object of a Data Protector integration, such as Oracle or SAP DB.
Internet Information Services (IIS)	<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).
IP address	An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).
ISQL	<i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.
ITO	See OM .
Java GUI Client	The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities 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".
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.

LBO *(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation** A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

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

load balancing By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

Local Continuous Replication	<p>(<i>Microsoft Exchange Server specific term</i>) Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group. See also Cluster Continuous Replication and Exchange Replication Service.</p>
lock name	<p>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.</p>
log_full shell script	<p>(<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 <code>ALARMPROGRAM</code> configuration parameter defaults to the <code>INFORMIXDIR/etc/log_full.sh</code>, where <code>INFORMIXDIR</code> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the <code>ALARMPROGRAM</code> configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code>.</p>
logging level	<p>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</p>

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

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID

(Microsoft SQL Server specific term) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database

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

- *user_name* is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- *password* must be the same as the password specified in the Oracle password file (*orapwd*), which is used for authentication of users performing database administration.
- *service* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *user_name/password@service*, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API

(Lotus Domino Server specific term) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX

systems. An LVM system consists of several volume groups, where each volume group has several volumes.

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 StorageWorks Disk Array XP specific term)</i> An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. See also BC (HP StorageWorks Disk Array XP specific term), CA (HP StorageWorks Disk Array XP specific term), and HP StorageWorks Disk Array XP LDEV .
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 session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. 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.
medium ID	A unique identifier assigned to a medium by Data Protector.
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 <code>Appendable</code> , <code>Non-Appendable</code> , or <code>Appendable for incrementals only</code> .
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. <i>See also overwrite.</i>

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 StorageWorks Disk Array XP specific term)	See target volume .
mirror rotation (HP StorageWorks Disk Array XP specific term)	See replica set rotation .
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](#), [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 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.
mount point	The access point in a directory structure for a disk or logical volume, for example <code>/opt</code> or <code>d:</code> . On UNIX, the mount points are displayed using the <code>bdf</code> or <code>df</code> command.
MSM	The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.
MU number	<i>(HP StorageWorks Disk Array XP specific term)</i> Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror. See also first-level mirror .
multi-drive server	A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.
obdrindex.dat	See IDB recovery file .
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.
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.
offline backup	<p>A backup during which an application database cannot be used by the application.</p> <ul style="list-style-type: none">• For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.• For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process. <p>See also zero downtime backup (ZDB) and online backup.</p>

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	<p><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:</p> <ul style="list-style-type: none"> • the <code>onbar</code> 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	<p><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 <code>onconfig</code> file in the directory <code>INFORMIXDIR\etc</code> (on Windows) or <code>INFORMIXDIR/etc/</code> (on UNIX).</p>
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 that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.</p> <ul style="list-style-type: none"> • For simple backup methods (non ZDB), backup mode is required for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished. • For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process. <p>In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. See also zero downtime backup (ZDB), and offline backup.</p>

online redo log	<i>(Oracle specific term)</i> Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused. See also archived redo log .
OpC	See OM .
OpenSSH	A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.
Oracle Data Guard	<i>(Oracle specific term)</i> Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.
Oracle instance	<i>(Oracle specific term)</i> Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.
ORACLE_SID	<i>(Oracle specific term)</i> A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <code>ORACLE_SID</code> . The <code>ORACLE_SID</code> is included in the <code>CONNECT DATA</code> parts of the connect descriptor in a <code>TNSNAMES.ORA</code> file and in the definition of the TNS listener in the <code>LISTENER.ORA</code> file.
original system	The system configuration backed up by Data Protector before a computer disaster hits the system.
overwrite	An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files. See also merging .

OM	<p>HP Operations Manager software for UNIX provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OM management servers on HP-UX, Solaris, and Linux. Earlier versions of OM were called IT/Operation, Operations Center and Vantage Point Operations.</p> <p>See also merging.</p>
ownership	<p>The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: <code>root.sys@Cell Manager</code>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.</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 backup medium and on Cell Manager into <code>Data_Protector_home\Config\Server\dr\pls</code> directory on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/pls</code> directory on a UNIX Cell Manager with the filename <code>recovery.pls</code>.</p>
package	<p><i>(MC/ServiceGuard and Veritas Cluster specific term)</i> 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><i>(HP StorageWorks Disk Array XP specific term)</i> A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:</p>

- COPY - The mirrored pair is currently re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

parallel restore	Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.
parallelism	The concept of reading multiple data streams from an online database.
physical device	A physical unit that contains either a drive or a more complex unit such as a library.
post-exec	A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX. See also pre-exec .
pre- and post-exec commands	Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
prealloc list	A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec	A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX. See also post-exec .
primary volume (P-VOL)	<i>(HP StorageWorks Disk Array XP specific term)</i> Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the 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 <code>.edb</code> file and a streaming native internet content <code>.stm</code> 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
RAID	Redundant Array of Inexpensive Disks.
RAID Manager Library	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.
RAID Manager XP	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This

	instance translates the commands into a sequence of low level SCSI commands.
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	<p><i>(Oracle specific term)</i> A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:</p> <ul style="list-style-type: none"> • 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	<p><i>(Oracle specific term)</i> Recovery files are Oracle 10g 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.</p>
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.

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.
recycle	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 StorageWorks Disk Array XP specific term)</i> The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can 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.

See also [snapshot](#), [snapshot creation](#), [split mirror](#), and [split mirror creation](#).

- replica set** *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also [replica](#) and [replica set rotation](#).
- replica set rotation** *(ZDB specific term)* The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also [replica](#) and [replica set](#).
- restore chain** All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.
- restore session** A process that copies data from backup media to a client.
- resync mode** *(HP StorageWorks Disk Array XP VSS provider specific term)* One of two XP VSS hardware provider operation modes. When the XP provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.
See also [VSS compliant mode](#), [source volume](#), [primary volume \(P-VOL\)](#), [replica](#), [secondary volume \(S-VOL\)](#), [MU number](#), and [replica set rotation](#).
- RMAN (Oracle specific term)** See [Recovery Manager](#).
- RSM** The Data Protector Restore Session Manager controls the restore session. 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.
scan	A function that 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).
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 StorageWorks Disk Array XP specific term)</i> secondary volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. 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, 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 <code>omnimnt</code> , <code>omnistat</code> , and <code>omniabort</code> 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.
single instancing	<i>(IAP specific term)</i> The process of recognizing redundancy of data, at both a whole object and a chunk level. It computes a strong hash for each data chunk and uses it as a unique content address needed to determine whether attempts to store duplicates are being made. See also backup to IAP .
Site Replication Service	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server 2000/2003 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.

- SMB** See [split mirror backup](#).
- smart copy** (*VLS specific term*) A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management. See also [Virtual Library System \(VLS\)](#).
- smart copy pool** (*VLS specific term*) A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library. See also [Virtual Library System \(VLS\)](#) and [smart copy](#).
- SMBF** The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, and media management sessions. One binary file is created per session. The files are grouped by year and month.
- snapshot** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation. See also [replica](#) and [snapshot creation](#).
- snapshot backup**
(*HP StorageWorks VA and HP StorageWorks EVA specific term*) See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- snapshot creation** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for

use. However background copying processes normally continue after creation.

See also [snapshot](#).

- source (R1) device** *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also [target \(R2\) device](#).
- source volume** *(ZDB specific term)* A storage volume containing data to be replicated.
- sparse file** A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.
- split mirror** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents 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 StorageWorks Disk Array XP specific term)* See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- split mirror creation** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also [split mirror](#).

split mirror restore	<i>(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)</i> A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. 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	<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	The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.
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	<i>(HP StorageWorks Disk Array XP specific term)</i> A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).
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.
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.
storage volume	<i>(ZDB specific term)</i> A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, 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.
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.
Symmetrix Agent (SYMA)	<i>(EMC Symmetrix specific term)</i> The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

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	<i>(Sybase specific term)</i> The four system databases on a newly installed Sybase SQL Server are the: <ul style="list-style-type: none"> • master database (master) • temporary database (tempdb) • system procedure database (sybprocedure) • model database (model).
System State	<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.
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	<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.

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 database	<i>(Oracle specific term)</i> In RMAN, the target database is the database that you are backing up or restoring.
target (R2) device	<i>(EMC Symmetrix specific term)</i> An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. See also source (R1) device .
target system	<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 crashed system and a target system is that a target system has all faulty hardware replaced.
target volume	<i>(ZDB specific term)</i> A storage volume to which data is replicated.
Terminal Services	<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.
thread	<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.
TimeFinder	<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).
TLU	Tape Library Unit.

TNSNAMES.ORA	<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.
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 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.
transaction log table	<i>(Sybase specific term)</i> A system table in which all changes to the database are automatically recorded.
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.

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 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.
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 StorageWorks EVA specific term)</i> The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers. See also Command View (CV) EVA .
Virtual Device Interface	<i>(Microsoft SQL Server specific term)</i> This is a SQL Server programming interface that allows fast backup and restore of large databases.
virtual disk	<i>(HP StorageWorks EVA specific term)</i> A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality. 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 .

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) .
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 mount point	<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 .
VPO	See OM .
VSS	See Microsoft Volume Shadow Copy Service .
VSS compliant mode	<i>(HP StorageWorks Disk Array XP VSS provider specific term)</i> One of two XP VSS hardware provider operation modes. When the XP 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 of the disks. See also resync mode , source volume , primary volume (P-VOL) , replica , secondary volume (S-VOL) , and replica set rotation .
VxFS	Veritas Journal Filesystem.
VxVM (Veritas Volume Manager)	A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

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.
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).
XCopy engine	<i>(direct backup specific term)</i> A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases

the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also [direct backup](#).

ZDB

See [zero downtime backup \(ZDB\)](#).

ZDB database

(ZDB specific term) A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.

See also [zero downtime backup \(ZDB\)](#).

ZDB to disk

(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.

See also [zero downtime backup \(ZDB\)](#), [ZDB to tape](#), [ZDB to disk+tape](#), [instant recovery](#), and [replica set rotation](#).

ZDB to disk+tape

(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, 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. On split mirror arrays, split mirror restore can also be used.

See also [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), [instant recovery](#), [ZDB to disk+tape](#), 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

- advantages, 147
 - VSS integration, 146, 147
- architecture
 - MS Exchange Server integration, 95
 - MS Exchange Single Mailbox integration, 116
 - MS SPS integration, 66
 - MS SQL Server integration, 29
 - VSS integration, 145, 149
- audience, 15

B

- backing up MS SPS
 - previewing backups, 78
 - starting backups, 79
- backing up MS SQL Server
 - backup specifications, creating, 38
 - concepts, parallelism, 29
- backing up MS Exchange Server
 - backup options, 103
- backing up MS SQL Server
 - backup options, 41
 - scheduling backups, 43
- backing up VSS
 - backup flow, 148
- backing up Informix
 - full backups, 27, 93
 - incremental backups, 27, 93
- backing up MS Exchange Server
 - backup specifications, creating, 98
- backing up MS Exchange Single Mailbox, 120
 - backup options, 123
 - backup types, 115
 - backup specifications, creating, 120
 - backup specifications, modifying, 124
 - full backups, 115
 - incremental backups, 115
 - performance tuning, 136
 - previewing backups, 125
 - scheduling backups, 124
 - scheduling backups, example, 124
 - starting backups, 126
 - starting backups, example, 127
- backing up MS SPS, 73 - 80
 - backup options, 77
 - backup specifications, creating, 74
 - backup types, 74
 - backup flow, 69
 - backup specification, modifying, 77
 - differential backups, 74
 - full backups, 74
 - scheduling backups, 77
 - scheduling backups, example, 77
 - transaction logs backups, 74
- backing up MS SQL Server, 38, 44
- backing up VSS, 159 - 171
 - architecture, 145, 149
 - backup specifications, creating, 160
 - scheduling backups, 169
 - starting backups, 170

- backup
 - creating backup specifications, 160
 - problems, 190
 - scheduling, 169
 - troubleshooting, 190
- backup options
 - MS SPS integration, 77
- backup specifications, scheduling
 - MS SPS integration, 77
 - VSS integration, 169
- backup flow, 148
 - VSS integration, 148
 - MS SPS integration, 69
- backup options
 - MS Exchange Server integration, 103
 - MS SQL Server integration, 41
 - MS Exchange Single Mailbox integration, 123
- backup specifications, creating
 - VSS integration, 160
- backup specifications, creating
 - MS Exchange Server integration, 98
 - MS Exchange Single Mailbox integration, 120
 - MS SPS integration, 74
 - MS SQL Server integration, 38
- backup specifications, modifying
 - MS Exchange Single Mailbox integration, 124
 - MS SPS integration, 77
- backup specifications, scheduling
 - MS Exchange Single Mailbox integration, 124
 - MS SQL Server integration, 43
- backup types
 - MS Exchange Single Mailbox integration, 115
 - MS SPS integration, 74

C

- checking configuration
 - MS Exchange Single Mailbox integration, 119
 - MS SPS integration, 73
 - MS SQL Server integration, 36
- checks and verifications
 - EVA troubleshooting, 190
- Cluster Continuous Replication, 151
- concepts
 - MS Exchange Server integration, 95
 - MS SPS integration, 66
 - MS Exchange Server integration, 94
 - MS Exchange Single Mailbox integration, 116
 - MS SQL Server integration, 28, 30
 - VSS integration, 147
 - VSS, Microsoft Exchange Server 2007 writer, 151
- configuration files
 - MS SQL Server integration, 31
- configuring
 - Microsoft Exchange Server writer, 155
- configuring MS SQL Server
 - checking configuration, 36
- configuring MS Exchange Server, 95 - 97
- configuring MS Exchange Single Mailbox, 117 - 120
 - checking configuration, 119
- configuring MS SPS, 69 - 73
 - checking configuration, 73
- configuring MS SQL Server, 30, 37
 - configuration files, 31
- configuring VSS, 153
- conventions
 - document, 22
- creating
 - backup specifications, 160

- creating backup specifications
 - MS Exchange Server integration, 98
 - MS Exchange Single Mailbox integration, 120
 - MS SPS integration, 74
 - MS SQL Server integration, 38
 - VSS integration, 160

D

- differential backups
 - MS SPS integration, 74
- disaster recovery
 - MS SQL Server integration, 51
- document
 - conventions, 22
 - related documentation, 15
- documentation
 - HP website, 15
 - providing feedback, 25

E

- EVA
 - checks and verifications, 190
 - replica creation and reuse, 160
- examples, MS Exchange Server integration
 - restoring using CLI, 111
- examples, MS Exchange Single Mailbox integration
 - restoring, 135
 - scheduling backups, 124
 - starting backups, 127
- examples, MS SPS integration
 - scheduling backups, 77
 - starting interactive backups, 80
- examples, MS SQL Server integration
 - restoring using CLI, 51
- Exchange Replication Service, 165

F

- full backups
 - Informix integration, 27, 93
 - MS Exchange Single Mailbox integration, 115
 - MS SPS integration, 74

H

- help
 - obtaining, 24
- HP
 - technical support, 24

I

- incremental backups
 - Informix integration, 27, 93
 - MS Exchange Single Mailbox integration, 115
- Informix backup
 - incremental backups, 27, 93
- Informix backup
 - full backups, 27, 93
- Informix restore
 - restore options, 50
- instant recovery
 - problems, 192
 - troubleshooting, 192
- integrated authentication, MS SQL Server integration, 35
- interactive backups
 - MS Exchange Single Mailbox integration, 126
 - MS SPS integration, 79
 - VSS integration, 170
- introduction
 - MS SQL Server integration, 27
 - MS Exchange Server integration, 93
 - MS SPS integration, 65
 - MS Exchange Single Mailbox integration, 115
 - VSS integration, 143

L

limitations

VSS integration, 152

Local Continuous Replication, 151

M

Microsoft Exchange Server

configuring writer, 155

Microsoft Exchange Server 2007

writer backup specifics, 165

modifying backup specifications

MS Exchange Single Mailbox
integration, 124

MS SPS integration, 77

monitoring sessions

MS Exchange Single Mailbox
integration, 136

MS SPS integration, 89

MS SQL Server integration, 58

VSS integration, 187

MS Exchange Server backup

backup specifications, creating, 98

MS Exchange Single Mailbox backup

backup options, 123

backup types, 115

incremental backups, 115

starting backups, 126

MS SPS backup

backup types, 74

scheduling backups, 77

transaction logs backups, 74

MS SPS configuration, 69 - 73

MS SPS integration

concepts, 66

introduction, 65

monitoring sessions, 89

MS SPS troubleshooting, 89 - 92

MS Exchange Server integration

restore, 105 - 112

MS Exchange Server backup

backup options, 103

MS Exchange Server configuration, 95
- 97

MS Exchange Server integration

architecture, 95

concepts, 94

configuration, 95 - 97

introduction, 93

troubleshooting, 112 - 114

MS Exchange Server restore, 105

restore options, 109

using CLI, 111

using GUI, 106

MS Exchange Server troubleshooting,
112 - 114

MS Exchange Single Mailbox backup
full backups, 115

MS Exchange Single Mailbox backup,
120

backup specifications, creating, 120

backup specifications, modifying, 124

performance tuning, 136

previewing backups, 125

scheduling backups, 124

scheduling backups, example, 124

starting backups, example, 127

MS Exchange Single Mailbox
configuration, 117 - 120

checking configuration, 119

MS Exchange Single Mailbox
integration

architecture, 116

backup, 120 - 127

concepts, 116

configuration, 117 - 120

introduction, 115

monitoring sessions, 136

restore, 127 - 136

troubleshooting, 137 - 142 - 142

- MS Exchange Single Mailbox restore, 127 - 136
 - examples, 135
 - restore options, 133
 - using CLI, 133
 - using GUI, 127
 - MS Exchange Single Mailbox troubleshooting, 137
 - MS SPS restore
 - using CLI, 86
 - using GUI, 81
 - MS SPS backup, 73 - 80
 - backup specification, modifying, 77
 - backup flow, 69
 - backup options, 77
 - backup specifications, creating, 74
 - differential backups, 74
 - full backups, 74
 - previewing backups, 78
 - scheduling backups, example, 77
 - starting backups, 79
 - MS SPS configuration
 - checking configuration, 73
 - MS SPS integration
 - architecture, 66
 - backup, 73 - 80
 - configuration, 69 - 73
 - restore, 81 - 88
 - troubleshooting, 89 - 92
 - MS SPS restore, 81 - 88
 - restore flow, 69
 - restore options, 86
 - specifying restore destination, 82
 - using another device, 88
 - MS SQL Server backup, 38
 - MS SQL Server backup, 44
 - backup options, 41
 - backup specifications, creating, 38
 - concepts, parallelism, 29
 - scheduling backups, 43
 - MS SQL Server configuration, 30, 37
 - checking configuration, 36
 - configuration files, 31
 - MS SQL Server integration
 - architecture, 29
 - backup, 38, 44
 - concepts, 28, 30
 - configuration, 30, 37
 - disaster recovery, 51
 - introduction, 27
 - monitoring sessions, 58
 - performance tuning, 54
 - restore, 44, 54
 - troubleshooting, 58, 63
 - MS SQL Server restore, 44, 54
 - disaster recovery, 51
 - restore options, 48
 - using CLI, 51
 - MS SQL Server troubleshooting, 58, 63
- O**
- online backups
 - MS SPS integration, 65
- P**
- performance tuning
 - MS Exchange Single Mailbox integration, 136
 - performance tuning
 - MS SQL Server integration, 54
 - prerequisites
 - VSS integration, 152
 - previewing backups
 - MS Exchange Single Mailbox integration, 125
 - previewing backups
 - MS SPS integration, 78
 - problems, 189 - 193
 - backup, 190
 - instant recovery, 192
 - restore, 192

R

- related documentation, 15
- replica creation and reuse
EVA, 160
- restore, 171 - 189
 - problems, 192
 - troubleshooting, 192
- restore flow
 - VSS integration, 149
 - MS SPS integration, 69
- restore options
 - MS Exchange Server integration, 109
 - Informix integration, 50
 - MS SQL Server integration, 48
 - MS Exchange Single Mailbox integration, 133
 - MS SPS integration, 86
 - VSS integration, 177
- restoring MS Exchange Server, 105 - 112
- restoring MS SPS
 - restore flow, 69
 - restore options, 86
 - using another device, 88
- restoring VSS
 - using GUI, 172
- restoring MS SQL Server, 44, 54
 - using CLI, 51
- restoring VSS
 - restore flow, 149
- restoring Informix
 - restore options, 50
- restoring MS Exchange Server
 - restore options, 109
- restoring MS Exchange Server
 - using CLI, 111
 - using GUI, 106
- restoring MS Exchange Single Mailbox, 127
 - examples, 135
 - restore options, 133
 - using CLI, 133
 - using GUI, 127

- restoring MS SPS, 81 - 88
 - specifying restore destination, 82
 - using CLI, 86
 - using GUI, 81
- restoring MS SQL Server
 - disaster recovery, 51
 - restore options, 48
- restoring VSS, 171 - 189
 - restore options, 177
- running backups
 - See starting backups

S

- scheduling backups
 - MS Exchange Single Mailbox integration, 124
- scheduling backups
 - MS SPS integration, 77
 - MS SQL Server integration, 43
 - VSS integration, 169
- SQL Server authentication, MS SQL Server integration, 35
- starting backups
 - MS SPS integration, 79
 - VSS integration, 170
- starting backups
 - MS Exchange Single Mailbox integration, 126
- Subscriber's Choice, HP, 25

T

- technical support
 - HP, 24
- technical support
 - service locator website, 25
- transaction logs backups
 - MS SPS integration, 74
- troubleshooting
 - backup, 190
 - EVA checks and verifications, 190
 - instant recovery, 192

- troubleshooting MS SQL Server, [58, 63](#)
- troubleshooting MS Exchange Server, [112 - 114](#)
- troubleshooting MS Exchange Single Mailbox, [137 - 142](#)
- troubleshooting MS SPS, [89 - 92](#)
- troubleshooting VSS, [189 - 193](#)

V

- Volume Shadow Copy Service
 - VSS integration, [143](#)
- Volume Shadow Copy Service
 - VSS integration, [143](#)
- VSS
 - Exchange Server 2007 writer concepts, [151](#)
- VSS backup
 - backup specifications, creating, [160](#)
 - flow, [148](#)
- VSS integration
 - backup, [159 - 171](#)
 - configuring, [153](#)
 - monitoring sessions, [187](#)
 - restore, [171 - 189](#)
- VSS restore
 - restore options, [177](#)
- VSS troubleshooting, [189 - 193](#)
- VSS backup, [159 - 171](#)
 - architecture, [145, 149](#)
 - backup specifications, [160](#)
 - backup flow, [148](#)
 - scheduling, [169](#)
 - scheduling backups, [169](#)
 - starting backups, [170](#)
- VSS configuration, [153](#)

- VSS integration
 - prerequisites, [152](#)
 - advantages, [146, 147](#)
 - architecture, [145, 149](#)
 - backup, [159 - 171](#)
 - concepts, [147](#)
 - configuration, [153](#)
 - introduction, [143](#)
 - limitations, [152](#)
 - prerequisites, [152](#)
 - troubleshooting, [189 - 193](#)
- VSS restore, [171](#)
 - restore flow, [149](#)
 - using GUI, [172](#)

W

- websites
 - HP Subscriber's Choice for Business, [25](#)
 - HP, [25](#)
 - product manuals, [15](#)
- Windows authentication, MS SQL Server integration, [35](#)
- writer backup specifics
 - Microsoft Exchange Server 2007, [165](#)
- writers specifics, VSS integration
 - Microsoft Exchange Server, [163](#)
- writers specifics, VSS integration, [155 - 159](#)
 - MS Data Protection Manager 2006, [167, 184](#)
 - MS Exchange Server, [179](#)
 - MS Exchange Server 2007, [181](#)
 - MSDE, [177](#)

