HP Data Protector 8.00
Integration Guide for Microsoft Applications
SQL Server, SharePoint Server, and Exchange Server

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

Contents

Publication history

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

Table 1 Edition history

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

About this guide

This guide describes how to configure and use Data Protector 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

The Help and other guides provide related information.

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

Help

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

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

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

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

Windows systems: Open `DP_help.chm`.

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

Guides

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

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

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

You can also access the guides:

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

Data Protector guides are:

- *HP Data Protector Getting Started Guide*

  This guide contains information to get you started with using Data Protector. It lists installation prerequisites, provides instructions on installing and configuring a basic backup environment and procedures for performing backup and restore. It also lists resources for further information.

- *HP Data Protector Concepts Guide*

  This guide describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented Help.

- *HP Data Protector Installation and Licensing Guide*

  This guide describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This guide also gives details on how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.

- *HP Data Protector Troubleshooting Guide*

  This guide describes how to troubleshoot problems you may encounter when using Data Protector.

- *HP Data Protector Disaster Recovery Guide*

  This guide describes how to plan, prepare for, test, and perform a disaster recovery.

- *HP Data Protector Command Line Interface Reference*

  This guide describes the Data Protector command-line interface, command options, and their usage as well as provides some basic command-line examples. It is located in the following directory:

  Windows systems: `Data_Protector_home\docs\MAN`

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

  On UNIX systems, you can use the `omniintro` man page to display a list of the available Data Protector commands. You can then execute the `man CommandName` command to retrieve information about each Data Protector command.

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

  This guide gives a description of new features of HP Data Protector 8.00. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.

- *HP Data Protector Integration Guides*

  These guides describe how to configure and use Data Protector to back up and restore various databases and applications. They are intended for backup administrators and operators. There are six guides:

  - *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*

    This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server.

  - *HP Data Protector Integration Guide for Oracle and SAP*

    This guide describes the integrations of Data Protector with Oracle Server, SAP R/3, and SAP MaxDB.

- *HP Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*

  This guide describes the integrations of Data Protector with the following IBM applications: Informix Server, IBM DB2 UDB, and Lotus Notes/Domino Server.

- *HP Data Protector Integration Guide for Sybase and Network Data Management Protocol Server*

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

- *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*

  This guide describes the integration of Data Protector with the Microsoft Volume Shadow Copy Service. This guide also documents application writer specifics.

- *HP Data Protector Integration Guide for Virtualization Environments*

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

- *HP Data Protector Zero Downtime Backup Concepts Guide*

  This guide describes Data Protector zero downtime backup and instant recovery concepts and provides background information on how Data Protector works in a zero downtime backup environment. It is intended to be used with the task-oriented *HP Data Protector Zero Downtime Backup Administrator's Guide* and the *HP Data Protector Zero Downtime Backup Integration Guide*.

- *HP Data Protector Zero Downtime Backup Administrator's Guide*

  This guide describes how to configure and use the integration of Data Protector with HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP 3PAR StoreServ Storage, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

- *HP Data Protector Zero Downtime Backup Integration Guide*

  This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases.

- *HP Data Protector Granular Recovery Extension User Guide for Microsoft Exchange Server*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft Exchange Server. Graphical user interface of the Data Protector Granular Recovery Extension for Microsoft Exchange Server is integrated into the Microsoft Management Console. This guide is intended for Microsoft Exchange Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft SharePoint Server. The Data Protector Granular Recovery Extension is integrated into Microsoft SharePoint Server Central Administration and enables you to recover individual items. This guide is intended for Microsoft SharePoint Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for VMware vSphere*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for VMware vSphere. The Data Protector Granular Recovery Extension is integrated into VMware vCenter Server and enables you to recover individual items. This guide is intended for VMware vCenter Server users and Data Protector backup administrators.

- *HP Data Protector Deduplication*

  This technical white paper describes basic data deduplication concepts, principles of Data Protector integration with Backup to Disk devices and its use of deduplication. It also provides instructions how to configure and use deduplication in Data Protector backup environments.

- *HP Data Protector Integration with Autonomy IDOL Server*

  This technical white paper describes all aspects of integrating Data Protector with Autonomy IDOL Server: integration concepts, installation and configuration, Data Protector backup image indexing, full content search-based restore, and troubleshooting.

- *HP Data Protector Integration with Autonomy LiveVault*

  This technical white paper all aspects of integrating Data Protector with Autonomy LiveVault: integration concepts, installation and configuration, backup policy management, cloud backup, cloud restore, and troubleshooting.

Documentation map

Abbreviations

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

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

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

Map

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

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

Integrations

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

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

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

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

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

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: "Document conventions" (page 14)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

CAUTION: Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT: Provides clarifying information or specific instructions.

NOTE: Provides additional information.

TIP: Provides helpful hints and shortcuts.

Data Protector graphical user interface

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

Figure 1 Data Protector graphical user interface

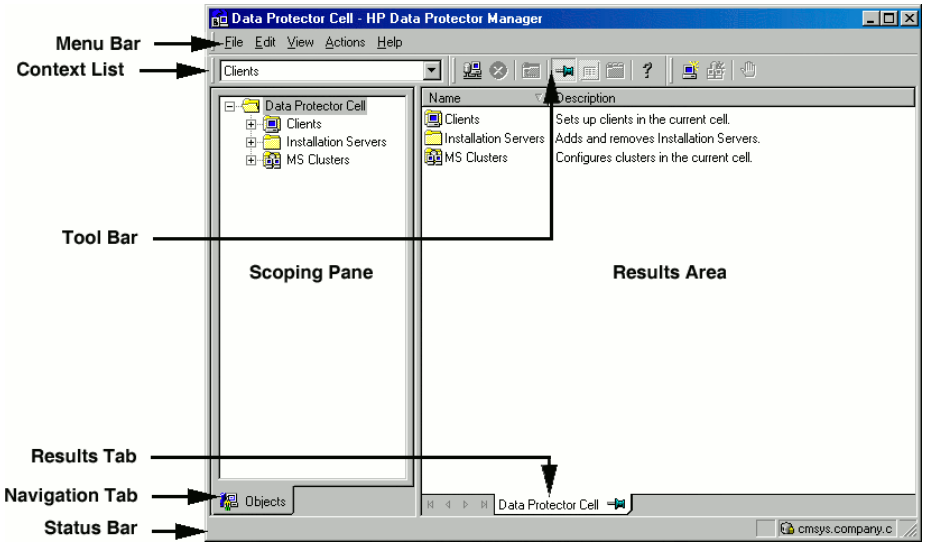

General information

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

HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

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

Documentation feedback

HP welcomes your feedback.

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

Part I Microsoft SQL Server

Data Protector integrates with Microsoft SQL Server through different integrations. These integrations provide different, but complementary functionality. Choose the appropriate integration depending on the desired functionality and the required platform coverage.

- ***Data Protector Microsoft SQL Server integration***

  Use this integration to back up all Microsoft SQL Server data or only individual databases.

  See "Data Protector Microsoft SQL Server integration" (page 18).

- ***Data Protector Microsoft SQL Server ZDB integration***

  Use this integration to back up Microsoft SQL Server data that resides on disk arrays. This integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Zero Downtime Backup Integration Guide*.

- ***Data Protector Microsoft Volume Shadow Copy Service integration***

  Use this integration to back up Microsoft SQL Server data using VSS writers.

  You can also use this integration to back up Microsoft SQL Server data that resides on disk arrays. This integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

NOTE: You can also back up Microsoft SQL Server databases using the common Data Protector file system backup functionality. Since you can only ensure data consistency by taking the Microsoft SQL Server offline, none of the benefits of the Data Protector integrations are available in this case.

1 Data Protector Microsoft SQL Server integration

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 Supported SQL Server online backup types

Full database backup	Includes all data regardless of the changes made after the last backup. In an availability group configuration, when you trigger a full backup of a database belonging to an availability group secondary replica, the backup type is automatically changed to a copy-only full 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. In a log shipping configuration, when a transaction log backup is triggered, the backup type is automatically changed to differential database backup.
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. Before you run a differential backup, make sure that a full backup exists. Otherwise, a restore from such a differential backup session fails. In an availability group configuration, when you trigger a differential backup of a database belonging to an availability group secondary replica, the backup type is automatically changed to a copy-only full backup.
Copy-only database backup[1]	A copy-only full backup is an independent full backup, which never truncates the transaction logs and does not affect an SQL Server restore chain. For this reason, it also cannot serve as a base of a differential backup. Run a copy-only full backup, if you do not want to influence a database backup.

[1] available by SQL Server 2008 or later

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 a different SQL Server, or to a different SQL Server instance. For detailed information, see "Restore options" (page 39).

This chapter provides information specific to this integration. For general Data Protector procedures and options, see the *HP Data Protector 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, copy-only full 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, standalone user databases, user databases belonging to availability groups, or certain 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 and availability group databases, see the SQL Server documentation. AlwaysOn Availability Groups solution is supported only on SQL Server 2012.

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

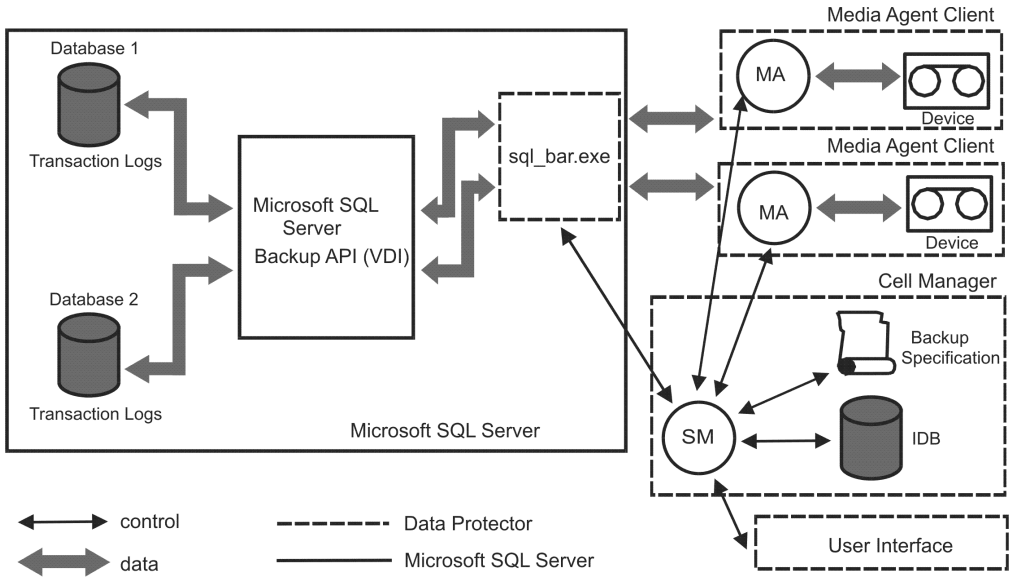

Table 4 Legend

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.
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 the 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.

"Database parallelism = 4, Overall Concurrency = 10" (page 20) shows a session in which each SQL Server database is backed up using a different number of concurrent streams.

Figure 3 Database parallelism = 4, Overall Concurrency = 10

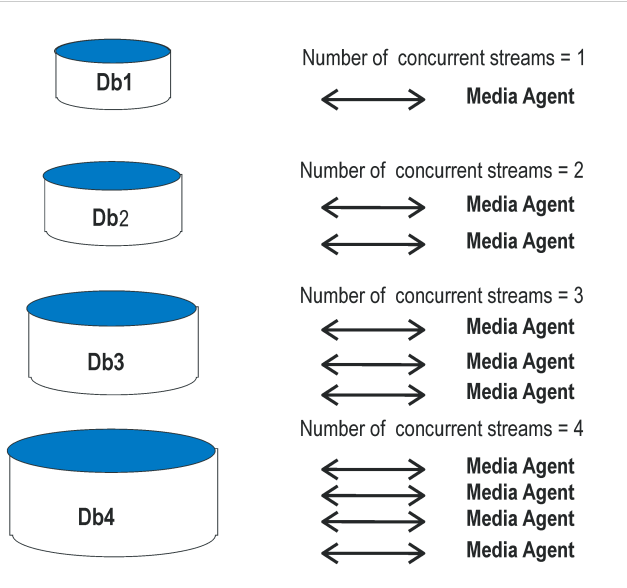

Configuring the integration

Prerequisites

- You need a license to use the SQL Server integration. For information, see the *HP Data Protector Installation and Licensing Guide*.
- Make sure that you correctly installed and configured SQL Server.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
  - For information on installing, configuring, and using SQL Server, see the SQL Server documentation.
- Make sure 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.

  In an availability group configuration, every availability group replica must have the `MS SQL Integration` component installed to enable backups across the availability group replicas.

Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the *HP Data Protector Help* index: "configuring devices" and "creating media pools". See also "Performance tuning" (page 43) for advanced options.
- On Windows Server 2003 systems, if you plan to use **Integrated authentication** to connect to an SQL Server instance, you need to restart the `Data Protector Inet` service under a Windows domain user account that has the appropriate SQL Server permissions for running backup and restore sessions. For information on changing the user account under which the `Data Protector Inet` service is running, see the *HP Data Protector Help* index: "Inet, changing account".

  However, for other supported Windows operating systems, you can use user impersonation instead. For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".
- Using the SQL Server Management Studio, add the user account which you will use for backing up and restoring SQL Server data to the fixed server role `sysadmin`. For instructions, see the SQL Server documentation.
- To test whether SQL Server and Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore. For instructions, see the *HP Data Protector Help*.

Data Protector SQL Server configuration file

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

HP-UX and Linux systems:

- For a standalone instance configuration

  `/etc/opt/omni/server/integ/config/MSSQL/`*`ClientName%InstanceName`*
- For an availability group configuration

  `/etc/opt/omni/server/integ/config/MSSQL/`*`ListenerName%AGName`*

Windows systems:

- For a standalone instance configuration

  `Data_Protector_program_data``\Config\Server\Integ\Config\MSSQL\`*`ClientName%InstanceName`*
- For an availability group configuration

  `Data_Protector_program_data`
  `\Config\Server\Integ\Config\MSSQL\`*`ListenerName%`* `\` *`AGName`*

  `ListenerName` is the name of an availability group listener, the virtual client used to connect to the SQL Server. *`AGName`* is the name of the SQL Server availability group corresponding to the selected listener.

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 content of the configuration file is:

```
Login='user';
Password='encoded_password';
Domain='domain';
Port='PortNumber';
```

IMPORTANT: To avoid backup problems, make sure that the syntax of your configuration file matches the examples. In an availability group configuration, also provide a port number used by the availability group listener to connect to the SQL Server. The default is 1433.

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

On Windows Server 2003 system, if you have restarted the `Data Protector Inet` service on the SQL Server system under a different user account, add this user to the Data Protector `admin` or `operator` Data Protector user group.

For information on adding users to the Data Protector groups, see the *HP Data Protector Help* index: "adding users".

Configuring an SQL Server cluster

In a cluster, all the nodes must be installed as Data Protector cluster-aware clients and the `Data Protector Inet` service on all nodes must run under a Windows domain user account that has also cluster administrator rights.

You must configure the `Data Protector Inet` service user impersonation for all cluster nodes. The Windows domain user account that is used must be given the following Windows operating system Security Policy privileges:

- Impersonate a client after authentication
- Replace a process level token

For more information, see the *HP Data Protector Help* index: "cluster-aware client", "Inet user impersonation", and the SQL Server cluster documentation.

Configuring SQL Server instances

An SQL Server instance is configured during the creation of the first backup specification. The configuration consists of setting the user account that Data Protector should use to connect to the SQL Server instance. The specified login information is saved to the Data Protector SQL Server instance configuration file on the Cell Manager.

If your SQL Server supports an AlwaysOn Availability Groups solution, you can configure availability groups instead of standalone instances. An availability group contains a set of read-write availability group primary replica databases and one to four sets of corresponding availability group secondary replica databases. For more information, see the SQL Server documentation.

NOTE: Make sure 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" (page 25).

Prerequisites

- SQL Server must be online during configuration.
- Make sure that the SQL Server Browser service is running.
- Configuration must be performed for every SQL Server instance separately.

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. For availability group configurations, select the availability group listener of the corresponding availability group. Note that you must first import the availability group listener as a virtual client by selecting **Virtual Host** in the Clients context. For details on how to import clients, see the HP Data Protector Help index: "importing, client systems".

   In **Application database**, select or specify the name of the SQL Server instance. In an availability group configuration, the name of the SQL availability group which is connected to the selected availability group listener is listed automatically and cannot be changed.

   Windows Server 2008: If you intend to use **Integrated authentication** and you want that the backup session runs under the specified operating system user account, specify the **Specify OS user** option. For information on the **User and group/domain** options, press **F1**.

   Click **Next**.
5. In the Configure MS SQL Server 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**: Select this option to enable Data Protector to connect to the SQL Server instance with the following Windows domain user account:
     - ***Windows Server 2008:*** The account specified in the **User and group/domain** options in the previous step or in the Client selection page.
     - ***Other Windows systems:*** The account under which the `Data Protector Inet` service on the SQL Server system is running.

   Make sure that the user account you specify has the appropriate permissions for backing up and restoring the SQL Server databases.

   For an availability group configuration, you can also provide a port number used by the availability group listener. The default is 1433.

   See "Configuring SQL Server" (page 24).

   If configuring the integration in an availability group environment, see "Configuring SQL Server — an AlwaysOn Availability Groups solution" (page 24).

Figure 4 Configuring SQL Server

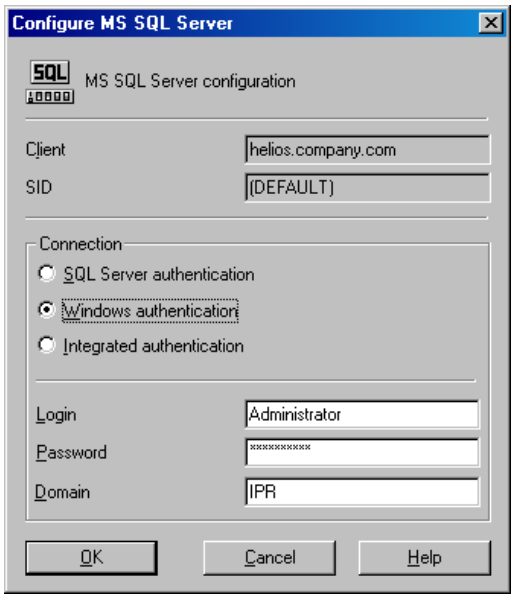

Figure 5 Configuring SQL Server — an AlwaysOn Availability Groups solution

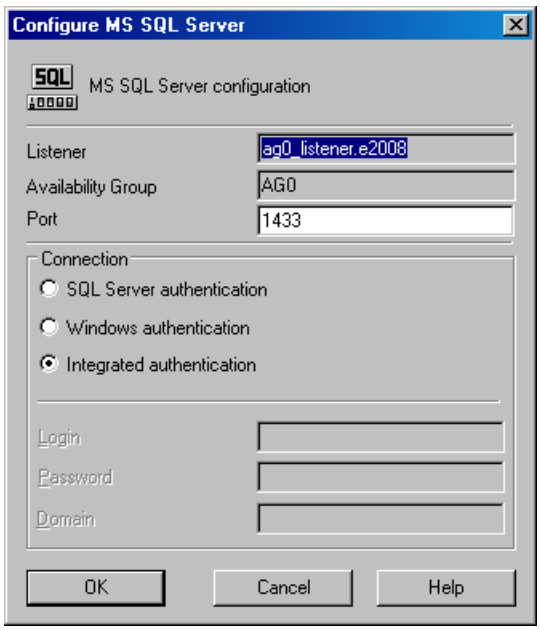

NOTE: It is recommended that the SQL Server system administrator configures the integration.

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 "Creating backup specifications" (page 27).

Using the Data Protector CLI

Execute:

- For a standalone instance configuration:

  ```
  sql_bar config [-appsrv:SQLServerClient] [-instance:InstanceName]
  [-dbuser:SQLServerUser -password:password | -dbuser:WindowsUser
  -password:password -domain:domain]
  ```

- For an availability group configuration:

  ```
  sql_bar econfig [-appsrv:ListenerName] [-ag:AGname]
  [-dbuser:SQLServerUser -password:password | -dbuser:WindowsUser
  -password:password -domain:domain]-port:PortNumber
  ```

Parameter description

`-appsrv:SQLServerClient`
: The client system on which the SQL Server instance is running. This option is not required if you run the command locally.

`-appsrv:ListenerName`
: The name of an availability group listener, the virtual client on which the SQL Server availability group is running.

`-instance:InstanceName`
: The SQL Server instance name. If you omit this option, the default SQL Server instance is configured.

`-ag:AGname`
: The SQL Server availability group name.

`-dbuser:SQLServerUser -password:password`
: The SQL Server user account (**SQL Server authentication**)

`-dbuser:WindowsUser -password:password -domain:domain`
: The Windows domain user account (**Windows authentication**)

`-port:PortNumber`
: The port number used by the availability group listener to connect to the SQL Server. The default is 1433.

NOTE: If no user account is specified, Data Protector uses **Integrated authentication**.

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

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 instances" (page 22).

5. Right-click SQL Server and select **Check Configuration**. See "Checking configuration" (page 26).

Figure 6 Checking configuration

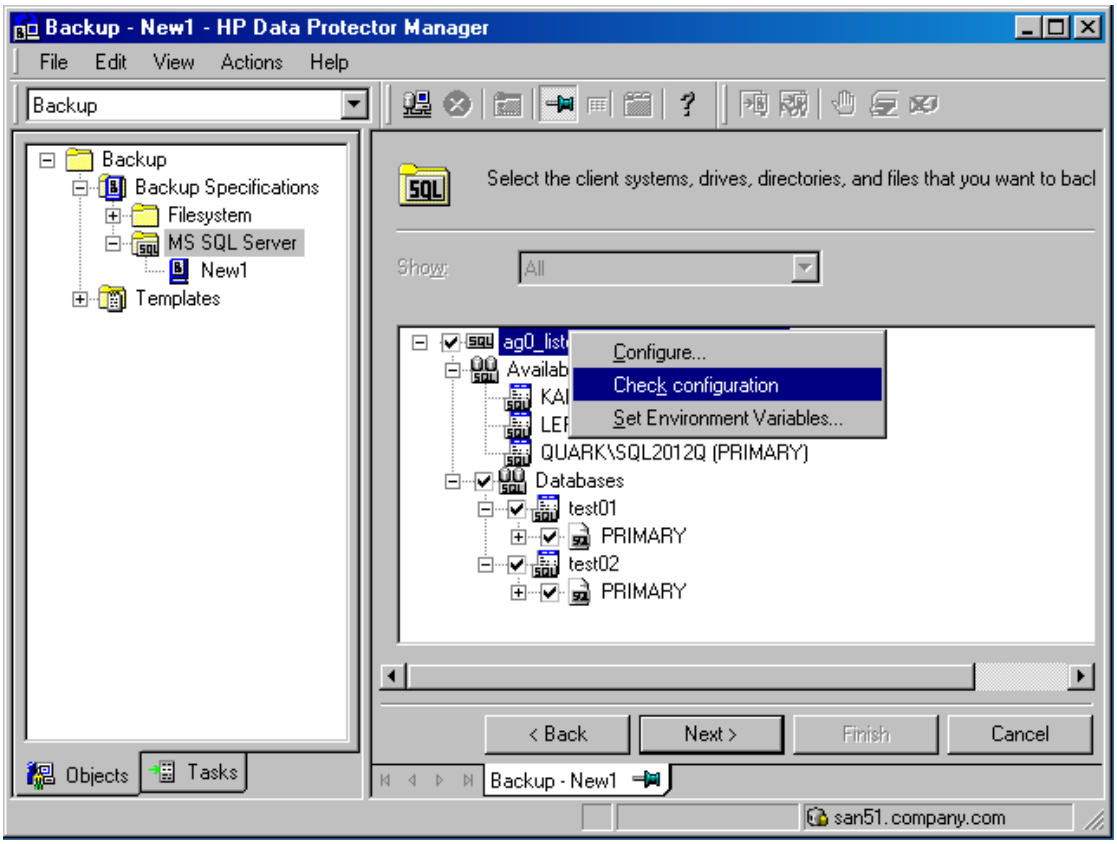

Using the Data Protector CLI

To change the configuration, run the command for configuring SQL Server instances again, entering different data. In an availability group configuration, run the command for configuring SQL Server availability groups, entering different data.

To check configuration for a standalone instance, run:

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

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

To check configuration for an availability group, run:

```
sql_bar chkconf -ag agname –appsrv:ListenerName
```

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

```
*RETVAL*8523
```

To get the information about the existing configuration for a standalone instance, run:

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

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

To get the information about the existing configuration for an availability group, run:

```
sql_bar getconf -ag agname –appsrv ListenerName
```

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.

Limitations

- Backup preview is not supported.
- If a database belonging to an availability group secondary replica has the `Readable Secondary` parameter set to `NO`, the Data Protector GUI cannot display database components such as file groups and data files. Therefore, you can only back up the entire database.
- When backing up a database belonging to an availability group, make sure that you either create a standalone instance backup specification or an availability group backup specification. Do not use both backup specification types as it may result in a broken restore chain.

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 "Selecting a blank Microsoft SQL Server backup template" (page 28).

Figure 7 Selecting a blank Microsoft SQL Server backup template

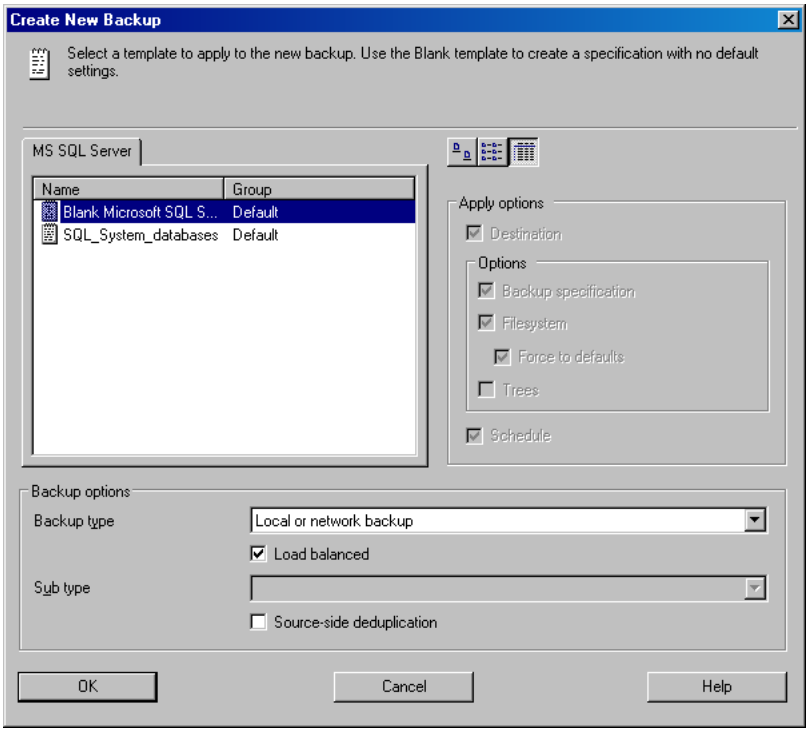

Click **OK**.

4. In **Client**, select an SQL Server. For cluster environments, select the virtual server of the SQL Server resource group. For availability group configurations, select the availability group listener of the corresponding availability group. Note that you must first import the availability group listener as a virtual client by selecting **Virtual Host** in the Clients context. For details on how to import clients, see the HP Data Protector Help index: "importing, client systems".

   In **Application database**, specify the name of the SQL Server instance. In an availability group configuration, the name of the SQL availability group which is connected to the selected availability group listener is listed automatically and cannot be changed.

   Windows Server 2008: If you intend to use **Integrated authentication** and you want that the backup session runs under the specified operating system user account, specify the **Specify OS user** option. For information on the **User and group/domain** options, press **F1**.

   Click **Next**.

5. If the client is not configured, the **Configure MS SQL Server** dialog box appears. Configure it as described in "Configuring SQL Server instances" (page 22).
6. Select the databases, file groups, or data files you want to back up.

Figure 8 Selecting backup objects — standalone instance backup

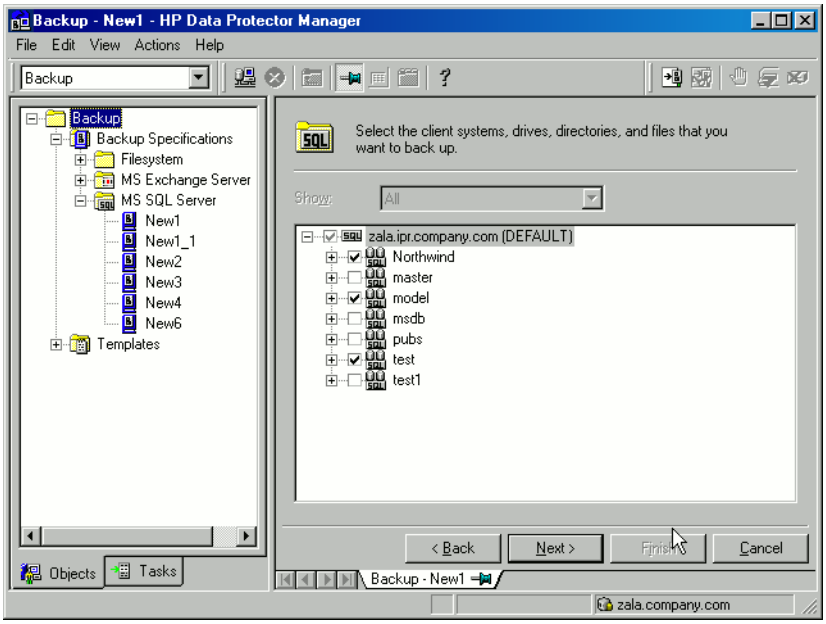

In an SQL Server availability group environment, when you are creating a standalone instance backup specification, Data Protector displays the name of the database together with the availability group name and its availability group replica role. For example, in the Data Protector GUI, the database named `DB1` belonging to the availability group primary replica named `AG1` is displayed as `DB1 [AG1 primary]`. Note that the name of the database together with the availability group name and its availability group replica role is visible only during the creation of the backup specification and not after it has already been saved.

Figure 9 Selecting backup objects — availability group backup

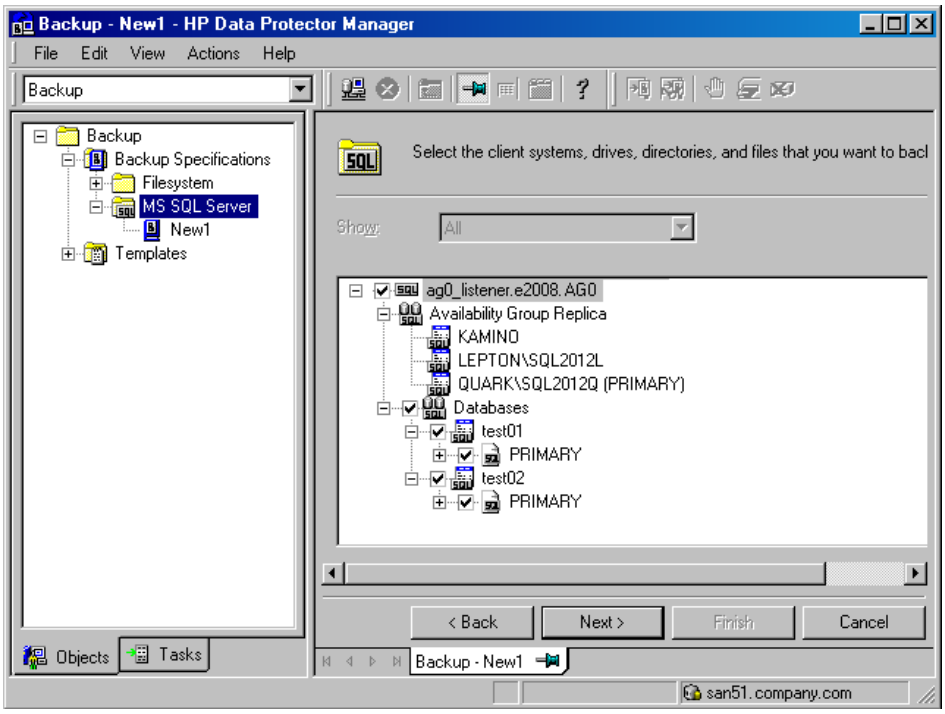

When you are creating an availability group backup specification, you can expand the availability group listener to display availability group replica clients and the databases belonging to the selected availability group. However, the clients are displayed for information only and cannot be selected. You can only select the databases to be backed up.

You can set the Data Protector Microsoft SQL Server-related environment variables by right-clicking the selected Microsoft SQL Server instance and selecting **Set Environment Variables**. In the Advanced dialog box, specify the desired variables and their values. Click **OK** to close the dialog box and store the settings into the Microsoft SQL Server configuration file. Note that environment variables override omnirc options that may be set client-wide in the `omnirc` file.

Figure 10 Setting environment variables

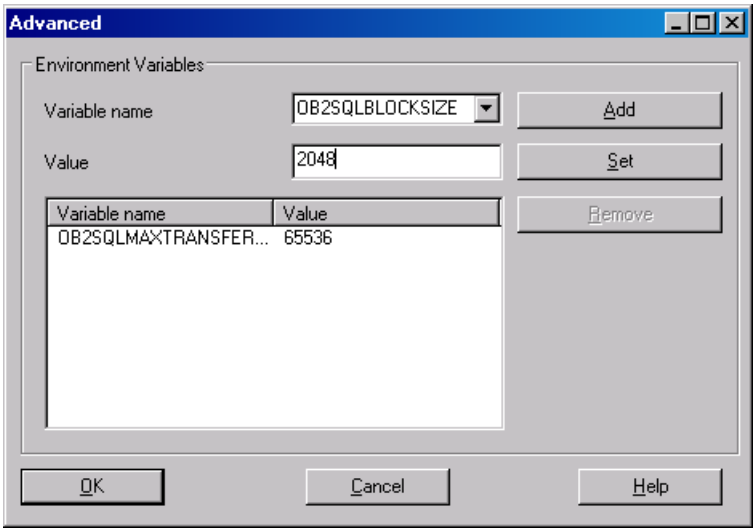

Click **Next**.

7. Select the devices. Click **Properties** to set the media pool and preallocation policy. The device concurrency is set to 1 and cannot be changed. 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 *HP Data Protector Help*.

   Click **Next**.

8. Select backup options.

   For information on **Backup Specification Options** and **Common Application Options**, see the *HP Data Protector Help*.

   For information on **Application Specific Option**, see "SQL Server-specific backup options" (page 31).

   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 11 Saving a backup specification

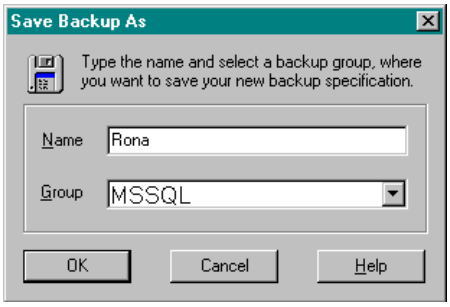

SQL Server-specific backup options

Specify SQL Server-specific backup options by clicking **Advanced** in the **Application Specific Options** group box and selecting the desired options by clicking **MS SQL integration** and **MS SQL backup preference** page.

Figure 12 Application-specific options

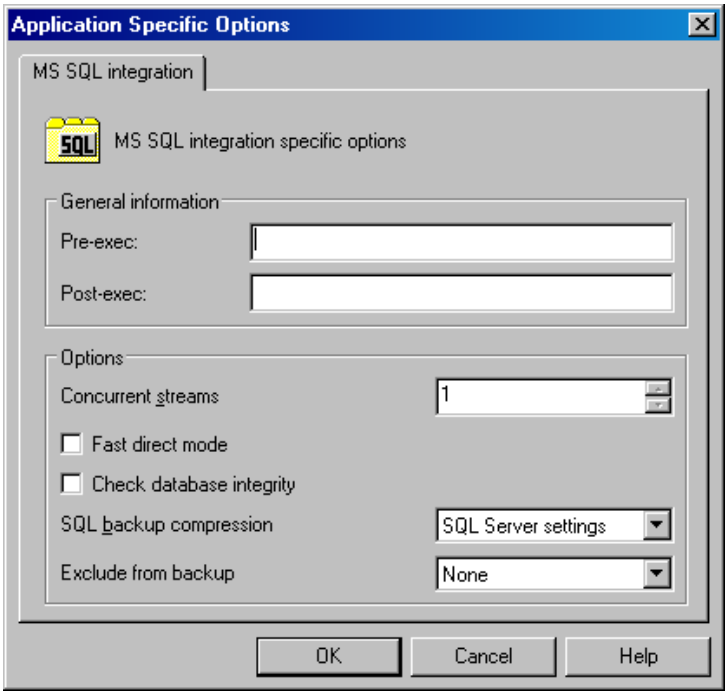

Figure 13 Application specific options — backup preferences

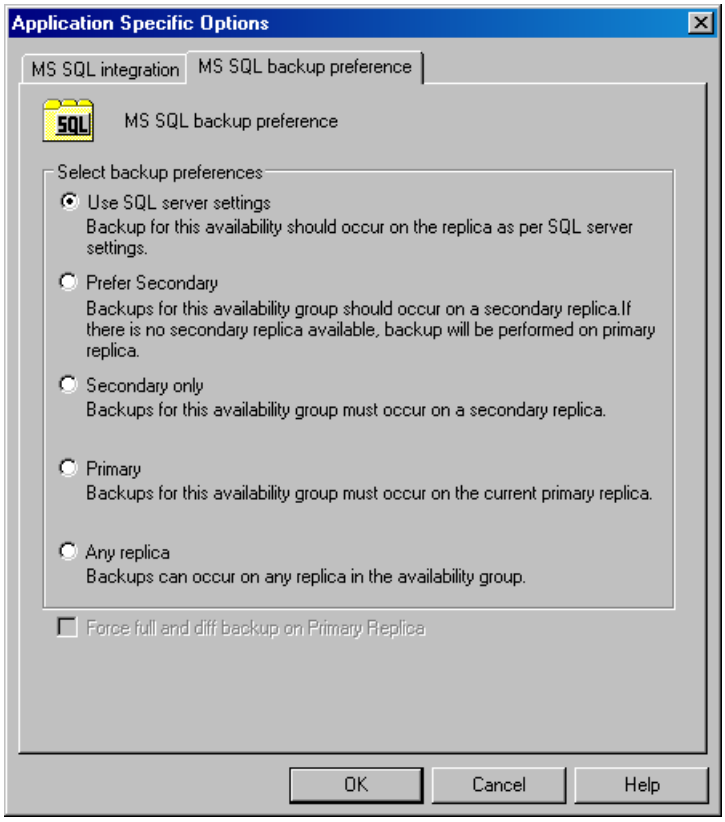

Table 5 SQL Server backup options

Pre-exec	Specifies a command with arguments or a script started by `sql_bar.exe` on SQL Server before backup. Resides in the *`Data_Protector_home`*`\bin` directory. Only the filename must be provided in the backup specification.	
Post-exec	Specifies a command with arguments or a script started by `sql_bar.exe` on SQL Server after backup. Resides in the *`Data_Protector_home`*`\bin` directory. Only the filename must be provided in the backup specification.	
Concurrent streams	Sets the number of concurrent streams used for backup.	
Fast direct mode	Used with locally connected devices to optimize performance. Must be combined with special device settings (see "Performance tuning" (page 43) for details).	
Check database integrity	Performs data integrity validation before backup. If the check fails, the session completes with warnings.	
SQL backup compression	Specify how Data Protector should handle the Microsoft SQL Server backup compression.	
	SQL Server settings (default)	Handles the backup compression according to the Microsoft SQL Server settings.
	Enable	Executes the backup compression regardless of the Microsoft SQL Server settings.
	Disable	Specifies that the backup compression should not be executed regardless of the Microsoft SQL Server settings.
Exclude from backup *(available for standalone*	Excludes specific databases from backup.	
	Availability Group Databases	Excludes databases belonging to any availability group from backup.

Table 5 SQL Server backup options *(continued)*

instance backup only)	**Standalone databases**	Excludes all standalone databases from backup.
	None (default)	Does not exclude any database from backup.
Select backup preferences *(available for availability group backup only)*	**Use SQL server settings** (default)	Performs backup according to the Microsoft SQL Server settings.
	Prefer Secondary	Performs backup of availability group databases on an availability group secondary replica. If there is no availability group secondary replica available, backup is performed on an availability group primary replica.
	Secondary only	Performs backup of availability group databases on an availability group secondary replica. If there is no availability group secondary replica available, backup fails.
	Primary	Performs backup of availability group databases on a primary replica.
	Any replica	Performs backup on any availability group replica in the availability group.
Force full and diff backup on Primary Replica	If selected, full and differential backups are always performed using the availability group primary replica, regardless of the selected backup preference. Prefer Secondary is used for transaction log backups only. If not selected, copy-only full backups are performed instead of full or differential backups when backing up a database belonging to an availability group secondary replica.	

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

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 or by clicking an object and then **Properties...**.

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

Figure 14 Object properties

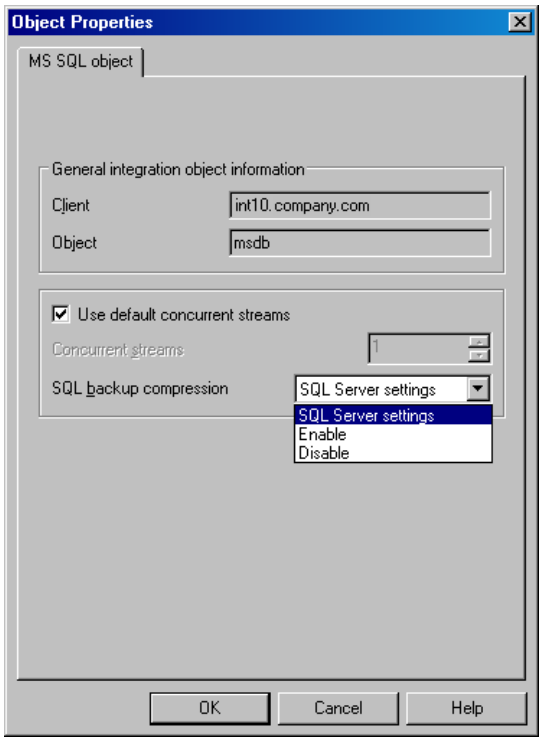

Table 6 Object-specific options

Use default concurrent streams	The number of concurrent streams is defined by Data Protector and all available devices are used.	
Concurrent streams	Sets the number of concurrent streams (devices). VDI supports up to 32 virtual devices per database.	
SQL backup compression	Specify how Data Protector should handle the Microsoft SQL Server backup compression.	
	SQL Server settings (default)	Handles the backup compression according to the Microsoft SQL Server settings.
	Enable	Executes the backup compression regardless of the Microsoft SQL Server settings.
	Disable	Specifies that the backup compression should not be executed regardless of the Microsoft SQL Server settings.
Exclude from backup *(available for standalone instance backup only)*	Excludes specific databases from backup.	
	Availability Group Databases	Excludes databases belonging to any availability group from backup.
	Standalone databases	Excludes all standalone databases from backup.
	None (default)	Does not exclude any database from backup.

Scheduling backups

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

Scheduling example

To schedule 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.

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 use 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 a different SQL Server or to a different SQL Server instance. For detailed information, see "Restore options" (page 39).

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

To recover the master database, start the SQL Server disaster recovery process. For more information, see the "Disaster recovery" (page 42).

Before you begin

- Verify that the databases to be restored are not being in use.
- In an availability group configuration, restore to a different client and instance is mandatory. Before restoring such a database, make sure that you do not select an availability group listener for the target client, and that the selected SQL Server instance exists on the target client. Also make sure that the database which you selected for the restore does not belong to any availability group.

Restoring using the Data Protector 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 SQL Server**, and then select the Microsoft SQL Server from which you want to restore. A list of backed up objects is displayed in the Results Area.
3. Select the backed up SQL Server database or data files you want to restore.

   To restore a file group, expand it and select all data files in it.

IMPORTANT: Before a data file can be restored, the active transaction log of the database must be backed up. In case the log has been corrupted, particular data files cannot be restored, and you can only restore the entire database.

See "Selecting backup objects for restore" (page 36) and "Selecting backup objects for restore in an availability group configuration" (page 36).

Figure 15 Selecting backup objects for restore

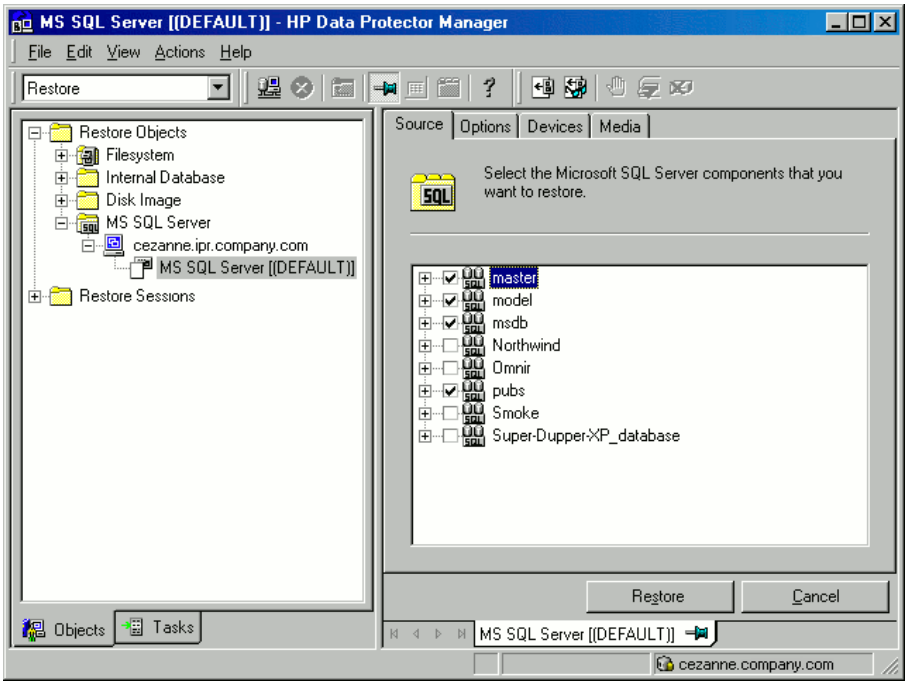

Figure 16 Selecting backup objects for restore in an availability group configuration

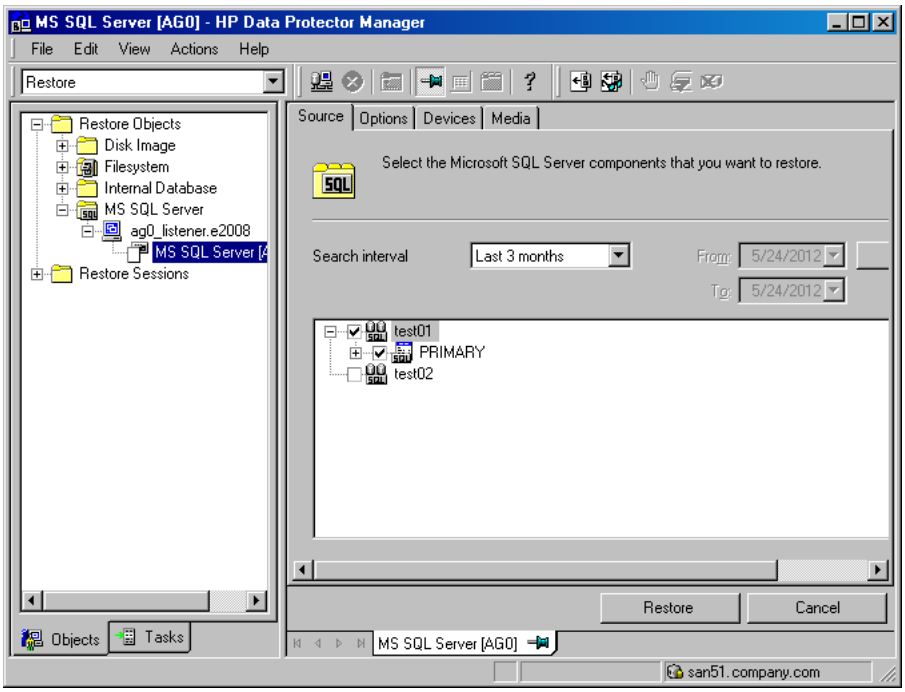

To select backup object-specific options, right-click the object and select **Properties**.

Figure 17 Selecting object-specific options

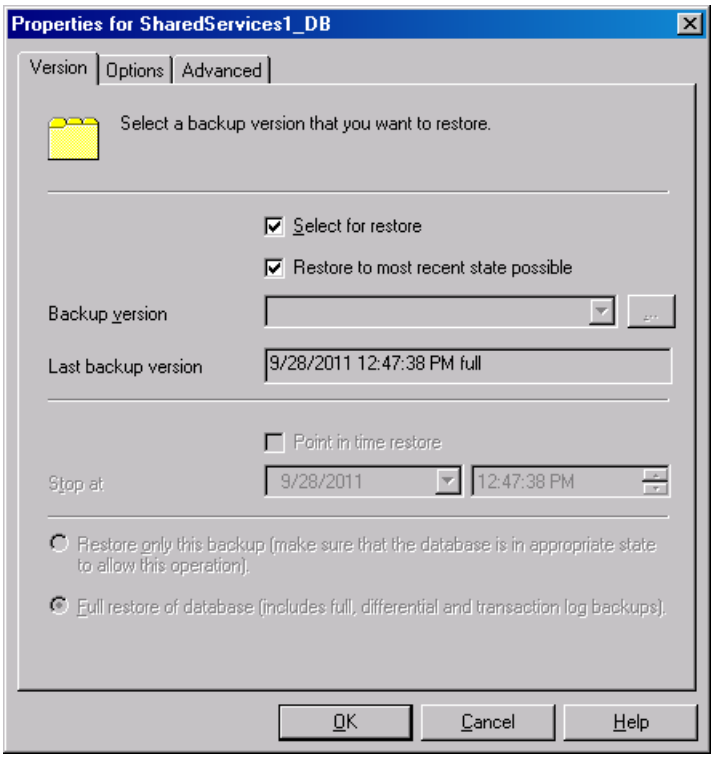

In the Version tab, select the backup version (backup date) which you want to use for restore or select the option **Restore to most recent state possible**. The latter always restores the chain of backups as if the **Full restore of database** option is selected. It includes the most recent full, differential, and transaction log backups.

Optionally, in the Advanced tab, select the **Restore database with new name** option and specify new restore locations.

IMPORTANT: In an availability group configuration, restoring a database to a different location is mandatory. However, the database can be restored with the same name, if it does not belong to any availability group.

Figure 18 Restoring a database with a different name and location

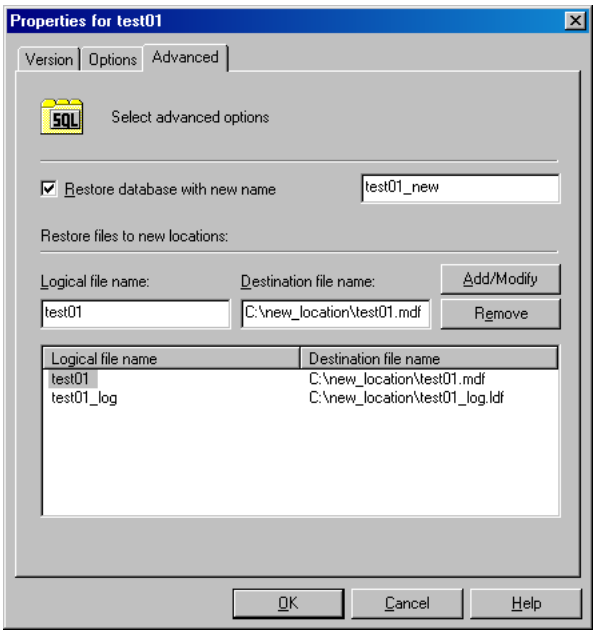

Select other restore options as appropriate. Note that some options are not available for restore of data files. See "Restore options" (page 39) for details.

Click **OK**.

4. In the **Options** property page, specify new locations for the databases, if you want to restore your data to a different client or instance.

IMPORTANT:

- 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.
- Make sure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.

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

- **Restore data**. Select to restore the whole database. This option is selected by default.
- **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.

Select **Enable tail log backup** to perform a tail log backup session, just before the restore session starts, using the backup specification selected in the drop-down list. This captures the logs from the tail that have not been backed up yet. Before selecting this option, make sure that:

- the option **Put database in single user mode - log off all users** is selected for all involved databases.
- the option **Restore data** is selected.

IMPORTANT: When restoring a database to a different client or/and instance, enabling tail log backup is not recommended. Therefore do not select **Enable tail log backup** if you are restoring a database backed up using an availability group configuration.

Figure 19 Restore options

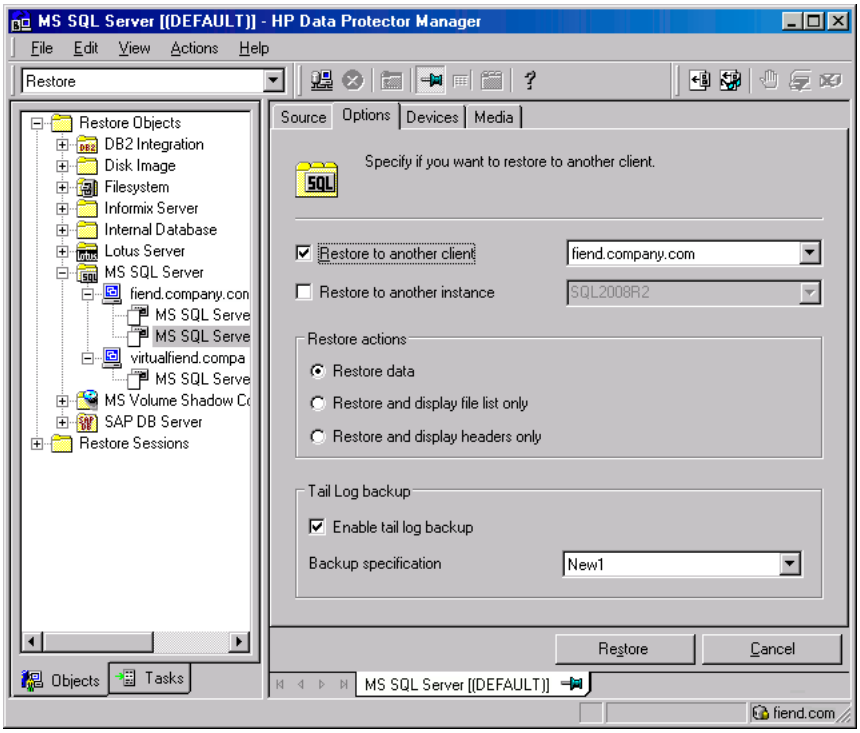

5. In the **Devices** page, select the devices to be used for the restore.

   For more information of how to select devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".

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 Microsoft SQL Server database restore options

Option	Description
Backup version	Specifies the backup session from which the selected objects will be restored.
Point-in-time restore	This option is only available for database objects. 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.
Stop at	This option is only available for database objects. Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable database recovery to a particular point in time, backup you restore from must be a transaction log backup. You cannot use this option 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**).
Restore only this backup	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.

Table 7 Microsoft SQL Server database restore options *(continued)*

Option	Description
Full restore of the database	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.
Force restore over the existing database	Select this option if a database with the same name but a different internal structure already exists at the target Microsoft SQL Server instance. If this option is not selected, the Microsoft SQL Server does not let you overwrite the existing database - the restore will fail. If you are restoring a data file from the PRIMARY group to an existing database, you must specify the option at the data file level. When using this option, make sure that the most recent logs are backed up before the restore.
Put database in single user mode - log off all users	Disconnects all users that are connected to the target Microsoft SQL Server database and puts the database in the single user mode. Note that if the database is not in the simple recovery mode, the **Force restore over the existing database** option should also be selected.
Recovery completion state	Enables selecting the database state after recovery. You may select from: • 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. This selection is only available for database objects.
Restore database with a new name	This option is only available for database objects. Restores the database under a different name. Specify the database logical filename and the destination filename (suboptions of **Restore files to new locations**).
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 to restore data to a different client, a different instance, or to make a database copy on the same client.
Restore to most recent state possible	Restores the entire backup chain (includes the most recent full, differential, and transaction log backups). This option is selected by default.

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 the existing database**.

Restoring to a different SQL Server instance or/and different 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 the configuration procedure, see "Creating backup specifications" (page 27).

Procedure

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

   IMPORTANT: Make sure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.
3. Specify new database locations.
4. Start restore. See "Restore" (page 35).

Restoring using the Data Protector CLI

Execute:

```
omnir -mssql -barhost ClientName [-destination ClientName] [-instance
SourceInstanceName] [-destinstance DestinationInstanceName] {—base
DBName [—session BackupID] [MSSQL_OPTIONS]... | —base DBName —datafile
GroupName/DataFileName —session BackupID [DATAFILE_OPTIONS]...}
MSSQL_OPTIONS
-asbase NewDBName {-file LogicalFileName1 PhysicalFileName1 [-file
LogicalFileName2 PhysicalFileName2]...}
-replace
-nochain
-recovery {rec | norec}
-standby File
-tail_log BackupSpecificationName
DATAFILE_OPTIONS
-replace
-nochain
-recovery {rec | norec}
```

NOTE:

- *BackupID* is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

  Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.

  The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.
- The *SourceInstanceName* is case-sensitive; it has to be the same as the name of the SQL Server instance that you specified in the backup specification. See Step 4.

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

Examples

To restore the database `RONA` running on the SQL Server `ALMA` to the same destination, execute:

```
omnir -msssql -barhost ALMA -base RONA
```

To restore the data file `DATAFILE_01` in the file group `FILEGROUP_02` of the database `RONA` running on the SQL Server `ALMA` to the same destination, execute:

```
omnir -MSSQL -barhost ALMA -base RONA –datafile FILEGROUP_02/DATAFILE_01
–session 2011/10/17-3
```

Disaster recovery

Disaster recovery is a complex process involving products from different vendors. Therefore, you need to check the instructions from the database or 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.
- When reinstalling SQL Server, make sure that you use original local settings. Before restoring to a different client, also make sure 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\rebuildm.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. Start the 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" (page 27).

3. 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. Stop the service.
   d. Enter `-m` as a start-up parameter and start the services.
4. Restore the master database using the Data Protector Manager.

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" (page 43).

Recovering user databases

To restore user databases, proceed as described in "Restore" (page 35).

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 can 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. Make sure 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 the default. 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 "Adjusted local device" (page 45).

     You can modify block size later; however, you must first calculate it using the formula above and then insert the value as shown in "Advanced options" (page 46).

   - 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 *HP Data Protector Help* index: "changing block size" for instructions.

Figure 20 Adjusted local device

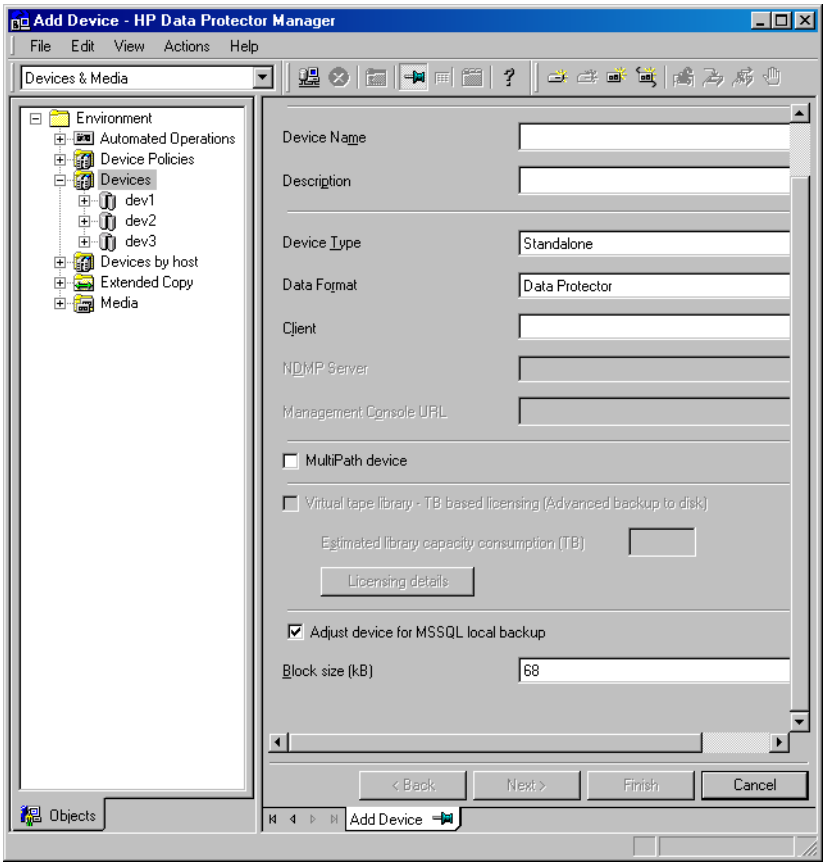

To modify block sizes of an existing device:

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 21 Advanced options

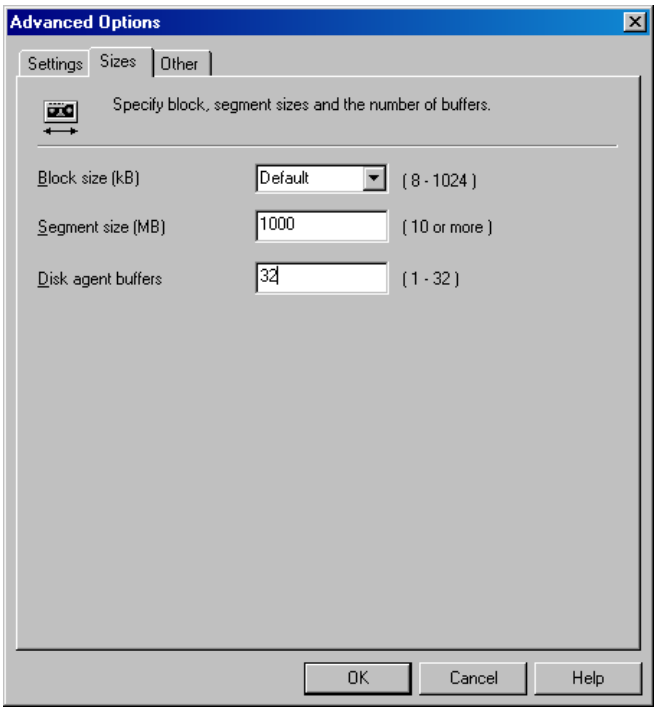

If **Fast direct mode** is activated and not all selected local devices in a backup specification are adjusted accordingly, you get the warning when saving the backup specification:

Figure 22 Block sizes 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 of 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 *HP Data Protector 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 *HP Data Protector 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" (page 48). 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

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

Checks and verifications

If your configuration, backup, or restore failed:

- Check that SQL Server services are running.
- Examine system errors reported in *Data_Protector_program_data*\log\debug.log on the SQL Server 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 *HP Data Protector 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 master database restore, the following error occurs when executing an SQL statement:

```
Error has occurred while executing an SQL statement.
Error message: 'SQLSTATE:[42000] CODE:(3108) MESSAGE:[Microsoft]
[ODBC SQL Server Driver][SQL Server]To restore the master database,
the server must be running in single user mode. For information on
starting in single user mode, see "How to: Start an Instance of SQL
Server (sqlservr.exe)" in Books Online.
```

Note that this behavior is expected when the master database is not restored in 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@computer.company.com "SQLSRV"
Time: 7/29/2011 8:19:22 PM
Error has occurred while executing SQL statement.
[Microsoft][ODBC SQL Server Driver][SQL Server]Backup or restore
operation terminating abnormally.'
[Critical] From: OB2BAR@computer.company.com "SQLSRV"
Time: 7/29/11 8:19:24 PM
Received ABORT request from SM => aborting
```

- SQL Server error log contains an entry similar to:

```
2011-07-29 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:

```
2011/07/30 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 with "The object was not open"

When backing up Microsoft SQL Server databases, the session fails with an error similar to the following:

```
[Critical]From : OB2BAR_Main@wemaoldb2dr "Aolins" Time:11/12/2011
02:01:34 AM Microsoft SQL Server reported the following error during
login : The object was not open
```

The error may appear if the SQL Server Browser service is not running.

Action

Proceed as follows:

1. Start the SQL Server Browser service.
2. Start a new backup session.

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

When you try to restore an SQL Server database from an object copy session, the restore fails.

An SQL Server database backed up using multiple streams (the **Concurrent streams** option set to more than 1) can only be restored if the backup objects created by the streams reside on separate media. During a Data Protector Microsoft SQL Server backup, each stream is always backed up to a separate medium. However, if you copy these backup objects on the same medium, using the object copy functionality, and start a restore from the object copy session, the restore fails.

Action

Before restarting the restore:

1. Increase the number of Disk Agent buffers for the device.
2. In the **Internal Database** context, find the objects belonging to the same backup (identified by the same backup ID).
3. 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.
4. 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" (page 20)).

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" (page 35).
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" (page 35).
13. Click **Restore** and then **Next** to select the Report level and Network load.
14. Click **Finish** to start the restore session.

Problem

Database restore fails

The restore session aborts with a major error similar to:

```
Error has occurred while executing a SQL statement.Error message:
'SQLSTATE:[42000] CODE:(3159) MESSAGE:[Microsoft][ODBC SQL Server
Driver][SQL Server]The tail of the log for the database "test2" has not
been backed up. Use BACKUP LOG WITH NORECOVERY to backup the log if it
contains work you do not want to lose. Use the WITH REPLACE or WITH
STOPAT clause of the RESTORE statement to just overwrite the contents
of the log. SQLSTATE:[42000] CODE:(3013) MESSAGE:[Microsoft][ODBC SQL
Server Driver][SQL Server]RESTORE DATABASE is terminating abnormally.
```

Action

To solve the problem , perform one of the following before restarting the restore session:

- Select the restore option **Enable tail log backup** (recommended).
- Perform a transaction log backup to obtain the most recent transaction logs.

Problem

Restore of a database in a log shipping configuration with the tail log backup enabled fails

In a Microsoft SQL Server log shipping configuration, Data Protector performs differential database backup instead of transaction log backup when the latter is run. The automatic backup type switch takes place also with tail log backup. In these circumstances, the database backup chain does not contain most recent transactions from the tail of the log. If tail of the log of the target database has not been backed up yet, Microsoft SQL Server does not allow restoring over this database.

Action

Perform one of the following and restart the restore session:

- Disable Microsoft SQL Server log shipping.
- Enable the option **Force restore over existing database** for all involved databases.

CAUTION: Tails of the logs of all involved databases will be lost.

Part II Microsoft SharePoint Server

Data Protector integrates with Microsoft SharePoint Server through integrations and extensions. They provide different, but complementary functionality. Choose the appropriate integration depending on your Microsoft SharePoint Server version and the desired functionality.

Microsoft Office SharePoint Server 2007 and Microsoft SharePoint Server 2010

- ***Data Protector Microsoft SharePoint Server 2007/2010 integration***

  See "Data Protector Microsoft SharePoint Server 2007/2010 integration" (page 55).

- ***Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution***

  Use this integration to back up Microsoft SharePoint Server using VSS writers. See "Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution" (page 83).

  You can also use this integration to back up Microsoft SharePoint Server data that resides on disk arrays. This integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions. See the *HP Data Protector Zero Downtime Backup Integration Guide*.

- ***Data Protector Granular Recovery Extension for Microsoft SharePoint Server***

  Use this extension to recover individual website items, such as a Calendar or Task items, or documents. The Data Protector Granular Recovery Extension for Microsoft SharePoint Server is a specialized extension that tightly integrates in to Microsoft Office SharePoint Server and provides you detailed control over what is recovered. The extension does not provide any backup solution but instead depends on the Data Protector SharePoint integrations for the backup.

  See the *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*.

NOTE: You can also back up Microsoft SharePoint Server components and the underlying Microsoft SQL databases using common Data Protector file system backup functionality, which operates on the file level. Since you can only ensure data consistency by taking the Microsoft SharePoint Server offline, none of the benefits of the Data Protector integrations are available in this case.

Table 8 Data Protector backup solutions for Microsoft Office SharePoint Server 2007

		SharePoint Server 2007/2010 VSS based solution	SharePoint Server 2007/2010 integration	Granular Recovery Extension
Backup features				
Granularity	Granularity level	Databases	Components[1]	N/A (only a restore solution, no backup)
	Backup of search components	Yes	Yes	
	Backup of the SSO (single sign-on) database	Yes	Yes	
Backup types	Full backup	Yes	Yes	
	Incremental backup (transaction log)	No	Yes	
	Differential backup	No	Yes	
Zero downtime backup (ZDB) supported		Yes	No	
Restore features				
Granularity	Granularity level	Databases	Components	Single item (document)
	Restore of search components	Yes	Yes	

Table 8 Data Protector backup solutions for Microsoft Office SharePoint Server 2007 *(continued)*

		SharePoint Server 2007/2010 VSS based solution	SharePoint Server 2007/2010 integration	Granular Recovery Extension
	Restore of the SSO (single sign-on) database	Yes	Yes	
	Restore of an entire farm	Yes	Yes	
	Individual client restore	No	Yes	
	Individual service restore	No	Yes	
	Restore of the configuration database	Yes	Yes	
	Individual content database restore	Yes	Yes	
Recovery to a point in time		No	Yes	Depends on the solution used for backup
Recovery to the latest state		No	Yes	
Restore to a new location	Restore to another farm	No	Yes	
	Restore as a new service	No	Yes	
	Restore into …	Yes	Yes	
Instant recovery supported		Yes	No	

[1] Web applications, individual content databases, search components, the configuration database, the Central Administration database, and the single sign-on database

Table 9 Data Protector backup solutions for Microsoft SharePoint Server 2010

		SharePoint Server 2007/2010 VSS based solution	SharePoint Server 2007/2010 integration	Granular Recovery Extension
Backup features				
Granularity	Granularity level	Databases	Components [1]	N/A (only a restore solution, no backup)
	Backup of search components	Yes	No	
	Backup of Secure Store Service	Yes	Yes	
Backup types	Full backup	Yes	Yes	
	Incremental backup (transaction log)	No	Yes	
	Differential backup	No	Yes	
Zero downtime backup (ZDB) supported		Yes	No	
Restore features				
Granularity	Granularity level	Databases	Components	Single item (document)
	Restore of search components	Yes	N/A	
	Restore of the Secure Store Service	Yes	Yes	
	Restore of an entire farm	Yes	No	
	Individual client restore	No	Yes	
	Individual service restore	No	N/A	
	Restore of the configuration database	Yes	Yes	
	Individual content database restore	Yes	Yes	

Table 9 Data Protector backup solutions for Microsoft SharePoint Server 2010 *(continued)*

		SharePoint Server 2007/2010 VSS based solution	SharePoint Server 2007/2010 integration	Granular Recovery Extension
Recovery to a point in time		No	Yes	Depends on the solution used for backup
Recovery to the latest state		No	Yes	
Restore to a new location	Restore to another farm	No	Yes	
	Restore as a new service	No	Yes	
	Restore into …	Yes	Yes	
Instant recovery supported		Yes	No	
SharePoint Server 2010 specific features				
FAST Search support		Yes	No	N/A

[1] Web applications, individual content databases, the configuration database, the Central Administration database, and the Secure Store Service

Combining different integrations to back up and restore Microsoft SharePoint Server 2010

Use the Data Protector Microsoft SharePoint Server 2007/2010 integration (**VDI based integration**) and the Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution (**VSS based solution**) as follows:

Backup

- Back up Web applications and the associated content databases using the VDI based integration.
- Back up all other components using the VSS based solution.
  1. Create the VSS based solution backup specifications using the Data Protector PowerShell command.
  2. Open the newly created backup specifications and exclude all the Web application content databases backed up with the VDI based integration from the VSS based solution backup specifications.

     NOTE: The names of the content databases are the same as in the VDI based integration backup specification.

  3. Start the backup specifications using the Data Protector PowerShell command on the client with the Data Protector Microsoft SharePoint Server 2007/2010 integration agent installed.

TIP: To start the VSS based solution backup sessions, specify the post-exec option in the VDI based integration backup specification.

IMPORTANT: If you add a new Web application, a new Microsoft SQL Server system, or you delete a content database, and so on, update the corresponding backup specifications accordingly.

Restore

- Perform disaster recovery using the VSS based solution. Ensure that the new configuration matches the original.
- Perform a restore of Web applications using the VDI based integration. A restore to a new location is possible.

2 Data Protector Microsoft SharePoint Server 2007/2010 integration

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Office SharePoint Server 2007 integration and the Data Protector Microsoft SharePoint Server 2010 integration (from now on, both integrations are called **Microsoft SharePoint Server 2007/2010 integration**, unless differences are pointed out). It describes concepts and methods you need to understand to back up and restore the following Microsoft SharePoint Server 2007/2010 objects (**objects**):

- The configuration database
- The Central Administration content database
- Web applications
- Search components (Microsoft Office SharePoint Server 2007 only):
  - Shared Services Provider (**SSP**)
  - Windows SharePoint Services (**WSS**) Help Search

  NOTE: Backup and restore of the Microsoft SharePoint Server 2010 Search components, SharePoint Service Applications (**SSA**) and SharePoint Foundation Help Search, are not supported by the Microsoft SharePoint Server 2007/2010 integration.

- Single sign-on (**SSO**) database (Microsoft Office SharePoint Server 2007 only)

  NOTE: Backup and restore of the Microsoft SharePoint Server 2010 Single sign-on database are not supported by the Microsoft SharePoint Server 2007/2010 integration.

Farms of arbitrary sizes, from a single-system to multi-system, are supported.

Backup

Data Protector integrates with Microsoft SharePoint Server 2007/2010 to back up objects online. During backup, the Microsoft SharePoint Server 2007/2010 and Microsoft SQL Server instances can be actively used (**online backup**).

You can run interactive and scheduled backups of the following types:

- Full
- Differential
- Incremental

For details on the backup types, see "Backup types" (page 61).

Restore

During restore, each object can be restored:

- To the latest state or to a certain point in time
- To the original location or to a new location

More specifically:

- Web applications can be restored:
  - Under a different name
  - To a different URL
- Content databases (Web application databases, SSP databases, SSO database) can be restored[1]:
  - To a different Microsoft SQL Server client
  - To a different Microsoft SQL Server instance
  - Under a different name
  - To a different directory
- SSP sites can be restored:
  - Under a different name
  - To a different Web application URL
  - To a different My sites web application URL
- SSP index files can be restored:
  - To a different Microsoft SharePoint Server client
  - To a different directory

This chapter provides information specific to the Microsoft SharePoint Server 2007/2010 integration. For limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HP Data Protector Help*.

Integration concepts

Data Protector integrates with Microsoft SharePoint Server 2007/2010 through the Data Protector Microsoft SharePoint Server 2007/2010 integration agent (`sharepoint_bar.exe`), which channels communication between the Data Protector Session Manager and the clients in the Microsoft SharePoint Server 2007/2010 environment. The Data Protector Microsoft SharePoint Server 2007/2010 integration agent uses the Data Protector Microsoft SQL Server integration agent for backup of SQL databases and the data movement agent (DMA) for backup of index files.

Whether your Microsoft SharePoint Server 2007/2010 environment consists of a single system or multiple systems (small, medium, or large farm), the architecture of the integration is basically the same.

"Microsoft SharePoint Server 2007/2010 integration" (page 57) shows how Data Protector integrates with a medium farm.

1. The configuration database and Central Administration content database can only be restored to the original location under the same name.

Figure 23 Microsoft SharePoint Server 2007/2010 integration

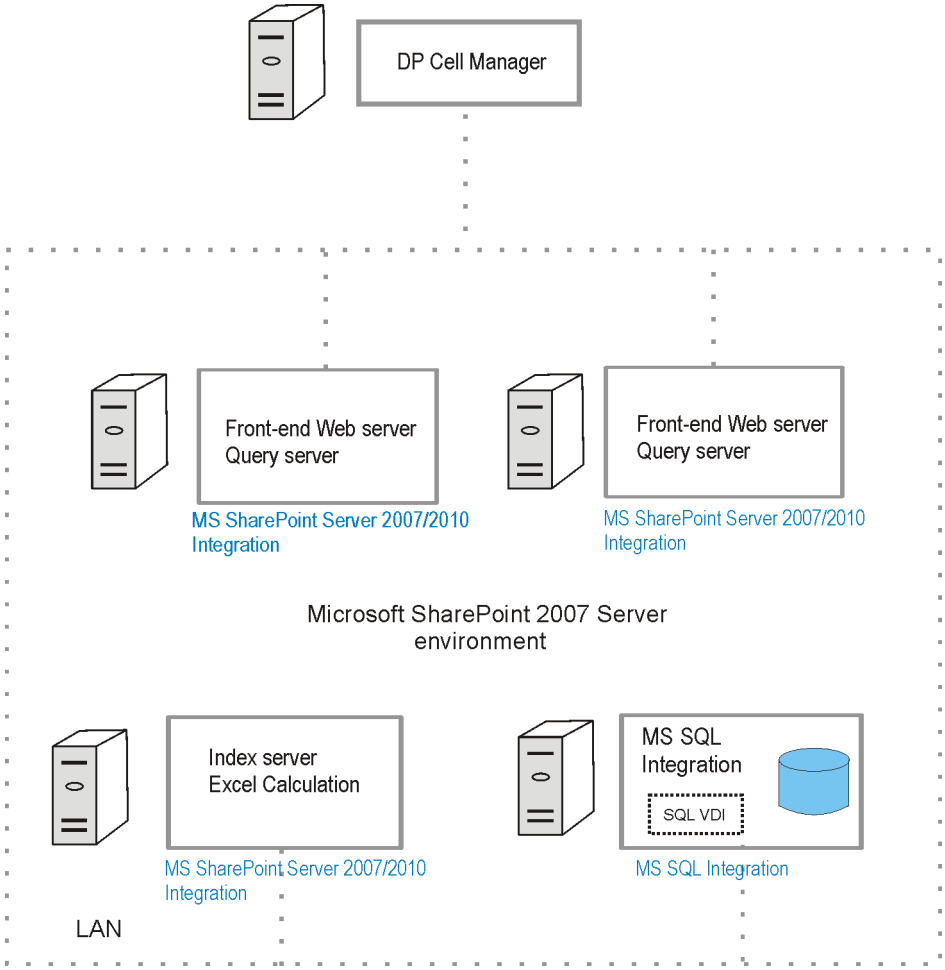

Table 10 Legend

MS SharePoint Server 2007/2010 Integration	A set of Data Protector executables that enables data transfer between Microsoft SharePoint 2007/2010 Server and Data Protector media
MS SQL Integration	A set of Data Protector executables that enables data transfer between Microsoft SQL Server and Data Protector media
SQL VDI	Microsoft SQL Server virtual device interface, through which Microsoft SQL Server and Data Protector exchange control and data
LAN	Local Area Network

Table 11 (page 57) briefly describes the Microsoft SharePoint Server 2007/2010 objects that you can back up and restore using the Data Protector Microsoft SharePoint Server 2007/2010 integration.

Table 11 Microsoft SharePoint Server 2007/2010 objects

Microsoft SharePoint Server 2007/2010 object	Description
Configuration database and Central Administration content database	The configuration database is a Microsoft SQL Server database that contains a configuration for an entire farm. The database itself resides on one Microsoft SQL Server system in the farm.

Table 11 Microsoft SharePoint Server 2007/2010 objects *(continued)*

Microsoft SharePoint Server 2007/2010 object	Description
	The Central Administration content database is a Microsoft SQL Server database that contains content for the Central Administration web application. The database resides on one Microsoft SQL Server system in the farm.
Content databases (Web application databases, SSP databases)	A Microsoft SQL Server database that stores content for a Web application. Each Web application can have one or more content databases. The content database contains content and metadata associated with site collections and sites/webs.
Web application	An entry point for individual sites, which hosts user content. A farm can have many Web applications.
SSP (Microsoft Office SharePoint Server 2007 only)	A search component which provides search and indexing services for user content in Web applications. Shared Services Provider (SSP) is a logical environment or a layer that contains all services you want to make available across your Web applications and sites. SSP provides services such as searching user profiles, site search, Excel services and audience. SSP has its own Microsoft SQL database to store all configuration data.
SSP index files	A folder that stores files. Each file is associated with user-defined information.
Windows SharePoint Services Help Search (Microsoft Office SharePoint Server 2007 only)	A search component which provides search capabilities of SharePoint help system.
Single sign-on database (Microsoft Office SharePoint Server 2007 only)	An 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.

Configuring the integration

Prerequisites

- Ensure that you have correctly installed and configured the Microsoft SharePoint Server 2007/2010 environment.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
  - For information on installing, configuring, and using Microsoft SharePoint Server 2007/2010, see the Microsoft SharePoint Server 2007/2010 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 Server 2007/2010 Integration` – on Microsoft SharePoint Server 2007/2010 systems (Microsoft SQL Server systems are excluded)
  - `MS SQL Integration` – on Microsoft SQL Server systems

  NOTE: If a system has both the Microsoft SQL Server and Microsoft SharePoint Server installed, install both Data Protector components on it.

Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the *HP Data Protector Help* index: “configuring devices” and “creating media pools”.
- To test whether a Microsoft SharePoint Server 2007/2010 and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every client in the farm. See the *HP Data Protector Help* for instructions.

Configuring user accounts

Backup and restore sessions are started by the `Data Protector Inet` service, which by default runs under the Windows local user account `SYSTEM`.

However, you must specify that the `Data Protector Inet` service starts the sessions under the Microsoft SharePoint Server 2007/2010 farm administrator Windows domain user account.

NOTE: When restoring the configuration database in the Microsoft SharePoint Server 2010 environment, the Data Protector Microsoft SharePoint Server 2007/2010 integration agent automatically uses the predefined credentials, User *PASSPHRASE* and Group *MSSPS*, saved in a Windows Registry.

Configure the user account as follows:

1. Ensure that the Microsoft SharePoint Server 2007/2010 farm administrator has been assigned the Windows local security policy user right **Replace a process level token**.
2. Add the Microsoft SharePoint Server 2007/2010 farm administrator to the Data Protector `admin` or `operator` user group. For details on adding users, see the *HP Data Protector Help* index: ”adding users”.
3. Save the Microsoft SharePoint Server 2007/2010 farm administrator and its password to a Windows Registry on all Microsoft SharePoint Server 2007/2010 systems and on all Microsoft SQL Server systems.

   NOTE: To restore the configuration database in the Microsoft SharePoint Server 2010 environment, save the predefined credentials, User *PASSPHRASE* and Group *MSSPS* to a Windows Registry on all Microsoft SharePoint Server 2007/2010 systems and on all Microsoft SQL Server systems.

   To save the user account, use:

   - The Data Protector GUI.

     For details, see the *HP Data Protector Help*.
   - The Data Protector CLI, by using `omniinetpasswd` or `omnicc` command.

     For details, see the `omnicc` and the `omniinetpasswd` man pages, or the *HP Data Protector Command Line Interface Reference*.

   NOTE: The `Data Protector Inet` service will start the session under this user account.

Examples

To save the user `jane` from the domain `HP` and with the password `mysecret` to a Windows Registry on all clients in the farm, log on to the Cell Manager and, execute:

```
omnicc -impersonation -add_user -user jane@HP -host Client1 -host
Client2 -host Client3 -passwd mysecret
```

Backup

You can back up the following Microsoft SharePoint Server 2007/2010 objects:

- The configuration database
- The Central Administration content database
- Web applications
- Search components (Microsoft Office SharePoint Server 2007 only):
  - Shared Services Provider (**SSP**)
  - Windows SharePoint Services (**WSS**) Help Search

NOTE: Backup of the Microsoft SharePoint Server 2010 Search components is not supported by the Microsoft SharePoint Server 2007/2010 integration.

- Single sign-on database (Microsoft Office SharePoint Server 2007 only)

NOTE: Backup of the Microsoft SharePoint Server 2010 Single sign-on database is not supported by the Microsoft SharePoint Server 2007/2010 integration.

Backup concepts

Before backing up Microsoft SharePoint Server 2007/2010 objects, you should consider the following specifics of each component.

- **Web application**

  The Data Protector Microsoft SharePoint Server 2007/2010 integration agent uses the Data Protector Microsoft SQL Server integration agent for backup of Web application content databases. Full, Differential and Incremental (transaction log) backup types are supported using the capabilities of the Data Protector Microsoft SQL Server integration agent. Web application settings are also backed up to simplify a restart of the service in case of a redirected restore or a disaster recovery.

- **Search components**

  This section is applicable only to Microsoft Office SharePoint Server 2007. Backup of Search components includes a backup of index files and associated Microsoft SQL Server databases. To ensure data consistency of the Search components, active crawlings must be paused and starting of new crawlings disabled during the backup. All parts of the Search components, all index files and their databases, must be backed up together. Individual parts of the Search components cannot be selected. Full, Differential and Incremental backup types are supported. Differential and Incremental backups use a timestamp strategy for the index files and the Data Protector Microsoft SQL Server integration agent native capabilities for the associated Microsoft SQL Server databases.

  The Data Protector Microsoft SharePoint Server 2007/2010 integration agent uses the Data Protector Microsoft SQL Server integration agent for backup of Microsoft SQL Server databases and the data movement agent (DMA) for backup of index files.

- **Configuration and Central Administration database**

  The configuration database and the Central Administration content database are synchronized and must be backed up together.

- **Single sign-on database**

  This section is applicable only to Microsoft Office SharePoint Server 2007. The Data Protector Microsoft SharePoint Server 2007/2010 integration supports a backup of an SSO database only. An associated encryption key can only be backed up through the Microsoft SharePoint

Server 2007/2010 user interface. You are notified with a warning message that the encryption key must be backed up manually on a specified Microsoft SharePoint Server 2007/2010 client. The Data Protector cannot disable or track the status of a re-encryption. Ensure that the SSO re-encryption is not running during the backup. For details of how to back up the encryption key, see the Microsoft SharePoint Server 2007/2010 documentation.

Backup types

The integration provides online backups of the following types:

Table 12 Backup types

Full	*Microsoft SQL Server database:* Performs a Microsoft SQL Server Full database backup; the complete database is backed up. *Index files:* Performs a Full filesystem backup of all index files.
Incremental	*Microsoft SQL Server database:* Performs a Microsoft SQL Server transaction log backup. Backs up transaction logs (`.log`) that have been created since the last transaction log backup of the Microsoft SQL Server database, and then truncates the transaction logs. *Index files:* Backs up only the index files that have been changed or created since the last backup of any type. **NOTE:** If the Microsoft SQL Server database is in the simple recovery mode (has no transaction logs), a Differential backup will be performed for the database instead. For the metadata[1] of the Microsoft SharePoint Server 2007/2010 components, a Full backup is always performed due to a small amount of data.
Differential	*Microsoft SQL Server database:* Performs a Microsoft SQL Server Differential backup of the database; backs up changes made to the database since the last Full backup. *Index files:* Backs up the index files that have been changed since the last Full backup.

[1] Metadata is defined as data providing information about backup of one or more Microsoft SharePoint Server 2007/2010 components. It can be stored and managed in a database, often called a registry or repository.

For details on the Microsoft SQL Server backup types, see the Microsoft SQL Server documentation.

NOTE: An Incremental or Differential backup cannot be performed, if a Full backup has not been performed.

Creating backup specifications

Create a backup specification using the Data Protector GUI (**Data Protector Manager**).

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SharePoint Server 2007/2010**, and click **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. Specify the Microsoft SharePoint Server 2007/2010 farm administrator **User name** and **Group/Domain name** under which a backup session should be performed.

   In **Client**, select any Microsoft SharePoint Server 2007/2010 system. The **Client** drop-down list contains all clients that have the Data Protector `MS SharePoint Server 2007/2010 Integration` component installed.

   Backup is started on the client that you specify here.

   Application database is selected automatically (by Microsoft SharePoint Server 2007/2010 integration).

   NOTE: The application database is equal to Microsoft SharePoint Server 2007/2010 configuration database

Click **Next**.

Figure 24 Selecting a client

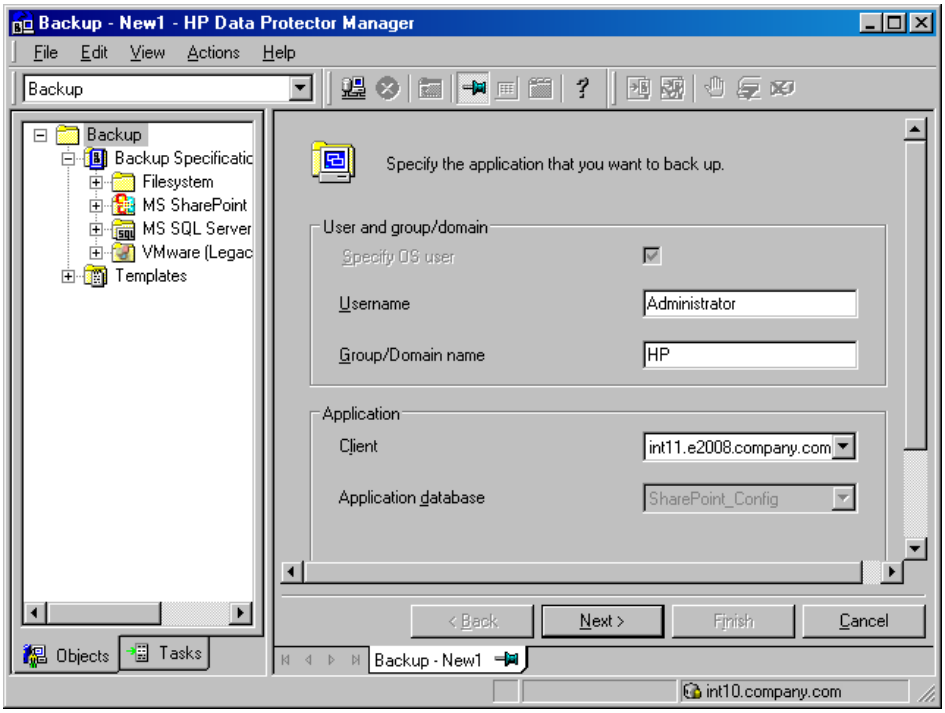

5. Select which objects to back up.

NOTE: If there are no components displayed, ensure that the user name and the domain name specified in step 4 are correct.

Figure 25 Selecting objects

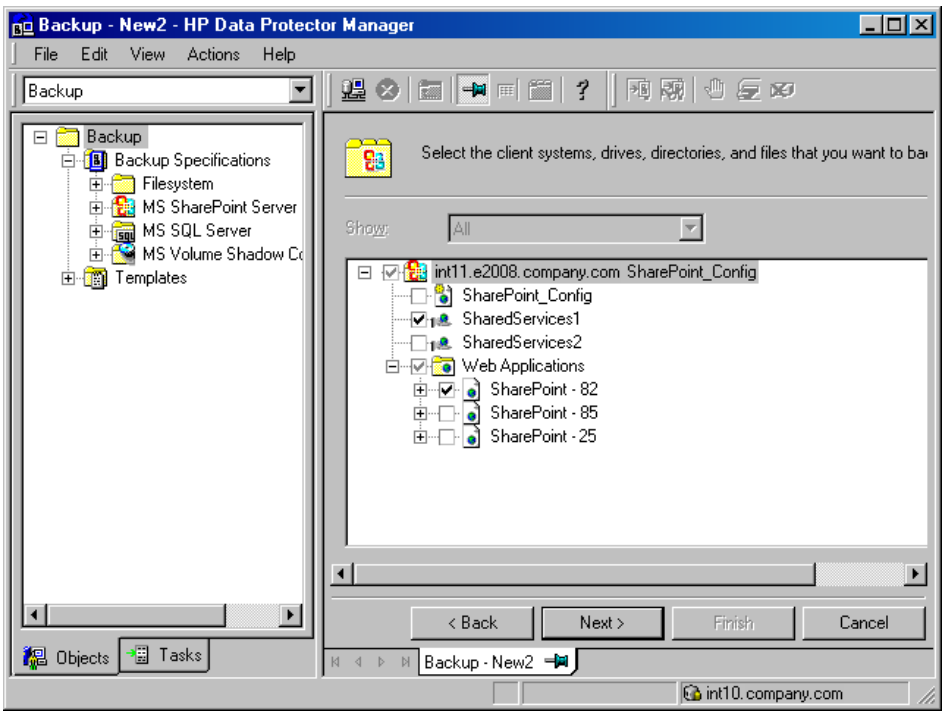

Click **Next**.

6. Select which devices to use for the backup.
   To specify device options, right-click the device and click **Properties**.
   Click **Next**.
7. Set backup options.
   For information on application-specific backup options, see "Application-specific backup options" (page 63).

Figure 26 Application-specific options

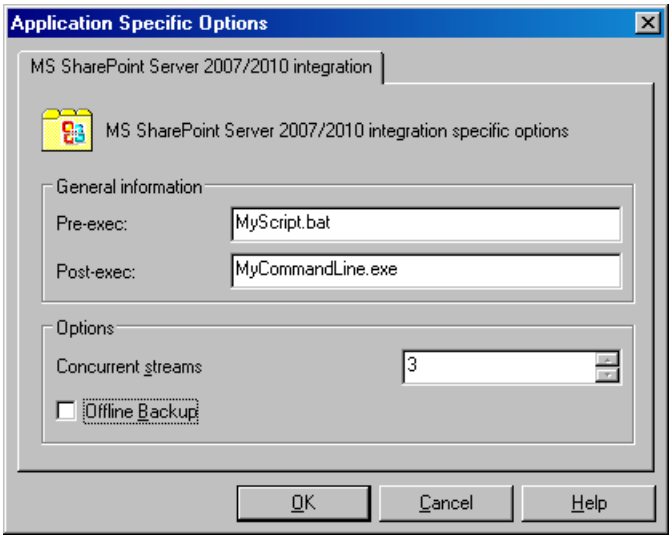

   Click **Next**.
8. Optionally, schedule the backup. See "Scheduling backup sessions" (page 64).
   Click **Next**.
9. Save the backup specification, specifying a name and a backup specification group.

TIP: Preview your backup specification before using it for real. See "Previewing backup sessions" (page 64).

Table 13 Application-specific backup options

Options	Description
Pre-exec, **Post-exec**	Specifies which command line to execute before (`pre-exec`) or after (`post-exec`) the backup. The command line is executed on the Microsoft SharePoint Server 2007/2010 system on which the backup session is started (that is the system on which the Data Protector Microsoft SharePoint Server 2007/2010 integration agent (`sharepoint_bar.exe`) is started). Type only the name of the command and ensure that the command is located in the *`Data_Protector_home`*`\bin` directory on the same system. Do not use double quotes.
Concurrent streams	Specifies how many parallel backup streams are used to back up Microsoft SQL Server databases. NOTE: Each Microsoft SQL Server database can be backed up in a separate backup stream. A maximum value which can be specified should equal to a number of devices that are selected for backup. If you change the number of devices, make sure to change the concurrency option as well.
Offline backup	Stops the Microsoft SharePoint Server 2007/2010 farm before starting a backup.

Table 13 Application-specific backup options *(continued)*

Options	Description
	NOTE: If selected, it enables you to avoid the restore limitation. For details, see "Backup concepts" (page 60).

Modifying backup specifications

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

Scheduling backup sessions

You can schedule a backup session to start automatically at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: "scheduled backups".

Scheduling example

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

1. In the **Schedule** property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See "Scheduling backup sessions" (page 64). Under **Session options**, select **Differential** from the **Backup type** drop-down list.

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

Figure 27 Scheduling backup sessions

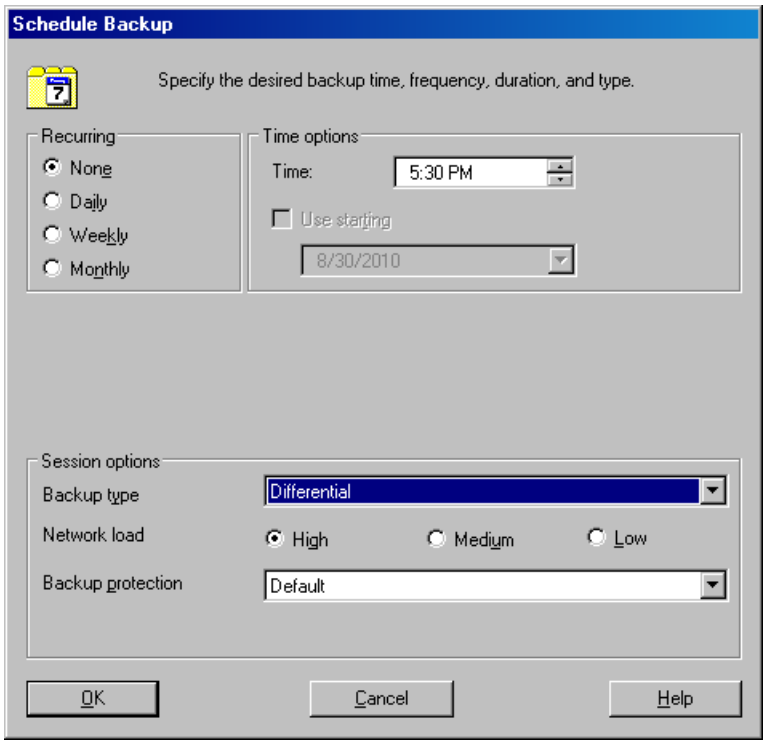

Previewing backup sessions

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 2007/2010**. 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

1. Log in to the Cell Manager or to any client with the Data Protector `User Interface` component installed under a user account that is configured as described in "Configuring the integration" (page 58).
2. Execute:

   ```
   omnib -mssharepoint_list BackupSpecificationName -test_bar
   ```

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

What happens during the preview?

The following are tested:

- Communication between the Microsoft SharePoint Server 2007/2010 system on which the backup session is started and the Data Protector Cell Manager
- If devices are correctly specified
- If necessary media are in the devices
- The syntax of the backup specification

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 Microsoft SharePoint Server 2007/2010 and Microsoft 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 2007/2010**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

1. Log in to the Cell Manager or to any client with the Data Protector `User Interface` component installed under a user account that is configured as described in "Configuring the integration" (page 58).

2. Execute:

   `omnib -mssharepoint_list` *`BackupSpecificationName`* `[-barmode` *`MSSharePointMode`*`] [`*`ListOptions`*`]`

   where *`MSSharePointMode`* is one of the following backup types:

   `full|diff|incr`

   If the `-barmode` option is not specified, a Full backup is performed.

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

   ### Examples

   To start a Full backup using the backup specification `myBackup`, execute:

   ```
   omnib -mssharepoint_list myBackup -barmode full
   ```

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

   ```
   omnib -mssharepoint_list myBackup -barmode diff
   ```

Preparing for disaster recovery

To be able to perform a disaster recovery, back up the following Microsoft SharePoint Server 2007/2010 objects:

Table 14 What must be backed up

Object	How to back up
Microsoft SharePoint Server 2007/2010 content databases	Back up the databases using the Data Protector Microsoft SharePoint Server 2007/2010 integration backup as described in this chapter.
Microsoft SQL Server configuration	Back up the master databases using the Data Protector Microsoft SQL Server integration backup. For details of how to recover the master database, see the disaster recovery subsection of the restore section in the Microsoft SQL Server chapter of this guide.
Encryption key[1]	Back up the encryption key as described in the Microsoft Office SharePoint Server 2007 documentation. NOTE: Applicable only to Microsoft Office SharePoint Server 2007.

Customizations
(from all the front-end Web server clients)
	Back up the customizations using the Data Protector filesystem backup.
	Normally, customization files are located in the following directories:
	• 12 Hive (Microsoft Office SharePoint Server 2007 integration):
	`Program Files\Common Files\Microsoft Shared\Web server extensions\12`
	• 14 Hive (Microsoft SharePoint Server 2010 integration):
	`Program Files\Common Files\Microsoft Shared\Web server extensions\14`
	• Internet Information Services (IIS) Virtual Directories:
	`\Inetpub\wwwroot\wss\VirtualDirectories`
	Contact the customization vendor to determine where exactly in the filesystem the customization files are located.
	For details of how to perform a filesystem backup, see the *HP Data Protector Help*.

Table 14 What must be backed up *(continued)*

Object	How to back up
	Alternatively, if the customizations are packed as solutions, these solution packages can be used for manual re-deployment.
IIS database (from all the front-end Web server clients)	Back up the database using the Data Protector filesystem backup. The IIS database is located in the client CONFIGURATION. For details of how to perform a filesystem backup, see the *HP Data Protector Help*.

[1] if it is used for the SSO service

Restore

You can restore Microsoft SharePoint Server 2007/2010 objects using the Data Protector GUI or CLI.

Restore concepts

Before restoring Microsoft SharePoint Server 2007/2010 objects, you should consider the following specifics of each component.

- **Web application**

  When restoring a Web application, you can select the entire Web application or individual content databases. Both can be restored to a new location. For details, see "Restore options" (page 75). If restored together, the Data Protector Microsoft SharePoint Server 2007/2010 integration reconnects the Web applications and their content databases after the restore to the new location. The restored content of the Web application needs to be re-crawled to become searchable.

- **Search components**

  Applicable only to Microsoft Office SharePoint Server 2007. Individual Search components (Shared Services Provider, Windows SharePoint Services Help Search) can be restored. To ensure data consistency, individual Search components must be restored in their completeness, not by their subcomponents (by index files or by associated Microsoft SQL Server databases).

  NOTE: You can view the Microsoft SharePoint Server objects, including the Search components, by Component or by Server in a specified time interval. Depending on the restore view type, individual parts of a Search service cannot be selected. For example, a Search database only cannot be selected for restore.

  After the restore, the Search components will be automatically resumed by the Data Protector Microsoft SharePoint Server 2007/2010 integration agent. To properly restore an SSP Index by Server View, ensure that the original index files are not accessible to any other Microsoft SharePoint Server service. All Microsoft SharePoint Windows services must be stopped. After the restore of index files, you need to clear the Microsoft SharePoint Server file system cache on all Microsoft SharePoint Server 2007/2010 clients in the farm on which Windows SharePoint Services Timer is running. Restart the previously stopped services. For details of how to clear the cache, see the Microsoft SharePoint Server 2007/2010 documentation (the Microsoft webpage: http://support.microsoft.com/kb/939308).

  NOTE:

  - Restore to a new location (redirected restore) is supported only for Shared Services Provider (SSP).
  - A SSP index name is changed during the restore process. Consequently, the next backup of the same SSP must be a Full backup.

- **Configuration and Central Administration database**

  The configuration database and the Central Administration content database contain a description of the state of the Microsoft SharePoint Server 2007/2010 farm, including client names. Consequently, only a restore to the original location is supported. To ensure data consistency, these databases must be restored together.

  NOTE: In case of a disaster recovery, the Data Protector Microsoft SharePoint Server 2007/2010 integration agent automatically disconnects all Microsoft SharePoint Server 2007/2010 clients, restores the databases and reconnects the farm clients to return the farm in a working condition. After the configuration database is restored, some of the Microsoft SharePoint Windows services on individual Microsoft SharePoint Server 2007/2010 clients remain disabled. You need to restart these services manually from a local client services console or by restoring an appropriate component; you are notified with a warning message. Microsoft SharePoint administration, timer and tracing services are started automatically.

- **Single sign-on database**

  This section is applicable only to the Microsoft Office SharePoint Server 2007. The Data Protector Microsoft SharePoint Server 2007/2010 integration supports a restore of SSO database only. An associated encryption key cannot be backed up by the Data Protector Microsoft SharePoint Server 2007/2010 integration agent and consequently cannot be restored. You are notified with a warning message that the encryption key must be restored manually on a specified Microsoft SharePoint Server 2007/2010 client. The Data Protector cannot disable or track a status of the re-encryption. Ensure that the SSO re-encryption is not running during the restore. For details, on how to restore the encryption key, see the Microsoft SharePoint Server 2007/2010 documentation.

  NOTE: A restore to a new location is supported.

Before you begin

- Ensure that the Microsoft SharePoint Server 2007/2010 and the Microsoft SQL Server instances are online and that the Microsoft SharePoint Server 2007/2010 services run under the Microsoft SharePoint Server 2007/2010 farm administrator account.
- If you plan to restore Microsoft SQL Server databases to another location:
  - Ensure that the destination Microsoft SQL Server system is part of the Microsoft SharePoint Server 2007/2010 environment and has the `MS SQL Integration` component installed.
  - Ensure that the destination Microsoft SQL Server instance exists, and is online.
- If you use the encryption key for the single sign-on service, note that the single sign-on database cannot be restored without the original encryption key.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects**, **MS SharePoint Server 2007/2010**, select the Microsoft SharePoint Server 2007/2010 client that served as an entry point to the Microsoft SharePoint Server 2007/2010 farm during backup, and then click **MS SharePoint Server 2007/2010 [Microsoft SharePoint Server 2007/2010]**.
3. In the **Source** page, select which Microsoft SharePoint Server 2007/2010 objects to restore. You can view the objects by **Component** or by **Server** in a specified time interval.

Figure 28 Selecting Microsoft SharePoint Server objects for restore (View by Component)

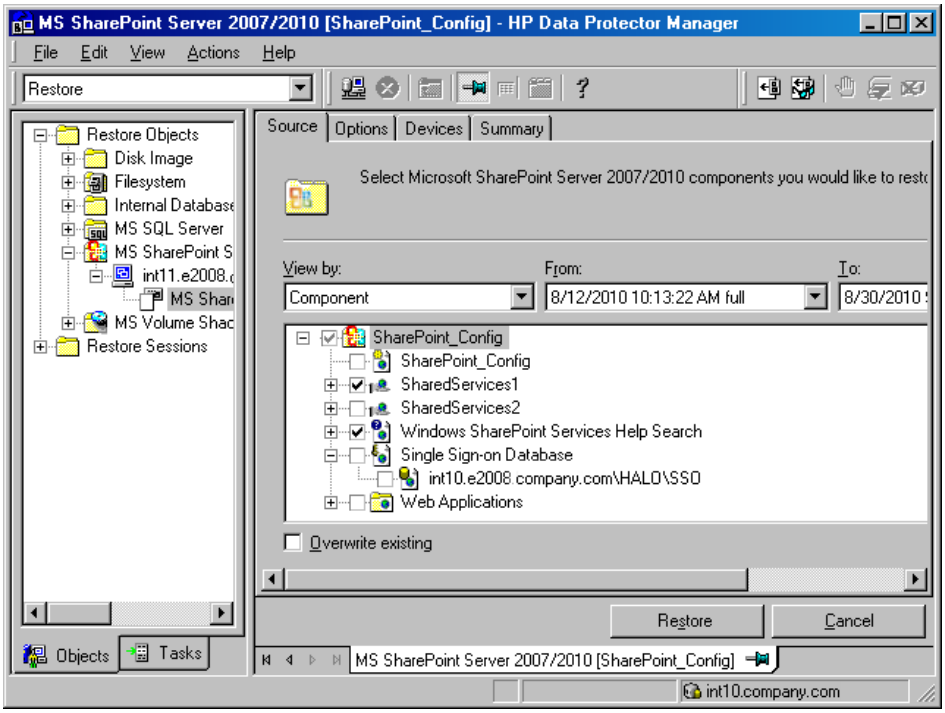

Figure 29 Selecting Microsoft SharePoint Server objects for restore (View by Server)

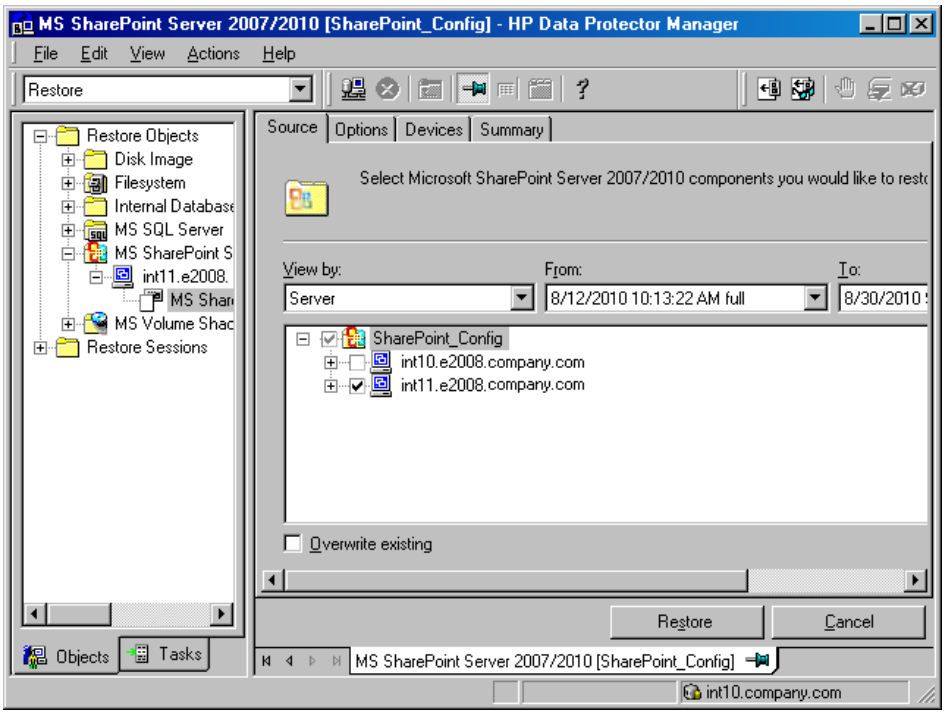

You can specify the restore destinations for each Microsoft SharePoint Server 2007/2010 object separately: right-click an object and click **Properties**. A Properties dialog box is displayed.

NOTE:

- The menu is available only if **Component** is selected in the **View by** drop-down list in the source page. The Properties dialog box of each component is pre-filled with the original data (names, locations, URLs).
- Applicable only if **Overwrite existing** is not selected.

  If **Overwrite existing** option is selected, the components are restored to the original location and with the same settings as when they were backed up. For details, see "Restore options" (page 75).

You can restore a Web application's settings under a different name or to a different URL. See "Specifying the restore destination for a Web application's settings" (page 70).

Figure 30 Specifying the restore destination for a Web application's settings

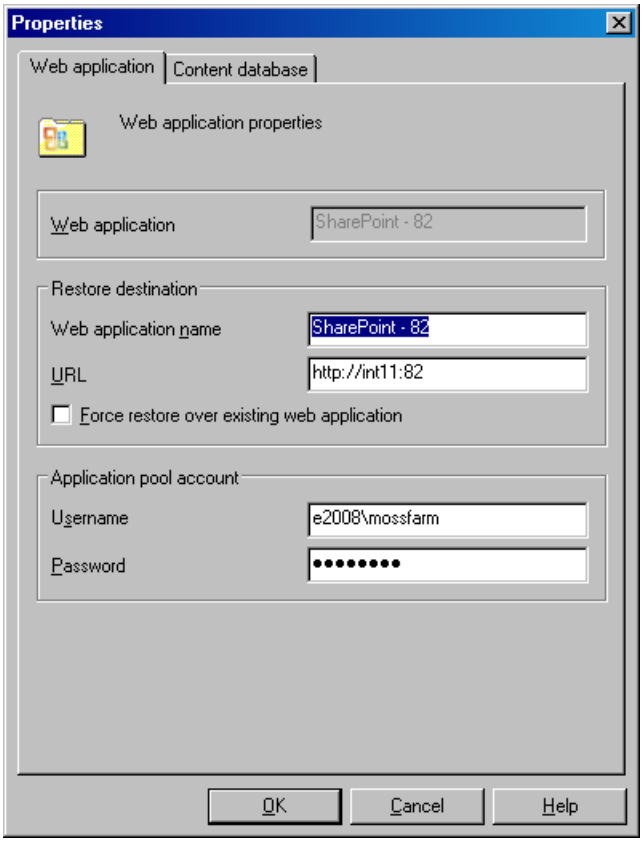

You can restore a content database to a different Microsoft SQL Server client, to a different Microsoft SQL Server instance, under a different name, or to a different directory. See "Specifying the restore destination for a Web application's content database" (page 71).

Figure 31 Specifying the restore destination for a Web application's content database

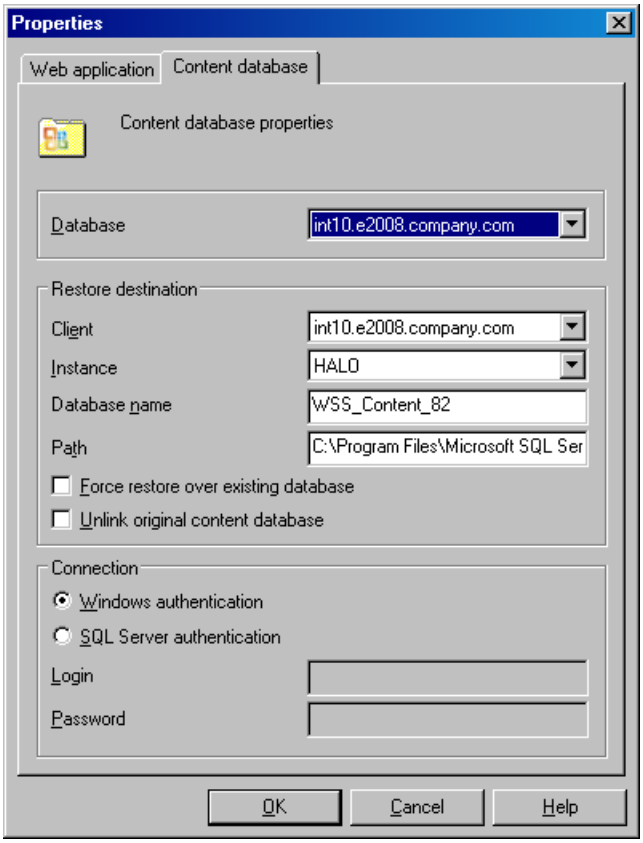

Figure 32 Specifying the restore destination for an SSO database

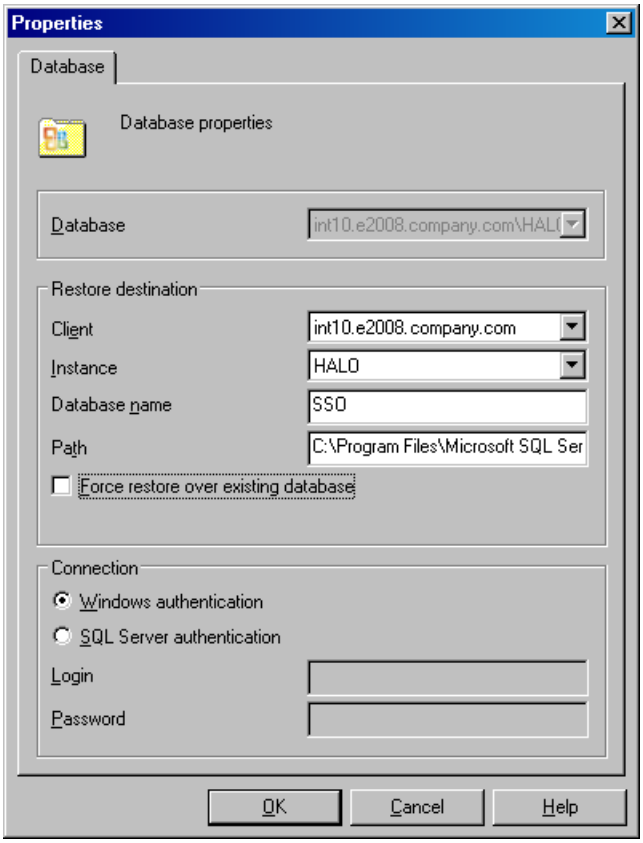

Figure 33 Specifying the restore destination for an SSP database

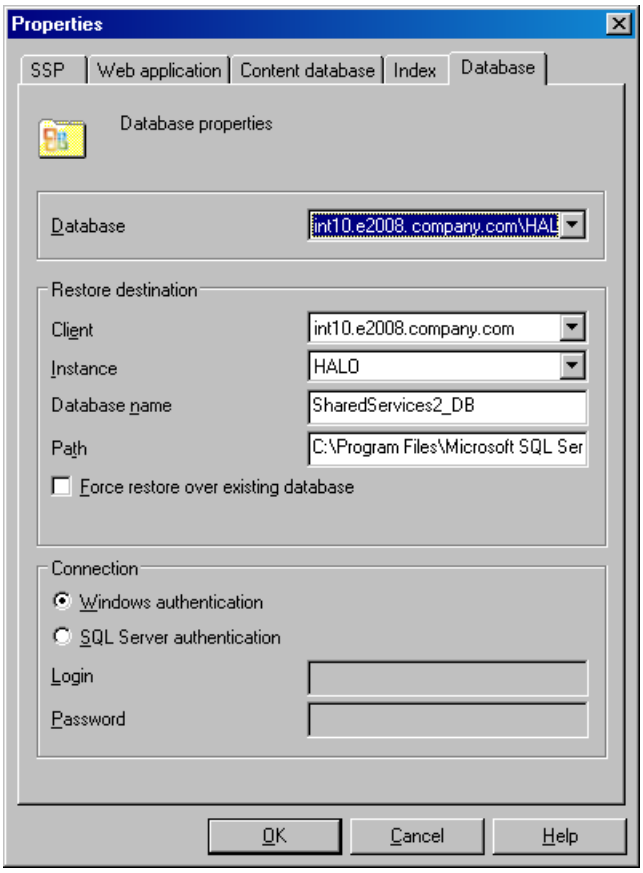

You can restore SSP sites under a different name, to a different Web application URL, or to a different My sites web application URL. See "Specifying the restore destination for an SSP" (page 73).

Figure 34 Specifying the restore destination for an SSP

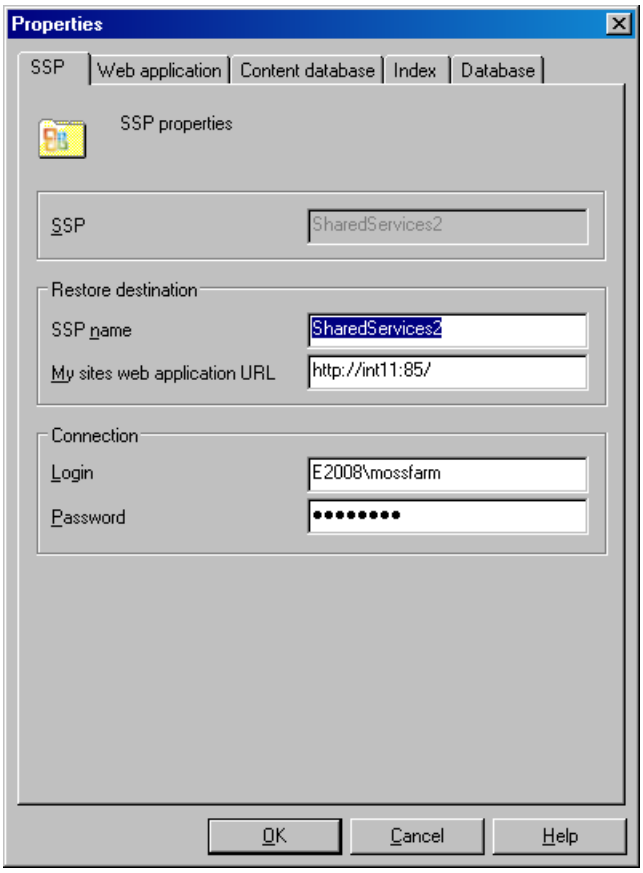

You can restore SSP index files to a different client or directory. See "Specifying the restore destination for an SSP index files" (page 74).

Figure 35 Specifying the restore destination for an SSP index files

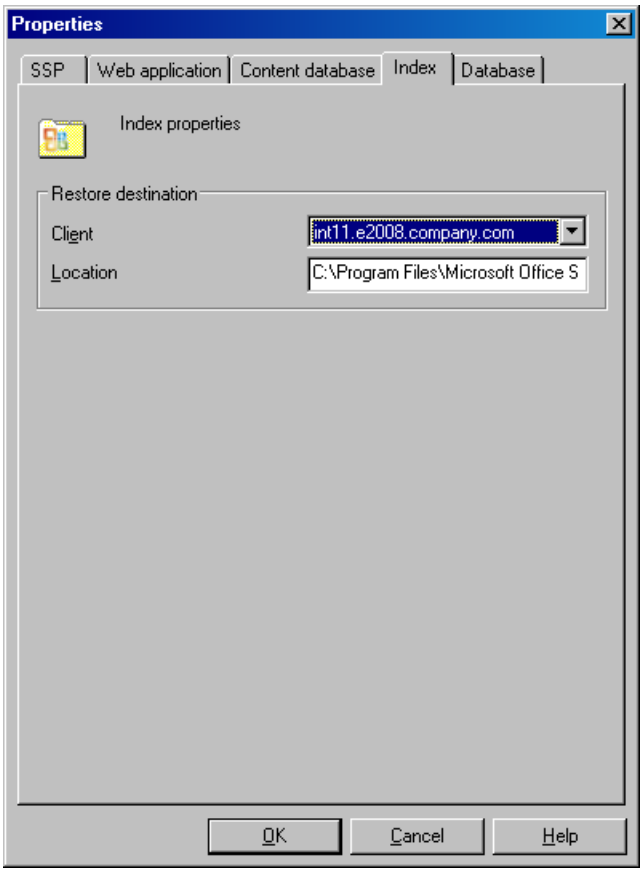

4. In the **Options** page, select the Microsoft SharePoint Server 2007/2010 specific restore options.

   You must specify the **Farm administrator user name** and **Farm administrator user group** options to perform a restore session under the Microsoft SharePoint Server 2007/2010 farm administrator Windows domain user account.

Figure 36 Restore options

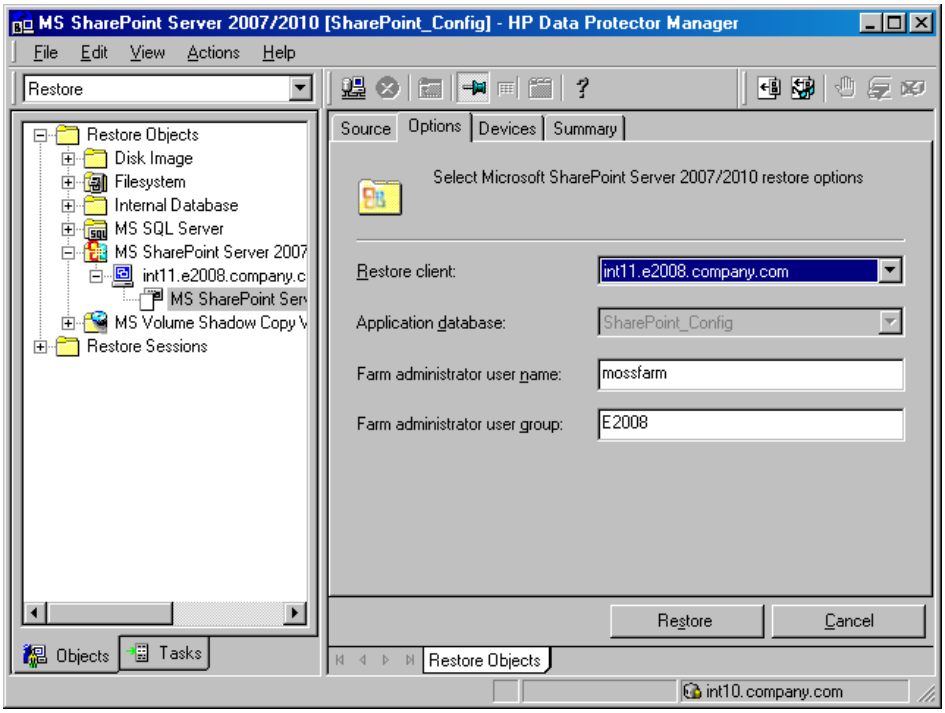

NOTE: When restoring the configuration database in the Microsoft SharePoint Server 2010 environment, the Data Protector Microsoft SharePoint Server 2007/2010 integration agent automatically uses the predefined credentials, User *PASSPHRASE* and Group *MSSPS*, saved in a Windows Registry.

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

   For more information on how to select devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".
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.

Restore options

Table 15 General restore options

Option GUI / CLI	Description
Restore client / `-destination`	Specifies the client on which the Data Protector Microsoft SharePoint Server 2007/2010 integration agent should be started. It also specifies to which farm the components are restored. The drop-down list contains all clients with the Data Protector Microsoft SharePoint Server 2007/2010 integration agent installed.
Application database	Shows the Microsoft SharePoint Server 2007/2010 configuration database name of the farm to which the selected client belongs.

Table 15 General restore options *(continued)*

Option GUI / CLI	Description
User name/User group / `-user`	Specifies the Windows domain user under which the Data Protector Microsoft SharePoint Server 2007/2010 integration agent should run. This user must be a farm administrator.
Overwrite existing / `-replace`	Overwrites all the existing redirection options specified for the selected components. A restore to the original location is performed.

Table 16 Web application options

Option GUI / CLI	Description
Web application / `-webapplication`	Shows the original Web application name.
Web application name / `-as`	Specifies the name under which the Web application should be restored.
URL / `-url`	Specifies the URL to which the Web application should be restored.
Force restore over existing web application / `-replace`	Overwrites the existing Web application residing at the target URL.
Username / `-poolusername` **Password** / `-poolpassword`	Specifies the Windows domain application pool user account under which the application pool should run. Note that each Web application has its own application pool.

Table 17 Web application – content database options

Option GUI / CLI	Description
Database / `-db`	Enables you to specify different options for different databases. The drop-down list contains the Web application databases that were backed up in the selected interval.
Client / `-tohost`	Specifies the Microsoft SQL Server client to which the database should be restored. The drop-down list contains the clients that have the `MS SQL Integration` component installed.
Instance / `-newinstance`	Specifies the Microsoft SQL Server instance to which the database should be restored. All created instances on the target client are listed.
Database name / `-as`	Specifies the name under which the database should be restored.
Path / `-todir`	Specifies the path to the directory to which the database files should be restored.
Force restore over existing database / `-replace`	Overwrites the existing database residing at the target Microsoft SQL Server instance. If a database with the same name as the one you are restoring already exists and has a different internal structure, Microsoft SQL Server does not let you rewrite the database unless you select this option.
Unlink original Content Database / `-unlink`	Removes the original content database from the farm. Available only when at least one of the original values of the restore redirection options is changed.

Table 17 Web application – content database options *(continued)*

Option GUI / CLI	Description
Windows or SQL authentication / `-sqllogin`	Specifies the authentication type which should be used to connect to the database.
Login and Password / `-sqlpassword`	Available only with the SQL authentication type selected. Specifies the Windows domain user account or the Microsoft SQL Server user account.

Table 18 SSP site options

Option GUI / CLI	Description
SSP / `-ssp`	Shows the original Shared Services Provider (Microsoft Office SharePoint Server 2007 only) name.
SSP name / `-as`	Specifies the name under which the Shared Services Provider should be restored.
Web application URL / `-url`	Specifies the URL of the Web application that should host the SSP administration webpage.
My sites web application URL / `-mysiteurl`	Specifies the URL of the Web application that should host personal sites and profiles.
Login / `-ssplogin`	Specifies the Windows domain user account under which the SSP timer job and web services should run.
Password / `-ssppassword`	Specifies a password for the login credential.

Table 19 SSP - index files options

Option GUI / CLI	Description
Client / `-tohost`	Specifies the Microsoft SharePoint Server 2007 client to which the index files of the selected SSP should be restored. The drop-down list contains all clients with the Data Protector `MS SharePoint Server 2007/2010 Integration` installed.
Location / `-todir`	Specifies the path to the directory to which the SSP index files should be restored.

Table 20 SSP - content database options

Option GUI / CLI	Description
Database / `-db`	See the description in "Web application – content database options" (page 76).
Client / `-tohost`	
Instance / `-newinstance`	
Database name / `-as`	
Path / `-todir`	
Force restore over existing database / `-replace`	
Windows or SQL Server authentication / `-sqllogin`	
Login and Password / `-sqlpassword`	

Table 21 SSO database options

Option GUI / CLI	Description
Database / `-db`	See the description in "Web application – content database options" (page 76).
Client / `-tohost`	
Instance / `-newinstance`	
Database name / `-as`	
Path / `-todir`	
Force restore over existing database / `-replace`	
Windows or SQL Server authentication / `-sqllogin`	
Login and Password / `-sqlpassword`	

Restoring using the Data Protector CLI

1. Log in to the Cell Manager or to any client with the Data Protector `User Interface` component installed under a user account that has been added to the Data Protector `admin` or `operator` user group.
2. Execute:

```
omnir -mssharepoint
 -barhost HostName
 [-destination RestoreClientName]
 -user User:Group
 [-session BackupID]
 [-replace]
 [-byserver ServerName [-byserver ServerName...]]
 -farmname FarmName
 [Component [Component...]]
 [GENERAL_OPTIONS]

Component
 -configdb |
 -webapplication WebApplicationName
     [WEB_APPLICATION_OPTIONS]
     [ContentDatabase [ContentDatabase...]] |
 -ssp SSPName [SSP_OPTIONS]
     [-index [INDEX_OPTIONS]]
     [Database [Database...]]
     [-webapp WebApplicationName
         [WEB_APPLICATION_OPTIONS]
         [ContentDatabase [ContentDatabase...]]] |
 -wsssearch [Database] |
 -ssodb [DB_OPTIONS]

ContentDatabase
-db DBName -host DBHostName [-unlink] [DB_OPTIONS]

Database
-db DBName -host DBHostName [DB_OPTIONS]

WEB_APPLICATION_OPTIONS
 -as WebApplicationName
 -url WebApplicationURL
 -poolusername Username [-poolpassword Password]
 -replace

DB_OPTIONS
 -sqllogin Username [-sqlpassword Password]
 -instance SourceInstanceName
 -as NewDBName
 -tohost DBHostName
 -newinstance DestinationInstanceName
 -todir NewDirectoryName
 -replace
```

```
SSP_OPTIONS
 -ssplogin Username [-ssppassword Password]
 -as SSPName
 -mysiteurl MySiteWebAppURL

INDEX_OPTIONS
 -tohost IndexServerHostName
 -todir NewDirectoryName
```

For a brief description of the options, see "Restore options" (page 75). For details, see the `omnir` man page or the *HP Data Protector Command Line Interface Reference*.

NOTE: A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the object *copy* session.

The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

Example

To restore a Web application content database from the latest session to another location, changing its name, a Microsoft SQL Server system, instance, and a data file path, execute:

```
omnir -mssharepoint -barhost wfe1.domain.com -webapplication "SharePoint
- 2224" -db "WSS_Content_2224" -as "WSS_new_DB" -tohost
mosssql2.domain.com -newinstance moss1 -todir "f:\program files\SQL\data"
```

Disaster recovery

Disaster recovery is a very complex process, involving different products from different vendors. Check the operating system and Microsoft SharePoint Server 2007/2010 instructions on how to prepare for it.

The following steps briefly describe the disaster recovery process:

1. Reinstall the operating system, the Microsoft SharePoint Server 2007/2010 environment and Microsoft SQL Server. Ensure that the configuration matches the original.
2. Install Data Protector in the newly configured environment.
3. Restore Microsoft SQL Server configuration, restoring the master database(s). For details, see the disaster recovery subsection of the restore section in the Microsoft SQL Server chapter of this guide.
4. Restore Microsoft SharePoint Server 2007/2010 databases from a Data Protector Microsoft SharePoint Server 2007/2010 integration backup as described in this chapter (at least the Configuration database and Central administration webpage content database).
5. Restore the IIS from a Data Protector filesystem backup (Windows CONFIGURATION – IIS database).

   For details of how to restore from a filesystem backup, see the *HP Data Protector Help*.
6. Restore the customizations from a Data Protector filesystem backup (or re-deploy manual solutions).

Monitoring sessions

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

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

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

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft SharePoint Server 2007/2010 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. On how to verify this, see the *HP Data Protector Help* index: "patches".
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://support.openview.hp.com/selfsolve/manuals for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- 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 the *HP Data Protector Help*.
- Check if your environment is set up correctly. To list the available systems and services, you can execute the `sharepoint_bar.exe -farmtree` command, which will list all servers that have persistent data that can be backed up. Unsupported services (FAST search, front-end Web Server system, and so on) are not listed.

Additionally, if your configuration or backup failed:

- Ensure that the Microsoft SharePoint Server 2007/2010 and Microsoft SQL Server instances are online.

Problems

Problem

Crawl status error occurs after the restore of the configuration database in the Microsoft SharePoint Server 2010 environment

The following message displays:

```
503 Service unavailable.
```

This happens because by disconnecting all Microsoft SharePoint Server 2007/2010 clients before the restore of the configuration database Search Service Application application pools are deleted from the IIS (Internet Information Services Virtual Directories / IIS database).

Action

Go to the Manage Service Applications in the Central Administration page and assign new application pools (the application pool for the Search Admin Web Service and the application pool for the Search Query and the Site Settings Web Service) to the Search Service Application.

Problem

Restore of the Shared Services Provider (SSP) fails with Session Manager aborting the session

During the restore of multiple SSPs, the Restore Session Manager aborts the session after 10 minutes.

The deletion of the SSPs can take longer than the Session Manager waits (default is 10 minutes) for the connection with clients.

Action

Set the global option `SmWaitForFirstRestoreClient` to the appropriate value or upgrade the resources of the farm clients.

Problem

Backup fails with "MS SQL integration not installed" message

When Microsoft SQL Server systems are configured with alias names and Microsoft SharePoint Server configuration uses the SQL Server system alias names, a backup session fails with an error similar to the following:

```
[Critical] From: OB2BAR_SHAREPOINT@Domain Database Time: Date Time 'MS
SQL' integration not installed on ''.
```

Action

1. Make sure that the Microsoft SQL Server Management Objects (SMO) is installed on each Microsoft SharePoint Server 2007/2010 system.
2. Install the Data Protector Microsoft SharePoint Server integration DPWIN_00574 or later patch on each Microsoft SharePoint Server 2007/2010 system.

Problem

Backup fails with "Required privilege not held by the client" message

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

```
[70:24] A system error occurred when starting the target script or an
agent module. The system error code reported is 1314 and the message
resolves to '[1314] A required privilege is not held by the client.'.
```

Action

1. Go to:

   Control Panel > Administrative Tools > Local Security Policy
2. Expand **Local Policies** and select **User Rights Assignment**.
3. The Windows domain user under which the Data Protector Inet service starts the backup (that is, the farm administrator) must be granted the **Replace a process level token** user right.

3 Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution

Introduction

This chapter explains how to configure and use the Data Protector Microsoft SharePoint Server 2007/2010 VSS based solution (**VSS based solution**). In reality, the solution is based on the Data Protector Microsoft Volume Shadow Copy Service integration (**VSS integration**). For details on the VSS integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

The chapter describes concepts and methods you need to understand to back up and restore Microsoft Office SharePoint Server 2007 and Microsoft SharePoint Server 2010 data that is stored in Microsoft SQL Server databases. For example:

- The configuration database (`SharePoint_Config`)
- Content databases (`SharePoint_AdminContent_Label`, `WSS_Content_Label`,...)
- Shared Services Provider databases (`SSP_DB`) (Microsoft Office SharePoint Server 2007)
- SharePoint Service Applications databases (`SSA_DB`) (Microsoft SharePoint Server 2010)
- Search databases (`SSP_Search_DB`)
- The Single Sign-On database (`SSO`)

In addition, you can also back up and restore Microsoft SharePoint Server search index files.

From now on, both Microsoft SharePoint Server versions are called **Microsoft SharePoint Server**, unless the differences are pointed out.

Backup

Microsoft SharePoint Server data that is stored in Microsoft SQL Server databases is backed up using one of the following Microsoft SQL Server VSS writers:

- `MSDE writer` (for Microsoft SQL Server 2000 databases)
- `SqlServerWriter` (for Microsoft SQL Server 2005/2008 databases)

Microsoft Office SharePoint Server 2007 search index files are backed up using the following VSS writers:

- `OSearch VSS writer`
- `SPSearch VSS writer`

Microsoft SharePoint Server 2010 search index files are backed up using the following VSS writers:

- `OSearch14 VSS writer`
- `SPSearch4 VSS writer`

Microsoft FAST Search Server 2010 search index files are backed up:

- using the Data Protector Disk Agent (standard filesystem backup with VSS enabled)

You can create and run backup specifications using the Data Protector PowerShell command which is described in "Backup" (page 85).

Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only `Full` backup type is supported.

Restore

Restore can be started using the Data Protector GUI or CLI as described in "Restore" (page 97).

Installation and configuration

Licensing

The Data Protector VSS based solution requires one online-extension license per each Microsoft SharePoint Server client participating in the backup and restore process. This means one online-extension license for each system on which the Data Protector `MS Volume Shadow Copy Integration` component is installed.

Installing the integration

For details on how to install a Data Protector cell, see the *HP Data Protector Installation and Licensing Guide*.

To be able to back up Microsoft SharePoint Server objects, install the following installation packages and Data Protector components:

- Service Pack 2 (Windows SharePoint Services 3.0 and Microsoft Office SharePoint Server 2007)
- Windows PowerShell 2.0 and the Data Protector `User Interface` component on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands and on which you install the Data Protector `MS Volume Shadow Copy Integration` component. See the next bullet.

  If not already available on your Windows system, you can download Windows PowerShell from the Microsoft webpage.
- The Data Protector `MS Volume Shadow Copy Integration` component on the Microsoft SQL Server system and the Microsoft SharePoint Server systems that have at least one of the following services enabled:

  Microsoft Office SharePoint Server 2007

  - `Windows SharePoint Services Database`
  - `Windows SharePoint Services Help Search`
  - `Office SharePoint Server Search`

  Microsoft SharePoint Server 2010

  - `SharePoint Foundation Database`
  - `SharePoint Foundation Help Search`
  - `SharePoint Server Search`
- The Data Protector `Disk Agent` component on each Microsoft FAST Search Server 2010 system for SharePoint (Microsoft SharePoint Server 2010)

Ensure that the `Volume Shadow Copy` service is started on all these clients.

Figure 37 Installing a medium farm (example)

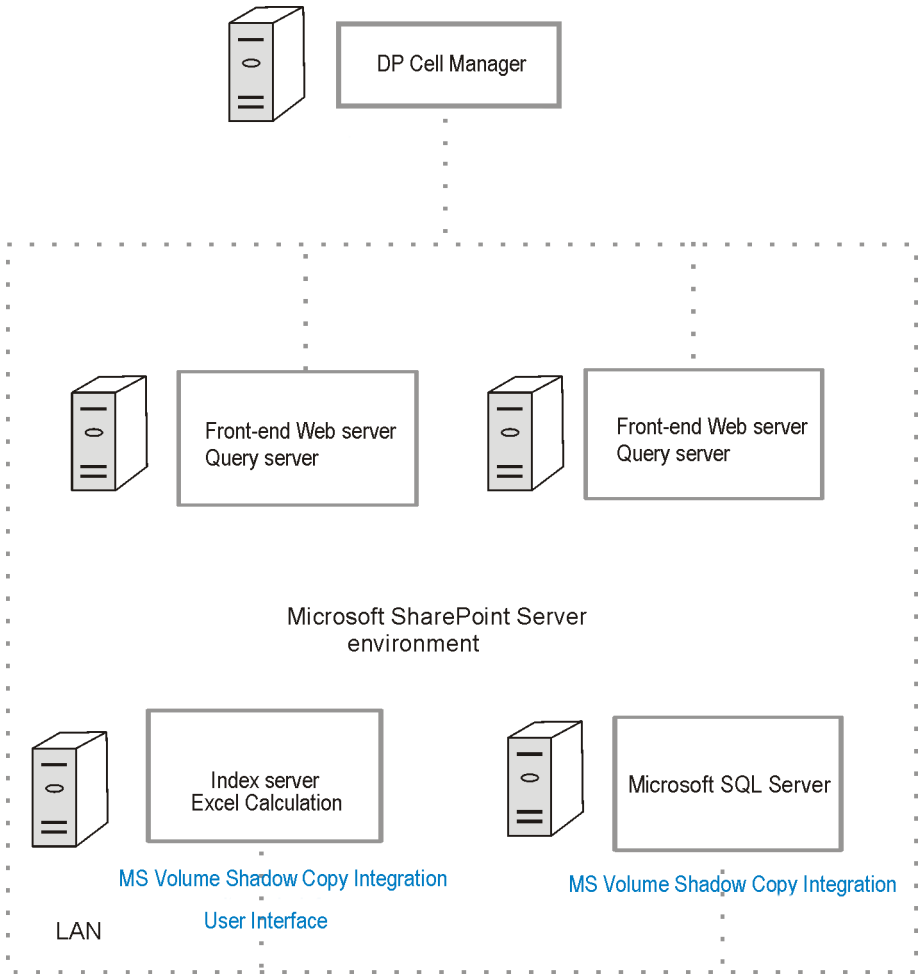

In Figure 37 (page 85), the Data Protector components that you need to install are colored blue.

Configuring the integration

For details on how to configure the Data Protector VSS integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Configuring user accounts

Create or identify a Windows domain user account that has Windows administrative rights on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands. This user must also be granted Microsoft SharePoint Server administrative rights and must be added to the Data Protector `admin` user group.

Backup

To back up Microsoft SharePoint Server data, create backup specifications and start backup sessions using the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`.

Prerequisites

- The Windows Remote Management service (which is used for starting and stopping Windows services remotely, and suspending and resuming FAST for Microsoft SharePoint Server 2010) must be configured on all systems.

  To configure and analyze the WinRM service, execute the `winrm quickconfig` command.

For more information, see the Windows Remote Management service documentation.

- In case of Microsoft SharePoint Server 2010 which uses Microsoft SQL Server 2008 for storing data, and Remote BLOB Storage (RBS) is used with the FILESTREAM provider, ensure that FILESTREAM access level is set to `Full access enabled` or `Transact-SQL access enabled`.

  For details of how to configure RBS and FILESTREAM, see the Microsoft SQL Server 2008 documentation.

Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only `Full` backup type is supported.

Recommendations

- Use the Data Protector PowerShell command to create backup specifications and not the Data Protector GUI.
- Use the Data Protector GUI to modify backup specifications (for example, to add backup devices).
- Use the simple mode for the SQL Server databases. In case you want to use the full mode anyway, ensure that you truncate the transaction logs. Otherwise, you may run out of disk space.
- Whenever you change the farm configuration, perform a new backup.
- In case you want to back up the Single Sign-On database, do not forget to back up the encryption key as described in: http://technet.microsoft.com/en-us/library/cc262932.aspx#Section32.

  Otherwise, you will not be able to restore the database.

How the command works

When you execute the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`, Data Protector first queries for information about the Microsoft SharePoint Server environment. Then it creates backup specifications.

The newly created backup specifications are named `SharePoint_VSS_backup_`*`ClientName`* and have the same backup device specified for use (the one that you specified at command runtime).

Once the backup specifications are created, the command starts backup sessions (one session for each backup specification).

Microsoft Office SharePoint Server 2007

In a Microsoft Office SharePoint Server 2007 environment, the command creates a separate backup specification for each Microsoft Office SharePoint Server 2007 system that has at least one of the following services enabled:

- `Windows SharePoint Services Database`
- `Windows SharePoint Services Help Search`
- `Office SharePoint Server Search`

For a system with the `Windows SharePoint Services Database` service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server

2005/2008) or `MSDE writer` (Microsoft SQL Server 2000) object selected (Figure 38 (page 87)).

Figure 38 Selection of Microsoft Office SharePoint Server 2007 databases

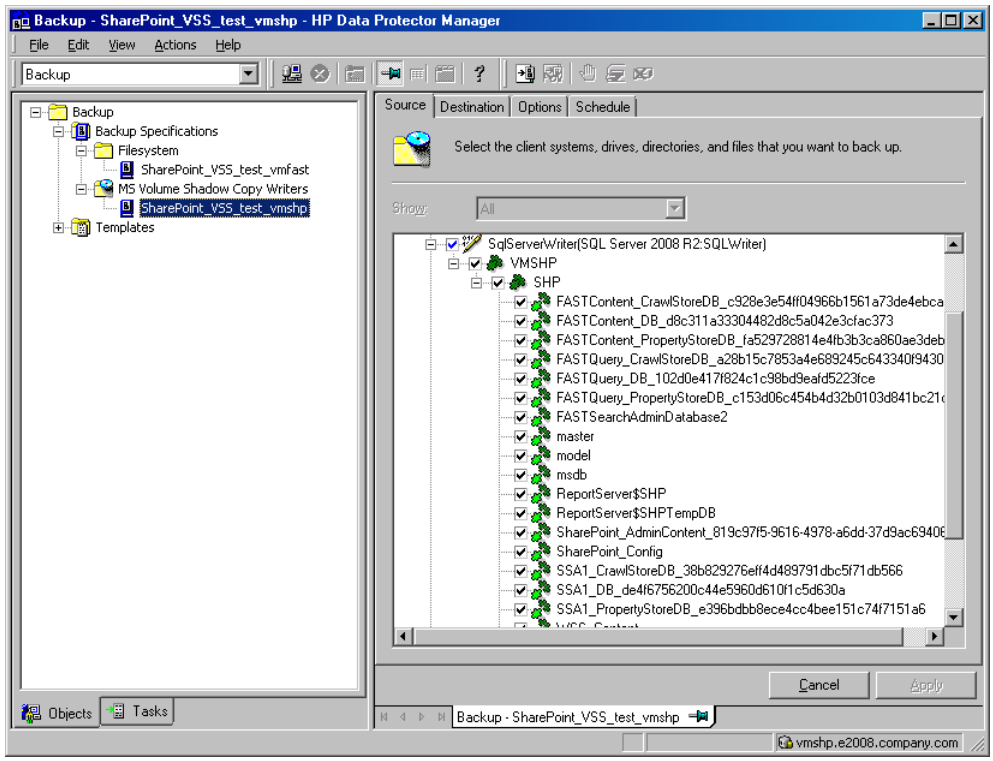

For a system with the `Windows SharePoint Services Help Search` and `Office SharePoint Server Search` services enabled, the command creates a backup specification that has the `SPSearch VSS Writer` and `OSearch VSS Writer` objects selected (Figure 39 (page 88)).

Figure 39 Selecting Microsoft Office SharePoint Server 2007 search index files

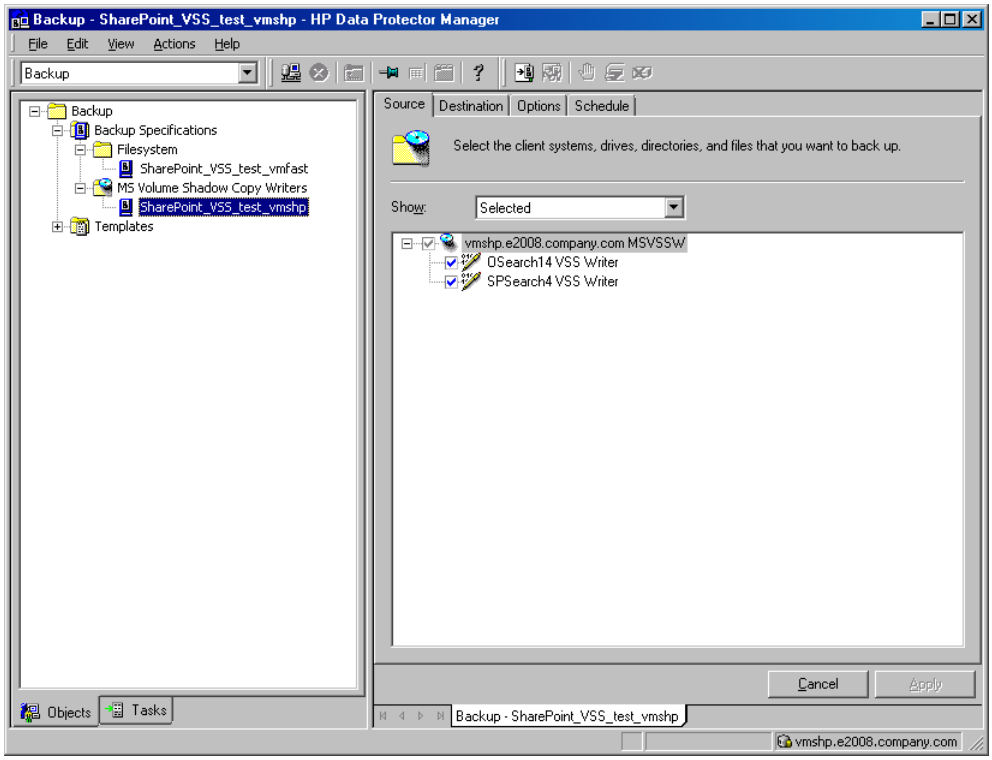

Microsoft SharePoint Server 2010

In a Microsoft SharePoint Server 2010 environment, the command creates a separate backup specification for each Microsoft SharePoint Server 2010 system that has at least one of the following services enabled:

- `SharePoint Foundation Database`
- `SharePoint Foundation Help Search`
- `SharePoint Server Search 14`
- `FAST Search Server 2010 for SharePoint (FAST Search)`

For a system with the `SharePoint Foundation Database` service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server 2005/2008) object selected.

For a system with the `SharePoint Foundation Help Search` and `SharePoint Server Search` services enabled, the command creates a backup specification that has the `SPSearch4 VSS Writer` and `OSearch14 VSS Writer` objects selected.

For a system with the `FAST Search Server 2010` service enabled, the command creates a filesystem backup specification that has the `FASTSearch` home folder selected, excluding `bin` and `lib` folders which contain FAST executables.

Considerations

- ***Microsoft Office SharePoint Server 2007:*** If the `Office SharePoint Server Search` service is enabled on two separate Microsoft SharePoint Server systems so that one is assigned the Query and the other the Indexing role, the command creates a backup specification only for the system with the `Indexing` role. It is not created for the one with the `Query` role. To restore index files on the `Query` system, copy the files from the `Indexing` system to the

Query system after the restore. For details, see the section "Restoring index files on the Query system" (page 102)".

- The command options enable you to split the process into two parts: first you create the backup specifications and then you start backup sessions. In this way, you can manually modify the newly-created backup specifications in the Data Protector GUI before the backup is actually started.
- If Microsoft SQL Server instances are used not only by Microsoft SharePoint Server but also by other database applications, modify the backup specifications so that only the databases that belong to Microsoft SharePoint Server are selected for backup. See the section "Modifying backup specifications" (page 93).
- If you have Microsoft SQL Server database mirroring enabled, a failover can occur and so a different Microsoft SQL Server system becomes active. Since the command creates backup specifications only for the currently active Microsoft SQL Server systems, it is advisable to update (recreate) the backup specifications before the backup is started.

The command syntax

```
SharePoint_VSS_backup.ps1 -help |-version
SharePoint_VSS_backup.ps1 -createonly CreateOptions
SharePoint_VSS_backup.ps1 -backuponly BackupOptions
SharePoint_VSS_backup.ps1 -resumefarm [-preview] | -resumecert

CreateOptions
-device DevName
[-overwrite]
[-prefix PrefixName]
[-excludeindex]

BackupOptions
[-outfile PathToFile]
[-prefix PrefixName]
[-preview]
[-reduce]
[-mode {full | incremental | incremental1 ... | incremental9}]
[-timeout Timeout]
```

IMPORTANT:

- The command must be executed from the *Data_Protector_home*\bin directory on the front-end Web Server system. Ensure that you are logged in under a user account that is configured as described in "Configuring user accounts" (page 85) and that you open the command prompt with administrative rights.
- Do not close the PowerShell console while the backup session is in progress. If you close the console during the backup, some actions are not performed: the backup sessions started do finish, but the farm does not resume the original state. To resume the farm, first execute the command with the –resumefarm option and then unquiesce the farm manually using the Microsoft SharePoint Server Central Administration or stsadm.

Option description

-help
: Displays the SharePoint_VSS_backup.ps1 command usage.

-version
: Displays the SharePoint_VSS_backup.ps1 version.

–createonly
: If this option is specified, Data Protector only creates backup specifications. Backup is not started.

`-backuponly`

If this option is specified, Data Protector only starts backup sessions using the existing backup specifications. The `-device` option is not required.

`-device` *`DevName`*

Specifies which Data Protector device to use for backup. You can specify only one device.

IMPORTANT: If only one device is used to back up a multi-system farm, the corresponding backup sessions cannot run in parallel. This prolongs the time during which the farm is in read-only mode. Specifically, the farm is in read-only mode from the moment when the backup sessions are started up until all VSS snapshots are created.

To enable backup sessions to run in parallel, select different or additional devices in each backup specification before the backup is started. See the section "Modifying backup specifications" (page 93).

`-overwrite`

By default, Data Protector does not create backup specifications if they already exist. If this option is specified, Data Protector overwrites the existing backup specifications with the newly-created ones. Not applicable if `-backuponly` is specified.

`-prefix` *`PrefixName`*

With this option specified, the backup specifications are created under a different name: `SharePoint_VSS_backup_`*`PrefixName_ClientName`*.

In case of backup, this option specifies which backup specifications to use: those which name contains *`PrefixName`*.

Non-ASCII characters in *`PrefixName`* are not supported.

`-outfile` *`PathToFile`*

If this option is specified, backup specification names, errors, sessions outputs, and `omnir` restore commands are written to the specified file.

`-preview`

If this option is specified, Data Protector displays information about the Microsoft SharePoint Server environment and describes the related actions without actually performing them.

`-reduce`

Applicable only to Microsoft SharePoint Server 2010. If this option is specified, the command excludes mirrored query components from backup to reduce the backup size.

`-excludeindex`

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). If this option is specified, Data Protector excludes `data_index` folder contained in the `FASTSearch` home folder from backup specification. This way, the backup is faster, but the restore is more time consuming. The option enables balancing between a backup size and a time to recovery.

`-mode {full|incremental|incremental1... |incremental9}`

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). With this option specified, either a `Full` or `Incremental` or leveled incremental backup can be started. By default, the `Full` backup is performed.

When the `incremental` option is specified and the `Full` backup does not exist, the option is ignored and the `Full` filesystem backup of the FAST Search index files is started.

`-resumecert`

Applicable only to Microsoft FAST Search Server 2010. If this option is specified, the FAST Search certificates for the content and the query connectors are reinstalled.

IMPORTANT: The `SharePoint_VSS_backup.ps1 -resumecert` command must be started on the Microsoft SharePoint Server system where the `SharePoint Server Search 14` service is enabled.

`-resumefarm`

To be used after restore. This option returns the farm to a working state by resuming all background activities and crawling, unlocking sites, and starting Microsoft SharePoint Server services.

IMPORTANT: The command with the `-resumefarm` option specified uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. To ensure its proper operation, an exception must be added to the Windows Default Firewall for `Remote administration`, which adds the WMI ports, or for the WMI directly. For details, see: http://support.microsoft.com/kb/154596.

`-timeout` *`Timeout`*

This option sets the timeout in minutes after which the crawl of the FAST Search index files is aborted and the farm is resumed. If not specified, the default timeout is 15 minutes.

Starting Windows PowerShell

1. Log in to the Microsoft SharePoint Server system where Windows PowerShell and `User Interface` component are installed, under a user account that is configured as described in "Configuring user accounts" (page 85).
2. Open the Windows PowerShell CLI. For example:

   Start > **Programs** > **Accessories** > **Windows PowerShell** > **Windows PowerShell**
3. In case you have Windows User Account Control (UAC) enabled, ensure that you open the CLI with administrative rights. Otherwise, you will not be able to run the Data Protector PowerShell command.

4. Ensure that the Windows PowerShell execution policy is set to `RemoteSigned` or `Unrestricted`.

   "Displaying the Data Protector PowerShell command syntax" (page 92) shows how the Windows PowerShell execution policy is set to `Unrestricted` and how the Data Protector PowerShell command syntax is displayed.

Figure 40 Displaying the Data Protector PowerShell command syntax

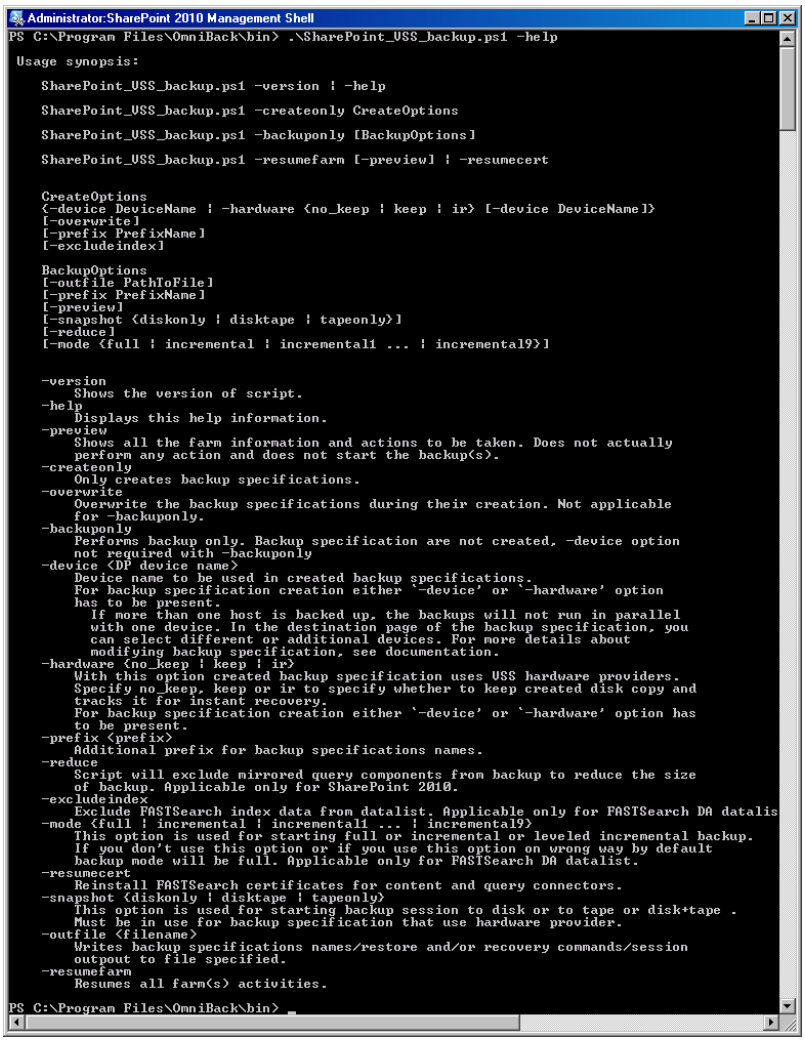

Creating backup specifications (examples)

1. To create backup specifications in which the backup device `filelib_writer1` is specified for use, execute:

   ```
   SharePoint_VSS_backup.ps1 –createonly –device filelib_writer1
   ```

2. To create backup specifications with the label `weekly` in their names and in which the backup device `dev1` is specified for use, execute:

   ```
   SharePoint_VSS_backup.ps1 –createonly –device dev1 –prefix weekly
   ```

3. Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010).

   To create filesystem backup specifications in which the backup device `dev1` is specified for use and with the `data_index` folder, contained in the `FASTSearch` home folder, excluded from the backup of the FAST Search index files, execute:

   ```
   SharePoint_VSS_backup.ps1 –createonly –device dev1 -excludeindex
   ```

Modifying backup specifications

To modify a backup specification, open the Data Protector GUI. In the Context list, select **Backup** and, under **MS Volume Shadow Copy Writers** or under **Filesystem** (if performing a standard filesystem backup of the FAST Search index files), click the name of the backup specification that you want to modify (see Figure 38 (page 87)).

Source page

If you want to modify the Source page of the backup specification (for example, you want to back up individual Microsoft SharePoint Server databases), consider the following:

- The configuration database and the Central Administration content database must both be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- **Microsoft Office SharePoint Server 2007**: The Shared Services Provider database (SSP_DB), Search database (`SSP_Search_DB`), and the associated search index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- **Microsoft SharePoint Server 2010**:
  - The SharePoint Service Applications, Search database (`SSA_Search_DB`), and the associated search index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
  - The FAST search index files and the FAST Content SSA crawl components must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.
- The Help Search database and the associated index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure the data consistency.

Otherwise, after restore, the Microsoft SharePoint Server data may not be consistent.

Destination page

In the Destination page of the backup specification, you can select different or additional devices and set the device and media options.

Options page

In the Options page of the backup specification, you can modify backup options. For the standard filesystem backup of the FAST search index files leave the **Use Shadow Copy** option specified to enable the use of the VSS.

Starting backup sessions (examples)

1. To preview the actions that are performed when a backup session is started, execute:
   `SharePoint_VSS_backup.ps1 –backuponly –prefix dev –preview`
   The following output is from a Microsoft Office SharePoint Server 2007 environment:

```
==============================
 Starting MOSS backup command
 02/10/2011 03:16:30
==============================
------------------------------
List of hosts and their services
------------------------------
 virtual20
 Application Server
   Windows SharePoint Services Help Search
   Windows SharePoint Services Database
   Information Management Policy Configuration Service
   Office SharePoint Server Search
   Shared Services Timer
   Office SharePoint Server Search Admin Web Service
   Excel Calculation Services
   Single Sign-on Service
   SSP Job Control Service
   Portal Service
   Office SharePoint Server Search
   Document Conversions Launcher Service
   Document Conversions Load Balancer Service
   Windows SharePoint Services Web Application
   Central Administration
   Windows SharePoint Services Incoming E-Mail
   Windows SharePoint Services Administration
   Windows SharePoint Services Timer

VIRTUAL21
 Application Server
   Windows SharePoint Services Help Search
   Office SharePoint Server Search
   Shared Services Timer
   Office SharePoint Server Search Admin Web Service
   Single Sign-on Service
   SSP Job Control Service
   Portal Service
   Office SharePoint Server Search
   Windows SharePoint Services Web Application
   Windows SharePoint Services Administration
   Windows SharePoint Services Help Search
   Windows SharePoint Services Timer
```

```
VIRTUAL23
 Application Server
   Windows SharePoint Services Help Search
   Office SharePoint Server Search
   Shared Services Timer
   Office SharePoint Server Search Admin Web Service
   Single Sign-on Service
   SSP Job Control Service
   Portal Service
   Office SharePoint Server Search
   Windows SharePoint Services Web Application
   Windows SharePoint Services Administration
   Windows SharePoint Services Help Search
   Windows SharePoint Services Timer
--------------------------------
SQL hosts list
 virtual20

Index hosts list
 virtual20
 VIRTUAL21
 VIRTUAL23

Help search hosts list
 VIRTUAL21
 VIRTUAL23
-------------------
SUSPENDING FARM
02/10/2011 03:16:43
-------------------
Farm SharePoint_Config

Service Windows SharePoint Services Help Search on host VIRTUAL21
-> Pausing background activity ...
   ... background activity paused.

Service Windows SharePoint Services Help Search on host VIRTUAL23
-> Pausing background activity ...
   ... background activity paused.

Web applications:
 Display name:  Recovery Web Application
 Alternate URL: http://virtual20:999

 Display name:  SharePoint - 123
 Alternate URL: http://virtual20:123
  Web site URL: http://virtual20:123/ssp/admin
  Root title:   Shared Services Administration: SSP1
  -> Setting lock state to readonly
   Crawled by:   , id
   Crawl status:
  -> Pausing background activity
...
Quiesce status is: Quiesced
-------------------
SUSPENDING END
02/10/2011 03:18:28
-------------------
```

```
-> Starting backups...

Starting backup:  omnib -msvssw_list SharePoint_VSS_backup_dev_virtual20 \ -barmode full
Starting backup:  omnib -msvssw_list SharePoint_VSS_backup_dev_VIRTUAL21 \ -barmode full
Starting backup:  omnib -msvssw_list SharePoint_VSS_backup_dev_VIRTUAL23 \ -barmode full

Waiting while VSS creates Volume Shadow Copies ...
       Please wait. DO NOT close PowerShell console!
       After shadow copies are created, the command will resume farm
       and display Data Protector backup session(s) output(s).
  SUCCESS: Volume Shadow Copy successfully created.
           Host   : virtual20
  SUCCESS: Volume Shadow Copy successfully created.
           Host   : VIRTUAL21
  SUCCESS: Volume Shadow Copy successfully created.
           Host   : VIRTUAL23

------------------
RESUMING FARM
02/10/2011 03:18:28
------------------
Service Windows SharePoint Services Help Search on host VIRTUAL21
-> Resuming background activity ...
   ... background activity resumed

Service Windows SharePoint Services Help Search on host VIRTUAL23
-> Resuming background activity ...
   ... background activity resumed

Web site URL: http://virtual20:123/ssp/admin
Root title:  Shared Services Administration: SSP1
 -> Reverting lock for site http://virtual20:123/ssp/admin to none
 -> Resuming background activity
...
--------------------
RESUMING END
02/10/2010 03:19:18
------------------
================================
 MOSS backup command finished
 02/10/2011 03:19:18
 Running time 00:02:48.3336122
================================
```

2. To start backup sessions using the existing backup specifications that have no prefix in their names, execute:

   ```
   SharePoint_VSS_backup.ps1 –backuponly
   ```

3. To start backup sessions using the existing backup specifications that have the prefix `weekly` in their names, execute:

   ```
   SharePoint_VSS_backup.ps1 –backuponly –prefix weekly
   ```

4. To start backup sessions using the existing backup specifications that have no prefix in their names and to save the output of the sessions and the associated restore commands to the file `c:\logs\shp.log`, execute:

   ```
   SharePoint_VSS_backup.ps1 –backuponly –outfile C:\logs\shp.log
   ```

5. To start `incremental` filesystem backup sessions of the FAST Search index files (Microsoft SharePoint Server 2010), execute:

   ```
   SharePoint_VSS_backup.ps1 -backuponly -mode incremental
   ```

Scheduling backup sessions

You can schedule backup sessions using the Windows system scheduler.

1. On the front-end Web server system, create a Windows PowerShell scheduled task. Go to: **Start** > **Settings** > **Control Panel** > **Scheduled Tasks** > **Add Scheduled Task**
2. Open advanced properties for the task.

Figure 41 Scheduling a backup session using the Windows scheduler

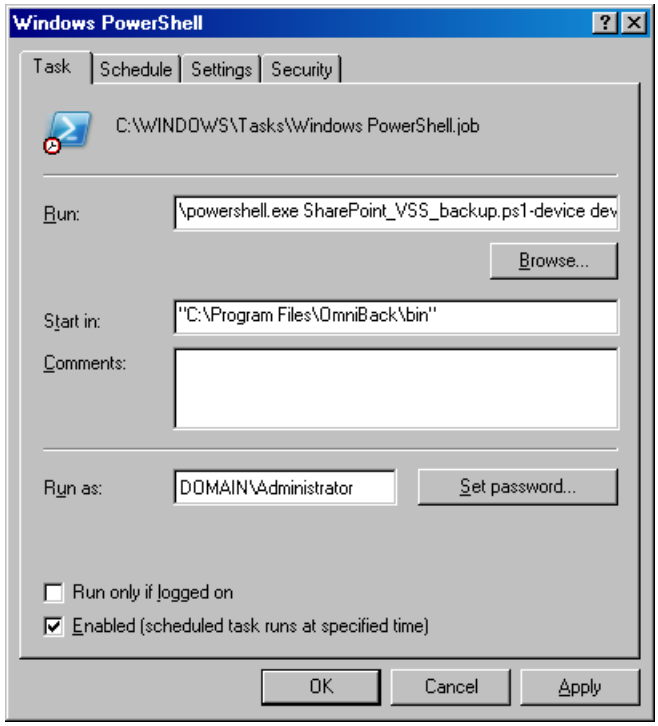

In **Run**, type:

Windows_PowerShell_home\powershell.exe SharePoint_VSS_backup.ps1[Options]

For details on Options, see "The command syntax" (page 89).

In **Start in**, type:

Data_Protector_home\bin

In **Run as**, type a Windows domain user account DOMAIN\UserName that is configured as described in "Configuring user accounts" (page 85).

Restore

To restore Microsoft SharePoint Server data:

- Stop Microsoft SharePoint Server services
- Restore the data.
- Return the farm to a working state.

For details, see the following sections.

Before you begin

- Stop and disable the following services:
  - `IIS Admin Service` (Only for Internet Information Services 6.0 on Windows Server 2003, when the whole farm is restored)
  - `Office SharePoint Server Search` (Microsoft Office SharePoint Server 2007)
  - `SharePoint Server Search 14` (Microsoft SharePoint Server 2010)

In addition, stop the following services:

- Microsoft Office SharePoint Server 2007
  - `Windows SharePoint Services Administration`
  - `Windows SharePoint Services Search`
  - `Windows SharePoint Services Timer`
- Microsoft SharePoint Server 2010
  - `SharePoint 2010 Administration`
  - `SharePoint Foundation Search V4`
  - `SharePoint 2010 Timer`
  - `SharePoint 2010 Tracing`
  - `FAST Search for SharePoint`
  - `FAST Search for SharePoint Monitoring`
- Put the Microsoft SQL Server instance offline if you plan to restore one of the following Microsoft SQL Server databases:
  - `master`
  - `model`
  - `msdb`
  - a database for which Microsoft SQL Server mirroring is enabled

NOTE:

- If you use `SqlServerWriter`, you can restore the `model` and `msdb` databases also when the Microsoft SQL Server instance is online. This is one advantage over `MSDE writer`.
- *Microsoft SQL Server mirroring*: If the original and mirror database reside in separate Microsoft SQL Server instances, put offline both Microsoft SQL Server instances.

Restoring data

You can restore Microsoft SharePoint Server data using the Data Protector GUI or CLI.

Considerations

- The configuration database and the Central Administration content database must both be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency. Since the configuration database and the Central Administration content

database contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.

- **Microsoft Office SharePoint Server 2007**: The Shared Services Provider database (`SSP_DB`), Search database (`SSP_Search_DB`), and the associated search index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
- **Microsoft SharePoint Server 2010**:
  - The SharePoint Service Applications, Search database (`SSA_Search_DB`), and the associated search index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
  - Since the FAST configuration database and the `FAST Search` home folder contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.
  - The FAST Search index files and the FAST Content SSA crawl components must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
- The Help Search database and the associated index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency.
- The following table shows which VSS restore modes are supported for which writers:

Table 22 VSS supported restore modes and writers

Writers	VSS restore modes	
	Restore to another client	**Restore files to temporary location**
MSDE writer SqlServerWriter	No	Yes (manual attach needed)
OSearch VSS writer OSearch14 VSS writer	Yes	No
SPSearch VSS writer SPSearch4 VSS writer	Yes	No

Prerequisites

- Applicable only to a Data Protector filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010). Before restoring the FAST Search index files, the **Overwrite** option must remain selected to ensure the data consistency. It is selected by default.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.

2. In the Scoping Pane, expand **MS Volume Shadow Copy Writers**, expand the client which data you want to restore, and then click **MS Volume Shadow Copy Writers**.

   If performing a filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010), expand **Filesystem**, expand the client which data you want to restore, and then click the filesystem object.

3. In the Source page, select the data that you want to restore.

Figure 42 Selecting Microsoft Office SharePoint Server 2007 databases for restore

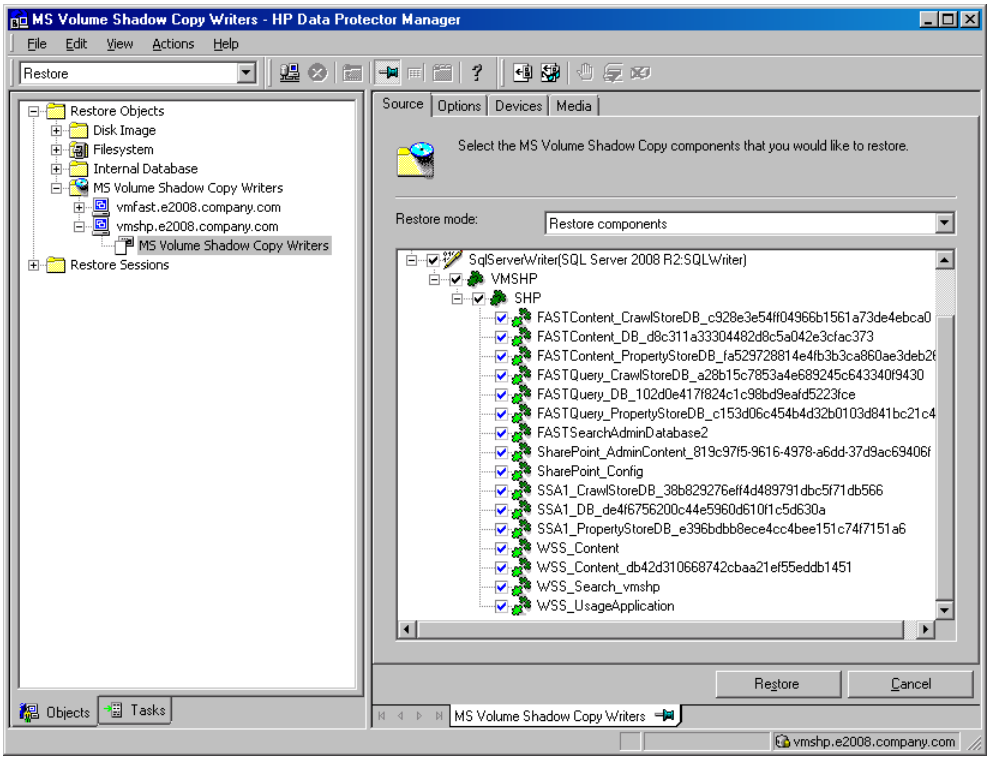

Figure 43 Selecting Microsoft Office SharePoint Server 2007 Search index files for restore

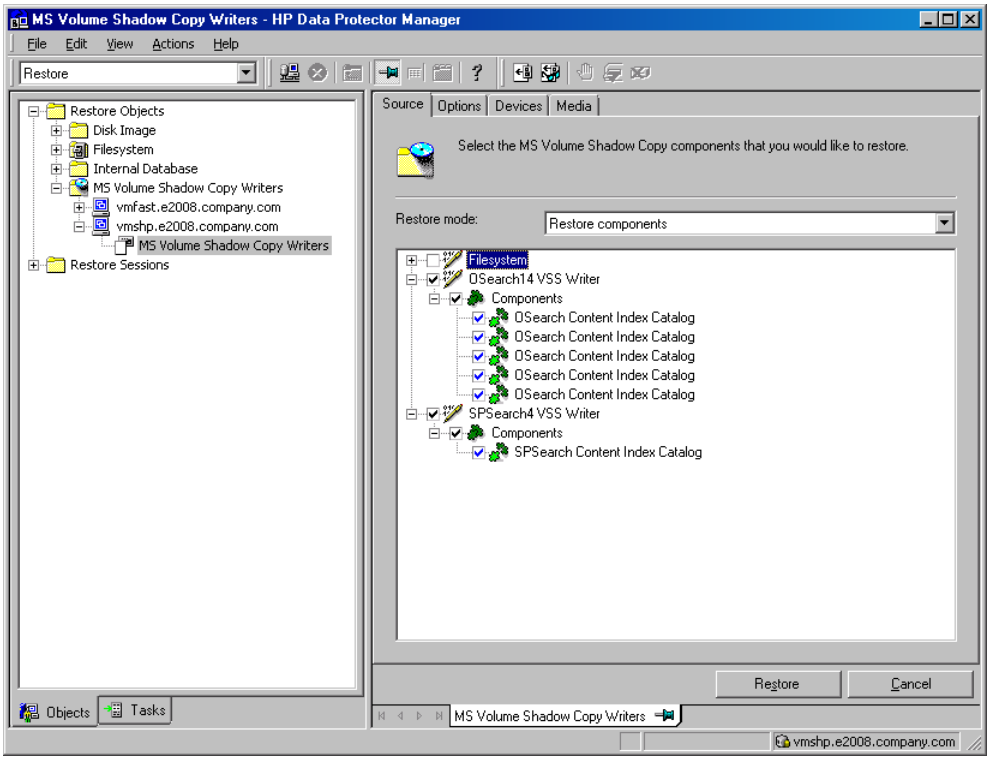

4. In the Options page, specify the restore options.
5. In the Devices page, select devices to use for restore.
6. Click **Restore**, review your selection, and click **Finish**.

Restoring using the Data Protector CLI

You can restore Microsoft SharePoint Server data using the Data Protector `omnir` command. For details, see the `omnir` man page or the *HP Data Protector Command Line Interface Reference*.

If you specified the `-outfile` option when you ran backup sessions, you can find the necessary `omnir` commands in the specified file. The following is an example of the `omnir` command from such a file.

```
 omnir -vss -barhost SHP-APP
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/master"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server2005:SQLWriter)/SHP-APP/model"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter/SHP-APP/msdb"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/
WSS_Content_SSPAdminAccounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/SSP_Accounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/
SHP-APP/SSP_Accounting_Search"
```

Limitations

The `omnir` command syntax should not contain more than 8191 characters. If you have so many `-tree` objects that the syntax exceeds 8191 characters, split the objects and run two separate sessions.

After the restore

After the restore:

1. Enable and start the service `IIS Admin Service` (only for IIS 6 on Windows Server 2003, when the whole farm was restored)
2. Enable the service `Office SharePoint Server Search or SharePoint Server Search 14`.
3. Bring the Microsoft SQL Server instances online (if offline).
4. Return the farm to a working state (that is, resume background activities and crawling, unlock sites, and start the Microsoft SharePoint Server services) by executing:

   `SharePoint_VSS_backup.ps1 -resumefarm`

NOTE:

- The command uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. Ensure its proper operation by adding an exception to the Windows Default Firewall for `Remote administration`, which adds the WMI ports, or for the WMI directly. For details, see: http://support.microsoft.com/kb/154596.
- If the FAST Search certificates for the content and query connectors are out of sync, you can reinstall them by executing:

  `SharePoint_VSS_backup.ps1 -resumecert`

  Start the command on the Microsoft SharePoint Server system where the `SharePoint Server Search 14` service is enabled.

Restoring index files on the Query system

This section is applicable for Microsoft Office SharePoint Server 2007 only. The `Office SharePoint Server Search` service is enabled on two separate Microsoft Office SharePoint Server 2007 systems, so that one is assigned the `Indexing` and the other the `Query` role.

To copy the newly restored index files from the `Indexing` system to the `Query` system, perform the following steps (depending on which Microsoft Office SharePoint Server 2007 and Windows Shared Services service pack you have):

- **Service Pack 1**
  1. On the `Query` system, stop and disable the service `Office SharePoint Server Search`.
  2. Copy the index files from the `Indexing` to the `Query` system.

     By default, index files are located in the `C:\Program Files\Microsoft Office Servers\12.0\Data\Office Server\Applications` directory.
  3. On the `Query` system, enable and start the service `Office SharePoint Server Search`.
- **Service Pack 2**

  On the `Query` system, execute:

  `stsadm -o search -reprovisionindex -ssp` *`SSPName`*

  for each Shared Services Provider separately.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft SharePoint Server VSS based solution.

For Microsoft Volume Shadow Copy troubleshooting information, see the troubleshooting chapter in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

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

Before you begin

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

Checks and verifications

If your browsing, backup, or restore failed:

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

After restore, you cannot connect to the Central Administration webpage

Problem

After restore, when you try to connect to the Microsoft SharePoint Central Administration webpage, an error similar to the following is displayed in your web browser:

Windows Internet Explorer:

```
Retrieving the COM class factory for component with CLSID
(BDEADEE2-C265-11D0-BCED-00A0C90AB50F) failed due to the following error
800703fa.
```

Mozilla Firefox:

```
An unexpected error has occurred.
```

Action

1. Restart Microsoft SharePoint Server services on all clients in the farm.
2. Open the Internet Information Services (IIS) Manager and restart all application pools.
3. In case an application pool fails to be restarted with the following error:

   ```
   Cannot Restore Application Pool. There was an error while performing
   this operation.
   ```

   wait for a few seconds and then restart the operation.
4. Delete browsing history in your web browser.
5. Log in to the Central Administration webpage.

Backup fails with the error Failed to resume Service Windows SharePoint Services Help Search

Problem

When you start backup sessions, an error similar to the following is displayed:

```
Service Windows SharePoint Services Help Search on host
MOSS07-INDEX
-> Resuming background activity ...
ERROR: Failed to resume Service Windows SharePoint Services Help Search
on host MOSS07-INDEX
Web site URL: http://moss07-web:2001
Root title: as
-> Resuming background activity
```

Action

Execute:

```
SharePoint_VSS_backup.ps1-resumefarm
```

After restore, a quiesce operation fails

Problem

After you have restored the configuration database and executed `SharePoint_VSS_backup.ps1-resumefarm`, the data in the Microsoft SharePoint Server file system caches on front-end Web Server systems is not consistent with the data in the newly-restored configuration database. When you try to quiesce the farm, the operation fails with the following error:

```
An unhandled exception occurred in the user interface. Exception
Information: An update conflict has occurred, and you must re-try
this action. The object SessionStateService Parent=SPFarm Name=<
farm_config_database_name > is being updated by < domain\username
>, in the w3wp process, on machine < servername >. View the tracing
 log for more information about the conflict.
```

Action

Clear the Microsoft Office SharePoint Server file system cache on all server systems in the farm. For details, see:

http://support.microsoft.com/kb/939308.

After restore, you cannot connect to the FAST Search Server

Problem

After restore, when you try to connect to the Microsoft FAST Search Server 2010 system for SharePoint, the operation fails.

FAST Query SSA search operations display an error similar to the following:

```
The search request was unable to connect to the Search Service.
```

Action

Execute:

```
SharePoint_VSS_backup.ps1 -resumecert
```

NOTE: The VSS based solution copies the FAST Search certificate `FASTSearchCert.pfx` from the FAST Admin Server system to the SharePoint Server system and installs it. Also, the SharePoint certificate is copied and installed to each FAST Search Server system. For details, see: http://technet.microsoft.com/en-us/library/ff381244.aspx.

The SharePoint_VSS_backup.ps1 script stops responding and the farm stays in read only mode

Problem

When starting a back up, the `SharePoint_VSS_backup.ps1` script stops responding when a crawl of the Microsoft SharePoint Server is being performed. The issue can appear due to external conditions such as a corrupted SSA index, the need to reissue the certificate manually and so on.

As a result, the farm stays in read-only mode.

Action

Normally, the crawl should be aborted automatically after 15 minutes. If this does not happen:

1. Abort the script by pressing **Ctrl-C**.
2. Manually resume the farm.

You can specify a different timeout after which the crawl is aborted and the farm is resumed by using the `-timeout` option.

Part III Microsoft Exchange Server

Data Protector offers different ways to back up Microsoft Exchange Server data online. Choose the appropriate backup and restore solution depending on your Microsoft Exchange Server version and the desired functionality.

Microsoft Exchange Server 2003

- ***Data Protector Microsoft Volume Shadow Copy Service integration***

  Use this integration to back up Microsoft Exchange Server 2003 data using VSS writers. You can back up all Microsoft Exchange Server data or individual storage groups.

  This integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Microsoft Exchange Server 2007

- ***Data Protector Microsoft Exchange Server 2007 integration***

  This integration operates on the level of Exchange Server databases. You can back up all Microsoft Exchange Server data or only particular databases: Information Store, Key Management Service, and/or Site Replication Service. Specific features such as local continuous replication (LCR) or cluster continuous replication (CCR) are not supported.

  See "Data Protector Microsoft Exchange Server 2007 integration" (page 108).

- ***Data Protector Microsoft Volume Shadow Copy Service integration***

  Use this integration to back up Microsoft Exchange Server 2007 data using VSS writers. You can back up all Microsoft Exchange Server data or individual storage groups. Use this integration to back up Exchange Server 2007 specific items such as local continuous replication (LCR) copies or cluster continuous replication (CCR) copies.

  Use this integration to back up *Microsoft Exchange Server 2007* data that resides on disk arrays. This integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

- ***Data Protector Microsoft Exchange Single Mailbox integration***

  This integration operates on the level of Microsoft Exchange Server logical objects. The smallest object that you can back up or restore is a Microsoft Exchange Server Mailbox or Public Folders item. In one session, you can back up data from only one Microsoft Exchange Server system.

  See "Data Protector Microsoft Exchange Single Mailbox integration" (page 158).

Microsoft Exchange Server 2010 and Microsoft Exchange Server 2013

- ***Data Protector Microsoft Exchange Server 2010 integration***

  This integration operates on the level of Microsoft Exchange Server logical objects. The smallest object that you can back up or restore is a Microsoft Exchange Server database. The integration also supports DAG environments, enabling you to back up database copies from different Microsoft Exchange Server systems in the same session. See "Data Protector Microsoft Exchange Server 2010 integration" (page 123).

- ***Data Protector Microsoft Exchange Server 2010 ZDB integration***

  Use this integration if your Microsoft Exchange Server data resides on disk arrays. In addition to the benefits already provided with the *Data Protector Microsoft Exchange Server 2010 integration*, this integration enables you to perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Zero Downtime Backup Integration Guide*.

- ***Data Protector Microsoft Volume Shadow Copy Service integration***

  This integration operates on the level of Microsoft Volume Shadow Copy Service writers. Since the Microsoft Exchange Server writer is only one of them, Microsoft Exchange Server data does not have the highest visibility. The smallest object that you can back up or restore is a Microsoft Exchange Server database file's (`.edb`) object or log files' object. In one session, you can back up data from only one Microsoft Exchange Server system. If your Microsoft Exchange Server data resides on disk arrays, you can also perform zero downtime backup (ZDB) and instant recovery (IR) sessions.

  See the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

- ***Data Protector Microsoft Exchange Single Mailbox integration***

  This integration operates on the level of Microsoft Exchange Server logical objects. The smallest object that you can back up or restore is a Microsoft Exchange Server Mailbox or Public Folders item. In one session, you can back up data from only one Microsoft Exchange Server system.

  See "Data Protector Microsoft Exchange Single Mailbox integration" (page 158).

- ***Data Protector Granular Recovery Extension for Microsoft Exchange Server***

  Use this extension to recover individual mailbox items, such as e-mail folders, calendar, contacts, or notes. The Data Protector Granular Recovery Extension for Microsoft Exchange Server is a specialized extension that integrates with Microsoft Exchange Server 2010 or Microsoft Exchange Server 2013 in different Exchange Server environments and provides you complete control over what is recovered. The extension does not provide any backup solution but depends on the Data Protector Microsoft Exchange Server 2010 integration for the backup and restore.

  See the *HP Data Protector Granular Recovery Extension User Guide for Microsoft Exchange Server*.

NOTE: You can also back up Microsoft Exchange Server data using common Data Protector filesystem backup functionality, which operates on the file level. The smallest object that you can back up or restore this way is a database file. To ensure data consistency, you must take the Microsoft Exchange Server offline before starting a backup session.

Table 23 Data Protector backup solutions for Microsoft Exchange Server 2010 and Microsoft Exchange Server 2013

Feature		Disk Agent (filesystem backup)	Volume Shadow Copy Service integration	Exchange Server 2010 integration	Exchange Single Mailbox integration[1]	Granular Recovery Extension
Crash-consistent backup		When selecting parts of the drive and VSS	No	No	No	N/A (only a restore solution, no backup)
Application-consistent backup		When selecting the entire drive and VSS	Yes	Yes	Yes	
Granularity	Select databases by type (active / passive)	No	No	Yes	No	Individual mailbox item (e-mail folder, calendar, contacts, and so on)
	Mailbox full and incremental backup	No	No	No	Yes	
	Database view	No	No	Yes	No	
	System view	Yes	Yes	Yes	Yes	
	DAG view	No	No	Yes	No	
	Single mailbox/mail recovery	With Microsoft tools	With Microsoft tools	With Microsoft tools	Yes	

Table 23 Data Protector backup solutions for Microsoft Exchange Server 2010 and Microsoft Exchange Server 2013 *(continued)*

Feature		Disk Agent (filesystem backup)	Volume Shadow Copy Service integration	Exchange Server 2010 integration	Exchange Single Mailbox integration[1]	Granular Recovery Extension
Backup types	Full (including log truncation)	No	Yes	Yes	No	N/A (only a restore solution, no backup)
	Copy	Yes	Yes	Yes	No	
	Incremental (including log truncation)	No	Only if active/ passive status did not change	Yes	No	
	Differential	No	Yes	Yes	No	
Are zero downtime backup and instant recovery supported?		Offline (filesystem) ZDB and IR	Yes	Yes	No	Depends on the solution used for backup
Database recovery to a point in time		Yes	Yes	Yes	No	
Database recovery to the latest state		Manually	Yes	Yes	No	
Extra licenses needed		None	• On-line Extension • Zero Downtime Backup (optional) • Instant Recovery (optional)	• On-line Extension • Zero Downtime Backup (optional) • Instant Recovery (optional)	• On-line Extension	• Granular recovery extension

[1] This is an MAPI-based integration.

4 Data Protector Microsoft Exchange Server 2007 integration

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server 2007 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 24 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 "Data Protector Microsoft Exchange Single Mailbox integration" (page 158).

This chapter provides information specific to this integration. For general Data Protector procedures and options, see the *HP Data Protector 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 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 44 (page 109).

Figure 44 Data Protector Exchange Server integration architecture

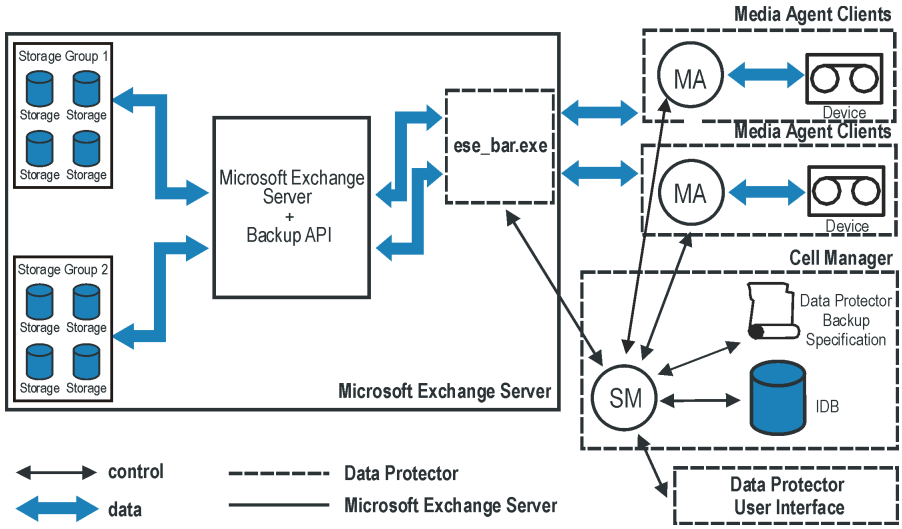

Table 25 Legend

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 to have the Data Protector online extension license-to-use (LTU) for Windows to be able to use the Data Protector Microsoft Exchange Server integration.

  For more information, see the *HP Data Protector Installation and Licensing Guide*.
- Ensure that you correctly installed and configured Exchange Server.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
  - For information on installing, configuring, and using 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 versions, backup objects belonging to a particular Exchange Server version cannot be restored to Data Protector clients on which a different Exchange Server version is installed.

Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the *HP Data Protector Help* index: "configuring devices" and "creating media pools".
- To test whether Exchange Server and Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore. For instructions, see the *HP Data Protector Help*.
- 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**.

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

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 45 Selecting a blank template

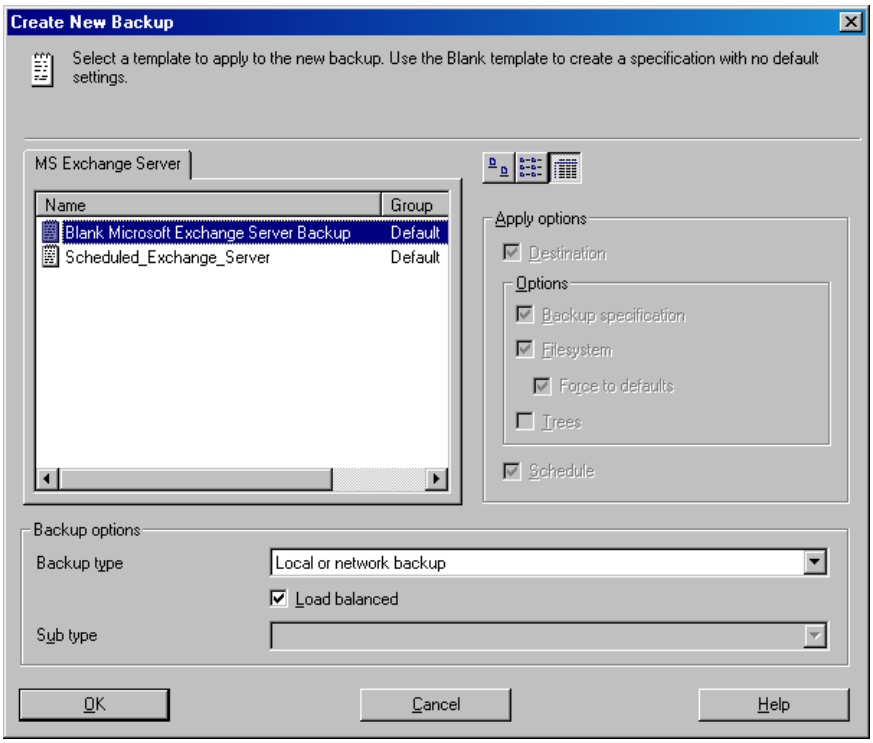

4. In **Client**, select Exchange Server. For cluster environments, select the virtual server of the Exchange Server resource group.

   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)

   For information on the **User and group/domain** options, press **F1**.

   Click **Next**.

Figure 46 Client name and application database

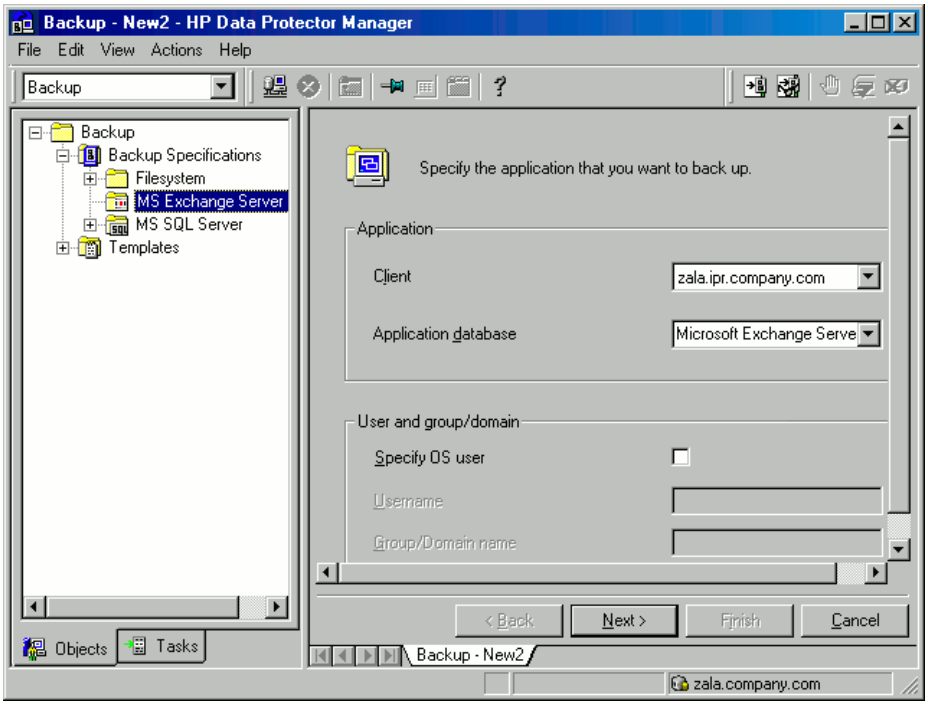

5. Select Exchange Server databases you want to back up.

Figure 47 Backup objects

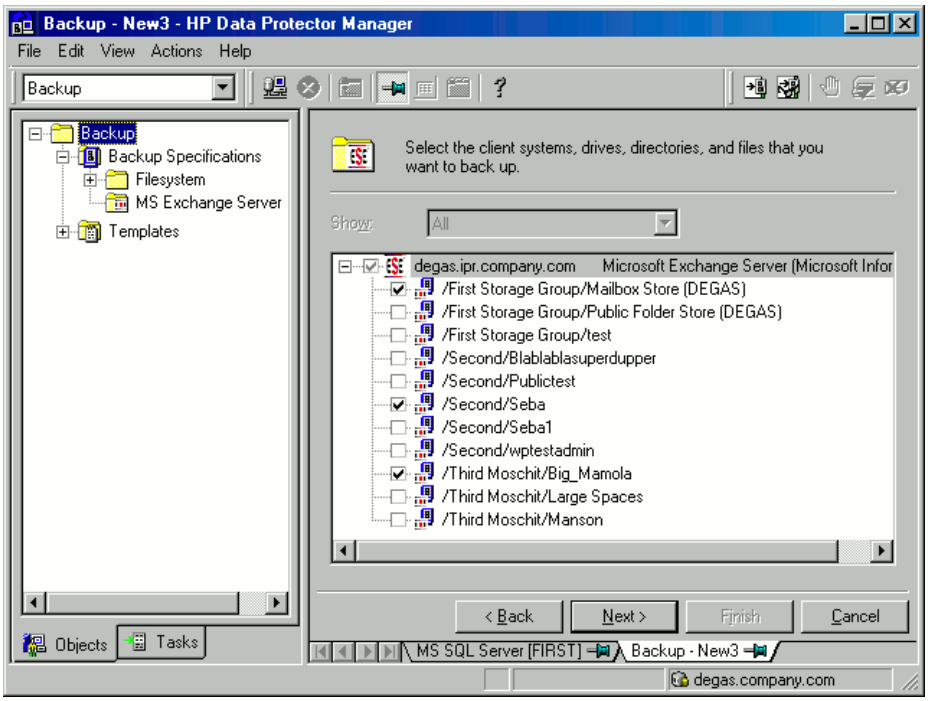

Click **Next**.

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 *HP Data Protector 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 48 Backup devices

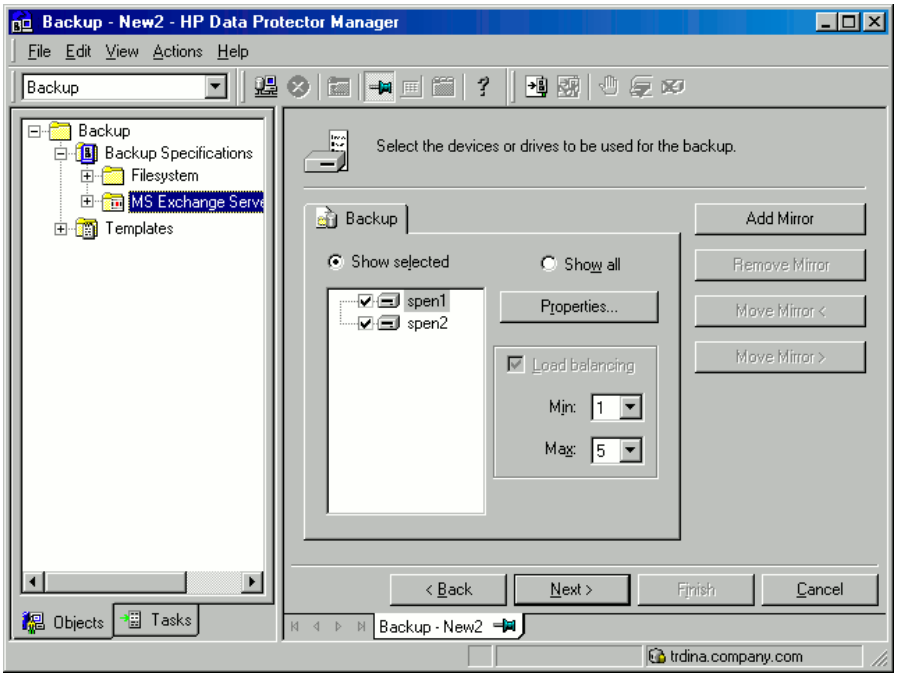

Click **Next** to proceed.

7. Select the backup options.

   For information on **Backup Specification Options** and **Common Application Options**, see the *HP Data Protector Help*.

   For information on **Application Specific Option**, see "SQL Server-specific backup options" (page 31) or the *HP Data Protector 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 49 Application-specific options

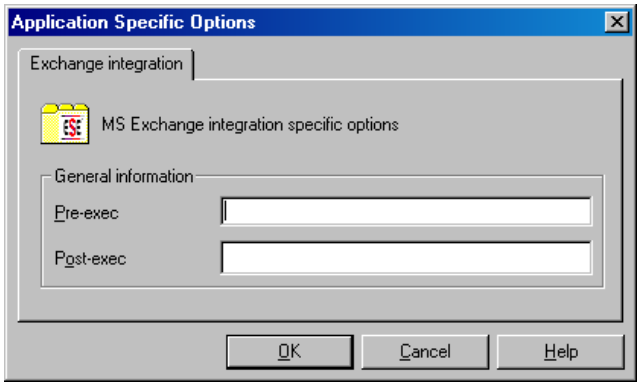

Table 26 Application-specific options

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.

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 *HP Data Protector Help* index: "scheduled backups".

Scheduling example

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

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.

NOTE: 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 use and select **Start Backup**.
3. Select **Backup type** and **Network load**. For information on these options, click **Help**.
4. Click **OK**.

Restore

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

Considerations

- If the Recovery Storage Group (RSG) exists on the Exchange Server system when a restore session is started, the restore of databases is redirected to the RSG. To prevent the restore of databases in such circumstances, set the omnirc option `OB2MSESE_CHECK_RSG` to `1`.

IMPORTANT: The database (store) must be dismounted before a restore.

To dismount 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 50 Dismounting a database

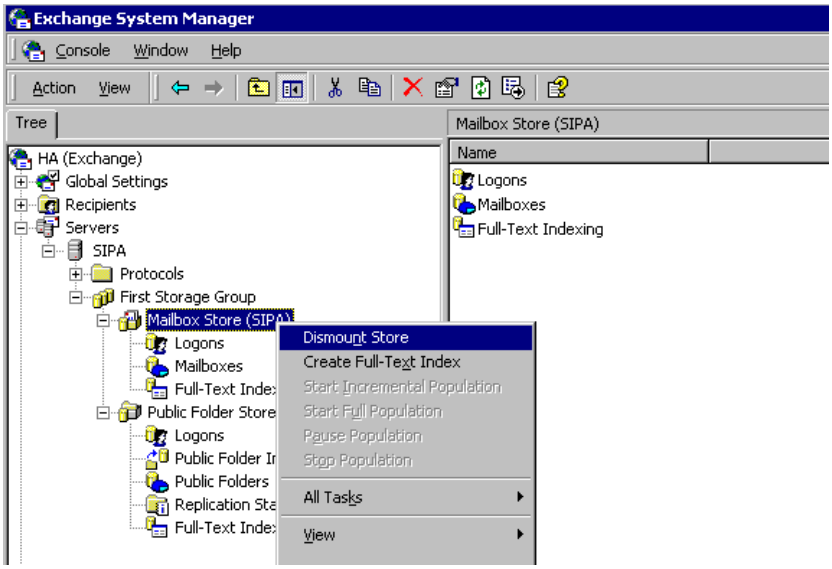

2. A warning appears. Click **Yes** to continue dismounting.

When dismounting completes, you may start restore.

After hard recovery, databases can be mounted automatically. See Table 27 (page 117) 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 51 Restore objects

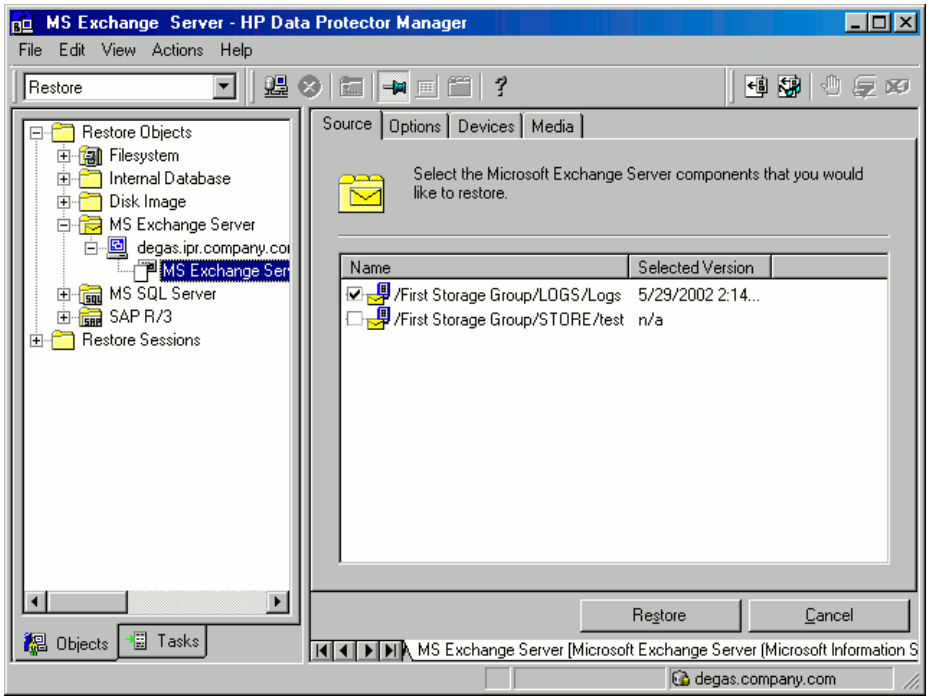

To select a backup version, right-click the object and select **Properties**.

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 52 Selecting a backup version

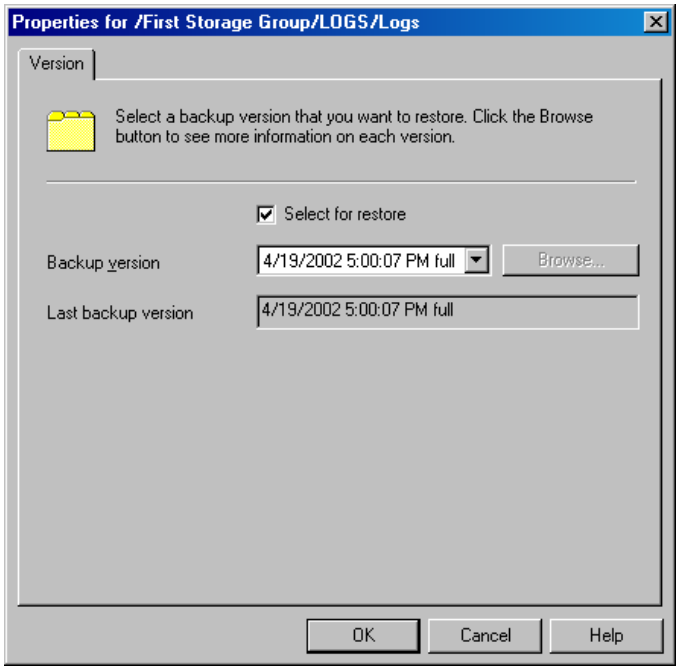

NOTE: Restoring databases to a certain state often requires a multiphase restore (multiple versions must be restored to retrieve data). During incremental backups, only transaction logs of storage groups are backed up (without information on physical location of the storage groups); therefore, you must 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`*** `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 27 (page 117) for details.
5. In the **Devices** page, select the devices to be used for the restore.

   For more information on how to select devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".
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 27 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 `eseutil /cc /t` 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 53 Restore options

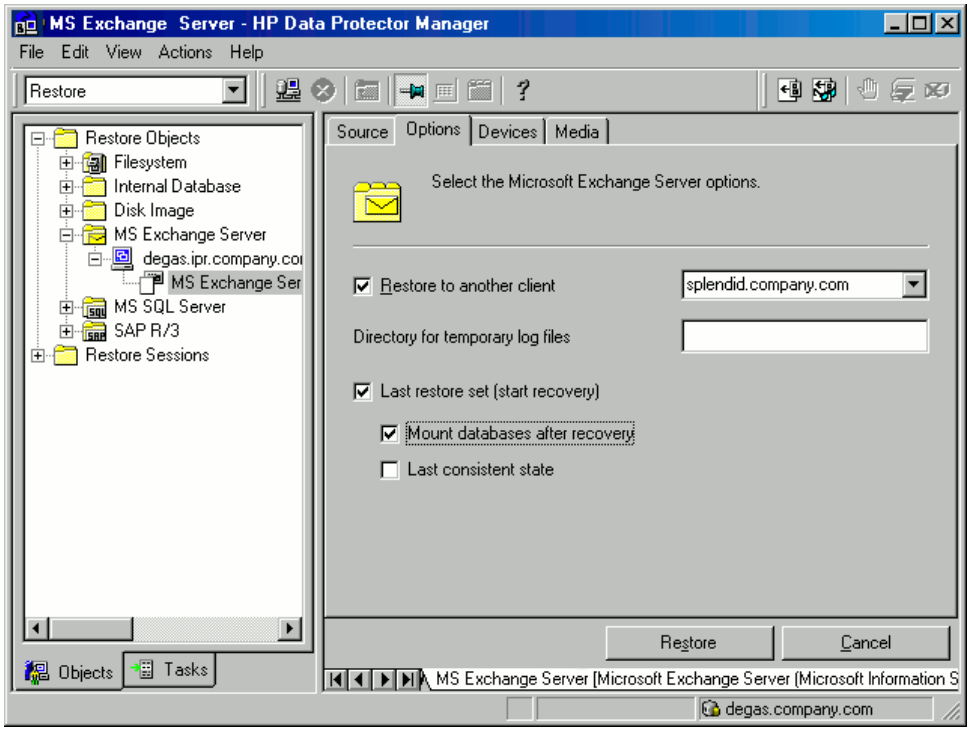

Restoring to another client

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 27 (page 117) 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

To perform a restore using CLI, execute the following command:

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

NOTE: A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the object *copy* session.

The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

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 backup ID `2010/07/07-13`, 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 "2010/07/07-13" -base "/First Storage
Group/STORE/Public Folder Store" -session "2010/07/07-13" -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" (page 120). 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. On how to verify this, see the *HP Data Protector Help* index: "patches".
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://support.openview.hp.com/selfsolve/manuals for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- 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 *HP Data Protector 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" (page 109).
- 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 session fails

During the restore session, the following error is displayed:

```
[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

Open the command prompt window and execute the `eseutil` command 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. Execute the following commands:

     ```
     Drive_letter:
     cd "\Temporary_log_files_directory_path\Storage Group Number"
     eseutil /cc /t
     ```

Problem

Restore of a database fails

Restore of an Exchange Server 2007 database ends abnormally after reporting the following error:

```
[Critical] From: OB2BAR_main@Hostname "Microsoft Exchange Server
(Microsoft Information Store)"  Time: Date Time
[151:214] Recovery SG 'RSG_name' is configured on the Microsoft
Exchange Server.
```

The problem occurs in two cases:

- If you try to restore the database to its original location when the Recovery Storage Group (RSG) exists on the Exchange Server system.

  The RSG may have been created manually or using the VSS integration agent for restoring a store to the RSG. Under such circumstances, Exchange Server redirects the restore of the database to the RSG instead of restoring the database to the original storage group.

- If you try to restore the database to the RSG when the omnirc option `OB2MSESE_CHECK_RSG` is set to `1`.

Action

To enable the restore of the database to the original storage group, perform one of the following:

- Using Exchange Management Console or Windows PowerShell, remove the RSG from the Exchange Server system.
- Add a registry key which will override redirection of restore to the RSG:
  1. Start Windows Registry Editor.
  2. In Registry Editor, expand the folder:

     ```
     HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
     MSExchangeIS\ParametersSystem
     ```

  3. Create a new `DWORD` value `Recovery SG Override` and set its value to `1`.

To enable the restore of the database to the RSG, set the omnirc option `OB2MSESE_CHECK_RSG` to `0`.

Problem

Restore of a database to the Recovery Storage Group (RSG) fails

Restore of an Exchange Server 2007 database to the RSG ends abnormally after reporting the following error:

```
ESE subsystem or operating system reported error for Mailbox:
0xc7fe1f42: Database not found.
```

The problem occurs when the database being restored is not properly linked to the RSG.

Action

To enable the restore of the database to the RSG, using Exchange Management Console or Windows PowerShell appropriately link the database to the RSG.

Problem

Restore of a database to the Recovery Storage Group (RSG) fails

Restore of an Exchange Server 2007 database to the RSG ends abnormally after reporting the following error:

```
ESE subsystem or operating system reported error for ():
0x3f3: The configuration registry key could not be opened.
```

The problem occurs when the database, which has been successfully restored to the RSG, cannot be mounted.

Action

To enable the database that has been restored to the RSG to be mounted, perform one of the following:

- In the Exchange Management Console, go to the **Database Recovery Management** tool and perform the task **Set up the 'Database can be overwritten by restore' flag**.
- In the Windows PowerShell, execute:

  ```
  Set-MailboxDatabase 'ExchangeServerName\RSGName\StoreName'
  -AllowFileRestore $true
  ```

5 Data Protector Microsoft Exchange Server 2010 integration

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server 2010 integration, where Data Protector integrates with Microsoft Exchange Server 2010 or Microsoft Exchange Server 2013 (hereinafter, both Exchange Servers are called **Microsoft Exchange Server**, unless differences are pointed out). It describes concepts and methods you need to understand to back up and restore Microsoft Exchange Server 2010 mailbox and public folder databases or Microsoft Exchange Server 2013 mailbox databases (**databases**).

Both standalone environments and Database Availability Group (**DAG**) environments are supported.

The Data Protector Microsoft Exchange Server 2010 integration is based on the Volume Shadow Copy Service (**VSS**) technology. For details on VSS concepts, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Backup

During backup, databases can be used actively (**online backup**). In DAG environments, you can back up active and/or passive database copies.

You can select among the following Microsoft Exchange Server backup types:

- Full
- Copy
- Incremental
- Differential

For details on the backup types, see "Backup types" (page 128).

Restore

During restore, each database can be restored using a different restore method. The following methods are available:

- Repair all passive copies with failed status
- Restore to the latest state
- Restore to a point in time
- Restore to a new mailbox database
- Restore files to a temporary location

This chapter provides information specific to the Microsoft Exchange Server 2010 integration. For additional limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HP Data Protector Help*.

Integration concepts

Data Protector integrates with Microsoft Exchange Server through the Data Protector Microsoft Exchange Server integration agent, which channels communication between the Data Protector Session Manager and the clients in the Microsoft Exchange Server environment. The agent communicates with Microsoft Exchange Server through the Microsoft Exchange Management Shell and uses VSS to back up data.

Supported environments

Data Protector supports Microsoft Exchange Server Database Availability Group environments (**DAG environments**) as well as environments with standalone Microsoft Exchange Server systems (**standalone environments**).

Standalone environments

In a standalone Microsoft Exchange Server environment, each Microsoft Exchange Server system stands on its own.

In one session, you can back up databases from only one Microsoft Exchange Server system. Data Protector sends backup and restore requests directly to the Microsoft Exchange Server system.

Figure 54 Standalone environment (example)

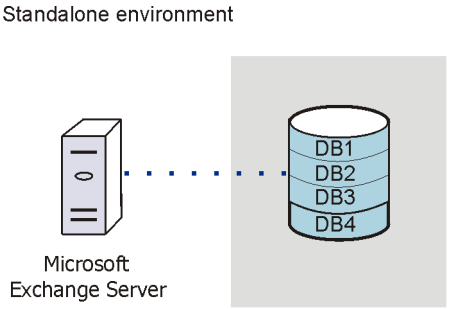

DAG environments

In a DAG environment, Data Protector communicates with the DAG using one of the Microsoft Exchange Server systems (the one that is currently active in the environment). All backup and restore requests are sent there.

In one session, you can back up active and/or passive database copies from different Microsoft Exchange Server systems that belong to the same DAG.

Figure 55 DAG environment (example)

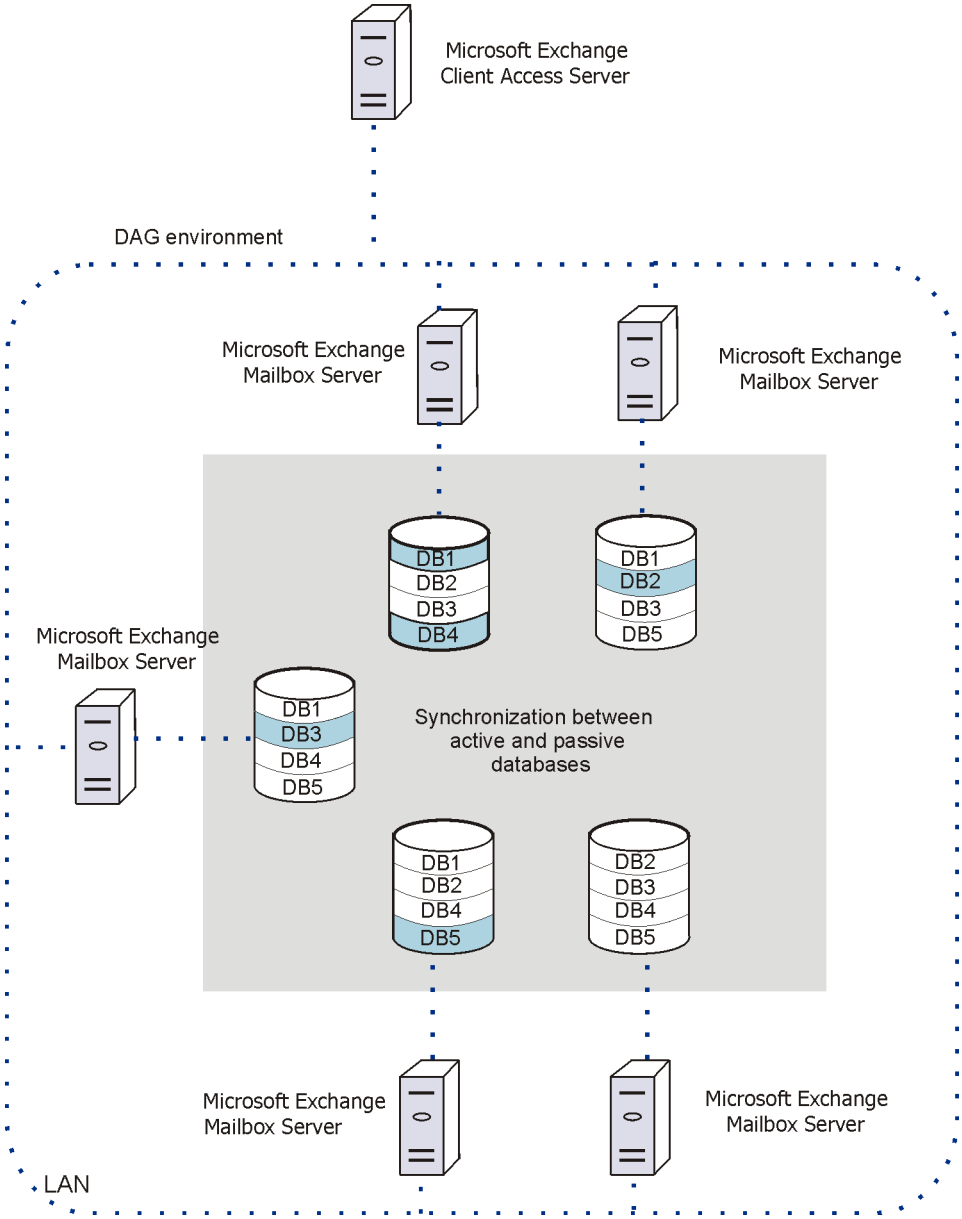

In Figure 55 (page 125), active databases are shaded in blue.

If a database has multiple passive copies, you can specify which particular passive copy you want to back up, using one of the following backup policies:

- minimize the number of hosts
- lowest activation preference
- highest activation preference
- shortest replay lag time
- longest replay lag time
- longest truncation lag time

You can also specify from which Microsoft Exchange Server systems database copies should not be backed up.

For a brief description of the activation preference number, replay lag time, and the truncation lag time, see " Microsoft Exchange Server parameters in DAG environments" (page 126).

Table 28 Microsoft Exchange Server parameters in DAG environments

Parameter	Description
Activation preference number	The activation preference number determines which passive copy is activated if multiple passive copies meet the same criteria; the copy assigned the lowest activation preference number is activated.
Replay lag time	The `ReplayLagTime` parameter plays a role when synchronizing a passive copy with the active copy. As soon as a log file at the active copy side is filled up, it is copied to the passive copy side. By default, the newly copied log is also applied to the passive copy database files. However, if the passive copy `ReplayLagTime` parameter is set to a value greater than `0`, the log is applied with a lag, creating a lagged database copy. The maximum value is `14` days.
Truncation lag time	The `TruncationLagTime` parameter specifies how long the Microsoft Exchange Replication service waits before truncating log files that have already been applied to the database files. The maximum value is `14` days.

Configuring the integration

Prerequisites

- Ensure that you have correctly installed and configured the Microsoft Exchange Server environment.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
  - For information on installing, configuring, and using Microsoft Exchange Server, see the Microsoft Exchange Server documentation.
  - In the Microsoft Exchange Server 2010 environment, if you intend to use the restore method **Restore to a point in time**, make sure that you have the Microsoft Exchange Server 2010 SP1 installed.
- If you intend to run Incremental and Differential backup sessions, make sure that circular logging is disabled.
- 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*.

  Ensure that the following Data Protector components are installed on all Microsoft Exchange Server systems:
  - `MS Exchange Server 2010/2013 2010+ Integration`
  - `MS Volume Shadow Copy Integration`

  In DAG environments, the DAG virtual system (host) must also be imported to the Data Protector Cell. On how to import a client to a Data Protector Cell, see the *HP Data Protector Help* index: "importing, client systems".
- For limitations, see “Limitations and recommendations” in the *HP Data Protector Product Announcements, Software Notes, and References*.

Limitations

- Due to incompatibility between Microsoft Exchange Server versions, backup objects belonging to a particular Exchange Server version cannot be restored to Data Protector clients on which a different Exchange Server version is installed.

Before you begin

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

Configuring user accounts

Windows domain user account for backup and restore sessions

Backup and restore sessions are started by the `Data Protector Inet` service, which by default runs under the Windows local user account `SYSTEM`. Consequently, a backup or restore session is performed using the same user account.

However, you can specify that the `Data Protector Inet` service should use a different Windows domain user account to start a session:

- To perform a backup session under a different user account, specify the **Specify OS user** option (see "Specifying view type" (page 131)) when creating a backup specification.
- To perform a restore session under a different user account, specify the **User name** and **Group/Domain name** options in the Options page (see "Restore options" (page 148)).

Before you specify a different Windows domain user account, configure the user account as follows:

1. Grant the user appropriate permissions to back up and restore Microsoft Exchange Server databases.
2. Add the user to the Data Protector `admin` or `operator` user group. For details on adding users, see the *HP Data Protector Help* index: "adding users".
3. Save the user and its password to a Windows Registry on the Microsoft Exchange Server system on which you plan to start the integration agent (`e2010_bar.exe`). To save the user account, use the Data Protector `omniinetpasswd` or `omnicc` command.

NOTE: The user account saved in the Windows Registry will be used by the `Data Protector Inet` service when needed.

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

Example

To save the user `jane` from the domain `HP` and with the password `mysecret` to a Windows Registry, log on to the Microsoft Exchange Server system and execute the following command:

```
omniinetpasswd -add jane@HP mysecret
```

User account for executing Exchange Management cmdlet operations

In the Microsoft Exchange Server 2013 environment, you need user credentials with specific Exchange Management Roles assigned to create a remote runspace for executing the Exchange Management cmdlet operations remotely. These operations are executed as part of Microsoft Exchange Server backup and restore operations.

Configure a valid Exchange domain user account, when creating a backup specification. The user account is saved in the user credentials specific configuration file located in the

`Data_Protector_program_data\Config\Server\Integ\Config\E2010` directory and named by the domain name. The saved user credentials will be used by Data Protector when needed.

For details, see "Specifying view type" (page 131).

For information on the Exchange Management cmdlet operations, see the Microsoft Exchange Server documentation.

Backup

When you back up a Microsoft Exchange Server database, the following files are automatically backed up:

- database files (`.edb`)
- transaction logs (`.log`)
- checkpoint files (`.chk`)

However, depending on the Microsoft Exchange Server backup type you select, not all files are always backed up. For details, see "Microsoft Exchange Server backup types" (page 128).

Backup types

Microsoft Exchange Server backup types

You can select among the following Microsoft Exchange Server backup types:

Table 29 Backup types

Full	Backs up the database files (`.edb`), transaction logs (`.log`), and checkpoint files (`.chk`), and then truncates the transaction logs.
Copy	Backs up the database files (`.edb`), transaction logs (`.log`), and checkpoint files (`.chk`), without truncating the transaction logs.
Incremental	Backs up the transaction logs (`.log`) that have been created since the last Full or Incremental backup, and then truncates the transaction logs.
Differential	Backs up the transaction logs (`.log`) that have been created since the last Full backup, without truncating the transaction logs.

NOTE:
An incremental or differential backup of a database cannot be performed:

- If a full backup has not been performed.
- If an incremental backup is started just after a differential backup has been performed, or the other way around.
- If Microsoft Exchange Server circular logging is enabled.

Backup parallelism

- During a backup session, copies of different databases are backed up in parallel, however, copies of the same database are not, due to a Microsoft Exchange Server VSS writers limitation.
- If multiple backup sessions that intend to back up the same database are started in parallel, only the session that first locks the database can back up the database; the other sessions cannot. In DAG environments, this also applies if backup sessions intend to back up different copies of the same database; only the session that first locks the database (that is, all its copies) can back up the database copies; the other sessions cannot.

NOTE: This behavior ensures that the construction of a restore chain is valid. For example, suppose that several Full backup sessions that intend to back up the same database are started in parallel. If all the sessions backed up the database, it might happen that the session given the latest SessionID is not the one that backed up the database last. For details on restore chains, see "Restore chain" (page 141).

Backup considerations

- *Backup strategy:*

  Choose one of the following strategies to back up your data:

  - Full
  - Full, Incremental, Incremental, ...
  - Full, Differential, Differential, ...
  - Full, Copy, Incremental, ..., Copy, Incremental, ...

IMPORTANT: An Incremental backup session cannot be followed by a Differential backup session, nor the other way around. You must first run a Full backup session.

- *Active copies as opposed to passive copies:*

  There is no difference between the active and passive copy, except in the currently active log file (at the active copy side), which is not copied to the passive copy side until the file is filled up (that is, reaches 1 MB). Consequently, if you back up a passive copy, the transactions in the currently active log file are not included.

- *Lagged database copies:*

  Backing up a lagged database copy is equivalent to backing up a non-lagged database copy. If you restore from the backup of a lagged database copy, files are not only restored, but logs are also applied to the database file, returning the database to its most recent state. However, restoring the logs and applying them to the database file is time-consuming and, therefore prolongs the restore session. Also note that you need enough disk space to restore all the necessary logs.

  On the other hand, restoring from the backup of a lagged database copy enables you to restore the database to a point in time before the backup was taken. Restore the database without performing database recovery and mounting. Then remove unwanted logs, and finally recover and mount the database.

- *Public folders*:

  In the Microsoft Exchange Server 2010 environment, backup of Microsoft Exchange Server public folders with activated replication is not supported.

- *Concurrent backup sessions:*

  Backup sessions that back up the same database cannot run in parallel.

Object operations considerations

- Object copy and object verification

  When copying or verifying Microsoft Exchange Server objects you need to select all Data Protector backup objects created in the same session. To make sure that you do not select only a few objects from the session, the Data Protector GUI does not list Microsoft Exchange Server backup objects for interactive object copy or object verification sessions in the Objects scope of the Object Operations context.

  Use the Session or the Media scope instead.

Creating backup specifications

Create a backup specification using the Data Protector GUI (**Data Protector Manager**).

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange 2010+ Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. In **Application system**, select the Microsoft Exchange Server system that you want to back up. In a DAG environment, select the DAG virtual system or a Microsoft Exchange Server system.

   NOTE: The **Application system** drop-down list contains all clients that have the Data Protector `MS Exchange Server 2010/2013 2010+ Integration` component installed. In a DAG environment, the list contains also the DAG virtual system (host).

   The backup session (that is, the integration agent `e2010_bar.exe`) will be started on the client that you specify here. If you select a DAG virtual system, the integration agent is started on the currently active Microsoft Exchange Server node. To find out which node is currently active, see "Tip" (page 154).

   NOTE: In the Microsoft Exchange Server 2010 environment, to back up public folders residing on a Microsoft Exchange Server system that is a part of a DAG environment, select the Microsoft Exchange Server system and not the DAG virtual system (host). If you select the DAG virtual system, you can back up only databases that belong to the DAG. The Microsoft Exchange Server public folders database is not the part of it.

   Click **Next**.
5. If you selected the DAG virtual system (host), specify **View Type** to define how Microsoft Exchange Server databases should be organized in the next page (source page):

   By Role

   All databases in the DAG are displayed.

   By Client

   All clients in the DAG are displayed, together with all the databases (active or passive) residing on them. Active databases have the label `(active)` appended at the end. Passive databases have no label.

Figure 56 Specifying view type

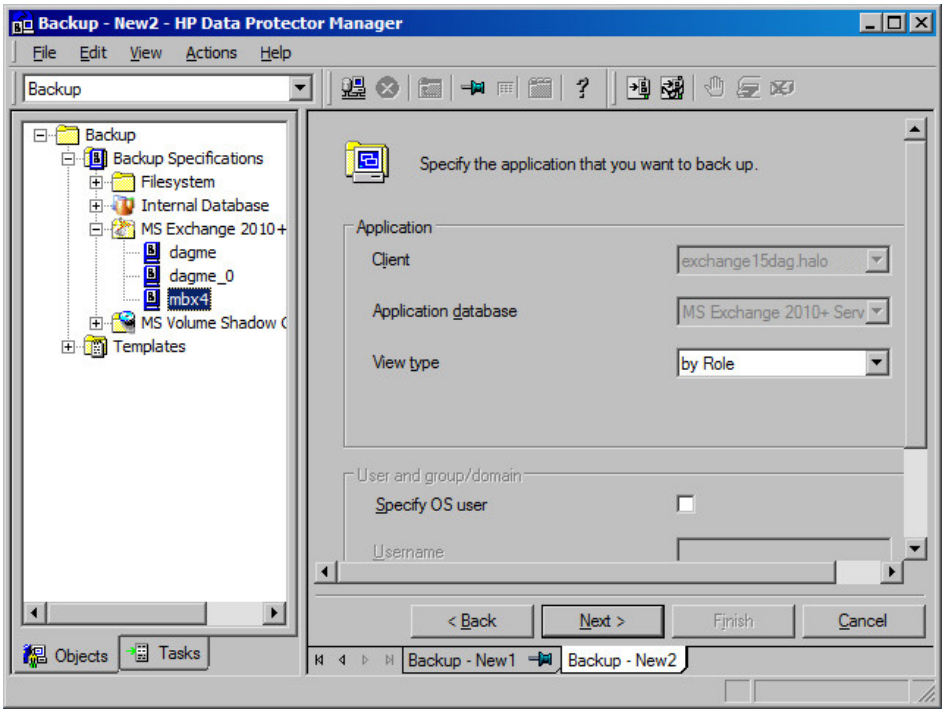

For information on the **User and group/domain** options, press **F1**.

NOTE: If no valid user credentials for executing the Exchange Management cmdlet operations remotely are specified, the Microsoft Exchange Server configuration dialog box is displayed. Enter the required user credentials and click **OK**.

Figure 57 User credentials for executing Exchange Management cmdlet operations

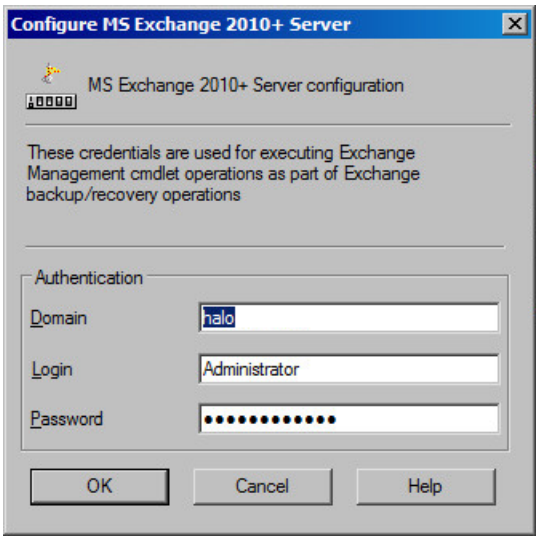

Click **Next**.

6. Select which Microsoft Exchange Server databases to back up.

Figure 58 Selecting databases (DAG environment – by role)

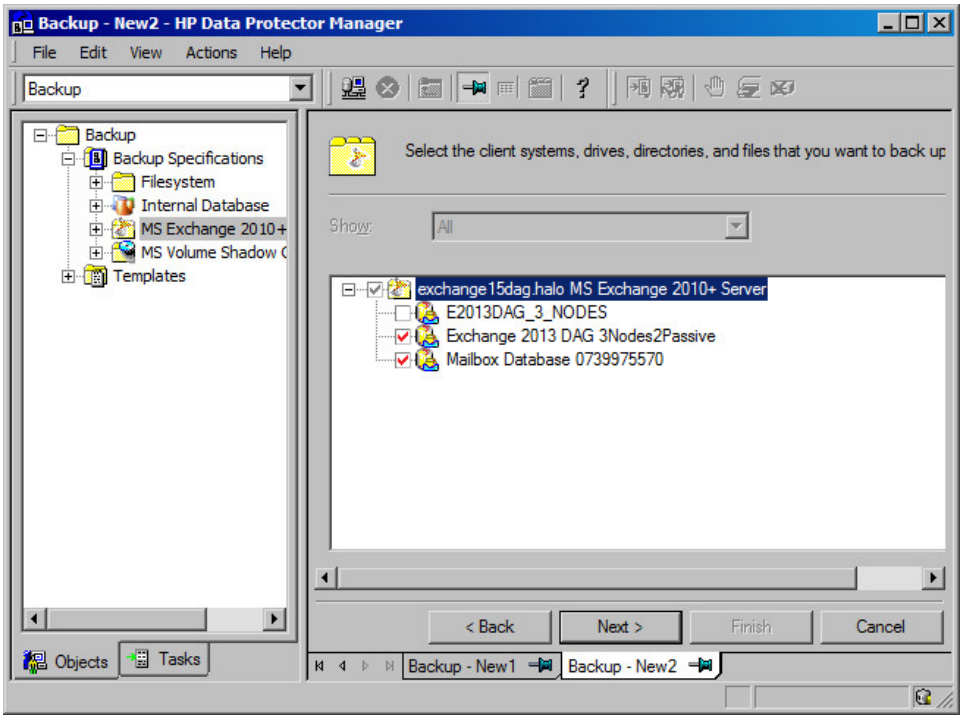

Figure 59 Selecting databases (DAG environment – by client)

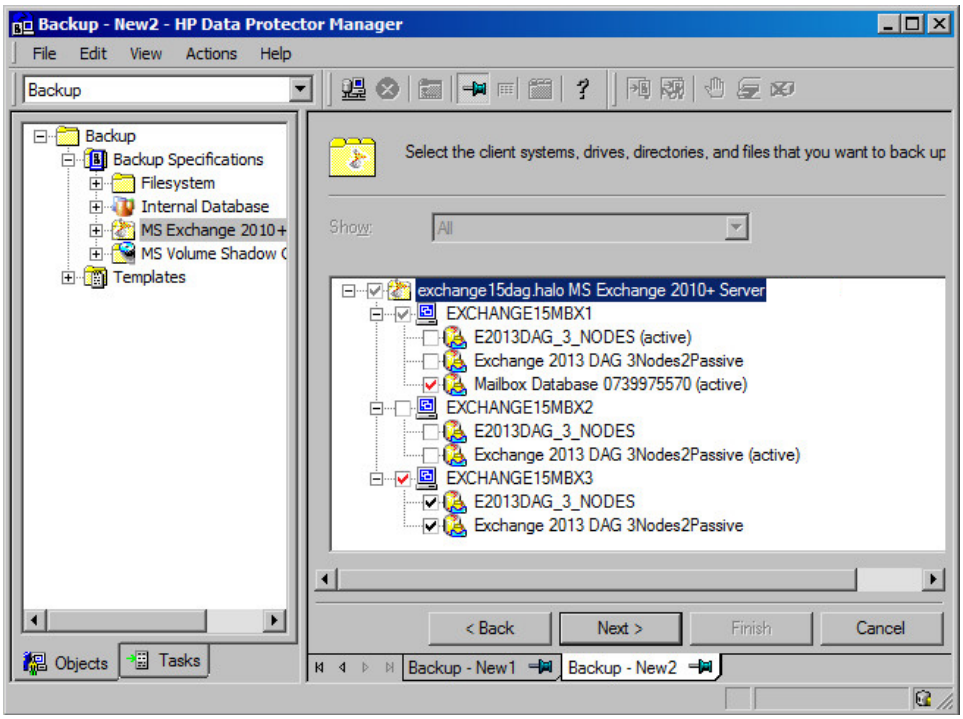

Figure 60 Selecting databases (standalone environment)

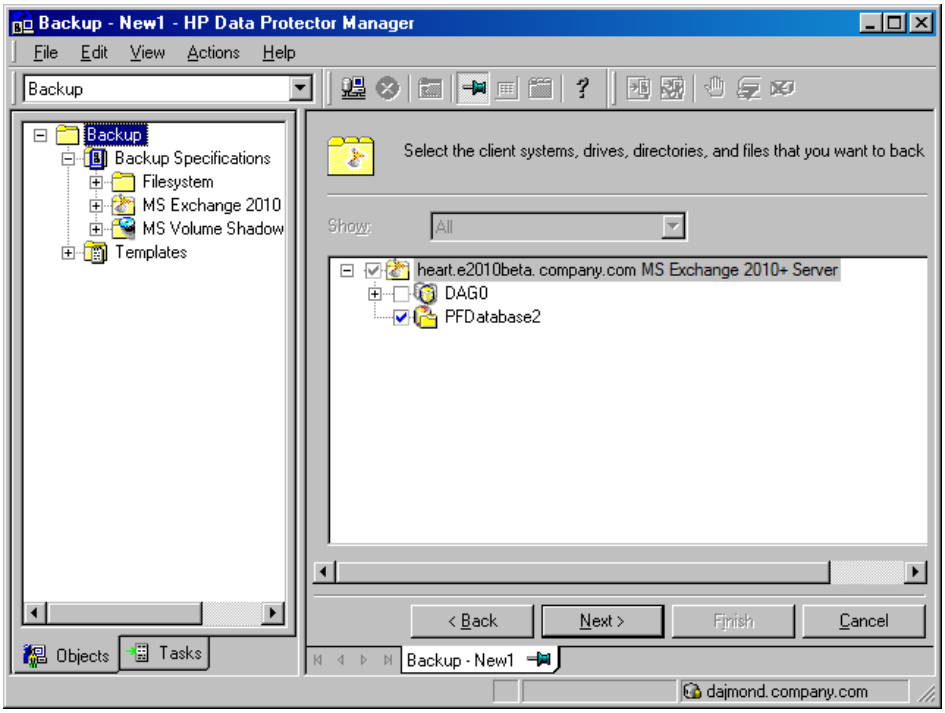

7. The following applies in DAG environments if you selected the **By Role** view type. Specify the backup policy options.

Figure 61 Backup policy options

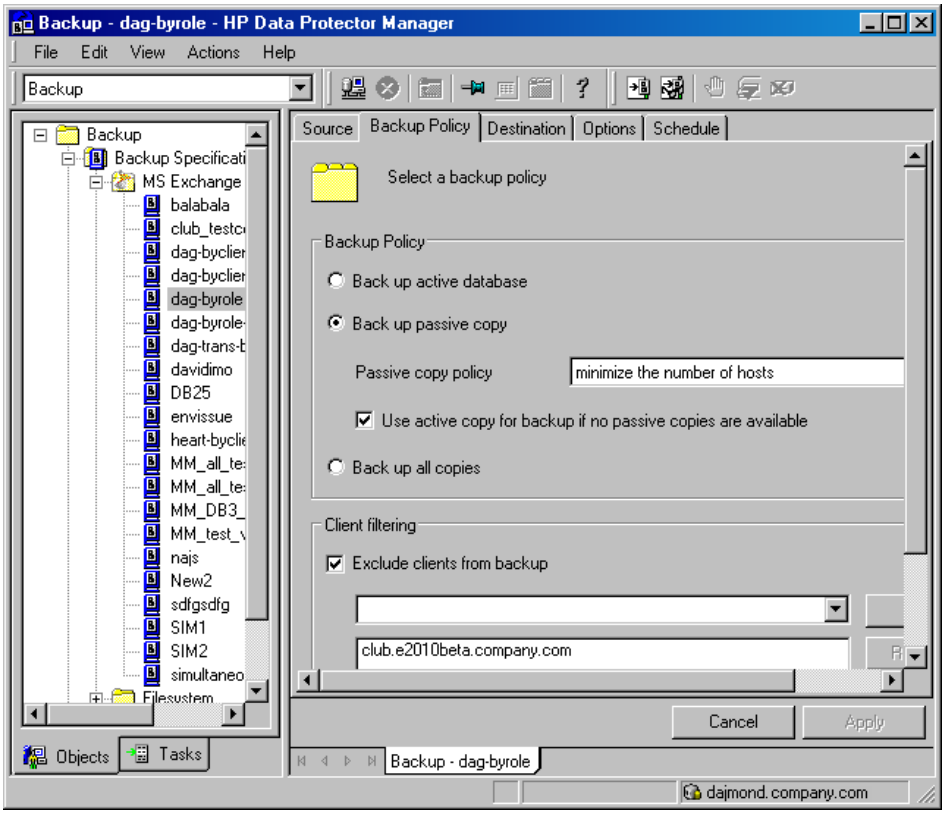

For details, see "Backup policy options" (page 134).

8. Select which devices to use for the backup.

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

   Click **Next**.

9. Set backup options.

Figure 62 Application-specific option

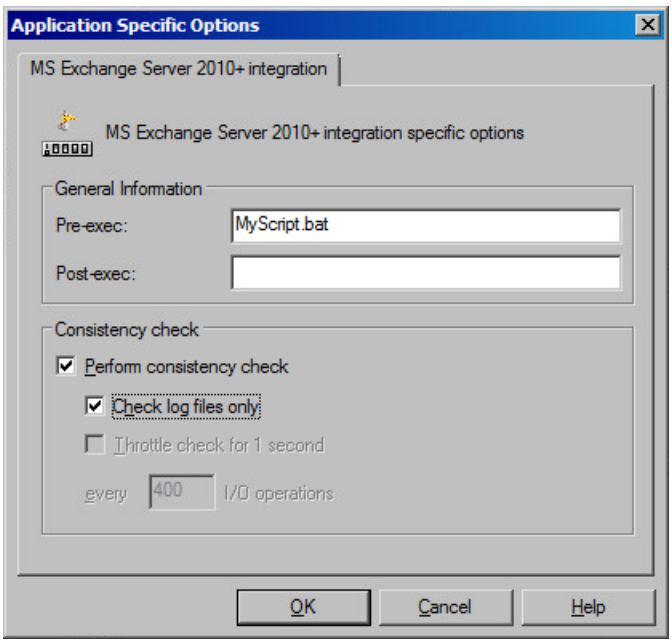

For information on application-specific backup options, see "Application-specific backup options" (page 135).

Click **Next**.

10. Optionally, schedule the backup. See "Scheduling backup sessions" (page 136).

    Click **Next**.

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

TIP: Preview your backup specification before using it for real. See "Previewing backup sessions" (page 137).

Table 30 Backup policy options

Options	Description	
Back up active database	If this option is selected, the active copy is backed up.	
Back up passive copy	If this option is selected, a passive copy is backed up. If a database has multiple passive copies, specify which particular copy you want to back up, using one of the following policies:	
	minimize the number of hosts (default)	If this option is selected, the minimum number of clients is involved in the backup. For example, if databases to be backed up have each a passive copy residing on the same client, they are all backed up from this client (and not one database from one client and another database from another client).
	lowest/highest activation preference	If this option is selected, the database copy with the lowest/highest activation preference number is backed up.

Table 30 Backup policy options *(continued)*

Options	Description	
	shortest/longest replay lag time	If this option is selected, the database copy with the shortest/longest replay lag time is backed up.
	longest truncation lag time	If this option is selected, the database copy with the longest truncation lag time is backed up.
	For a brief description of the activation preference number, replay lag time and transaction lag time parameters, see " Microsoft Exchange Server parameters in DAG environments" (page 126). For details, see the Microsoft Exchange Server documentation.	
	Use active copy for backup if no passive copies are available	Available if **Back up passive copy** is selected. If this option is selected, the active copy is backed up when no passive copy is available.
Back up all copies	Available if only one database is selected for backup. This option should only be used in ZDB environments. For details, see the *HP Data Protector Zero Downtime Backup Integration Guide*. Otherwise, it is enough that a single copy is backed up; you can restore different copies of a database from the backup of a single copy.	
Exclude clients from backup	Creates a list of clients. The database copies that reside on these clients are not backed up.	

Table 31 Application-specific backup options

Options	Description
Pre-exec, **Post-exec**	Specifies which command line to run on a Microsoft Exchange Server system before (`pre-exec`) or after (`post-exec`) the backup. The command line is executed only on the Microsoft Exchange Server system on which the backup session is started (that is the system on which the Data Protector Microsoft Exchange Server integration agent `e2010_bar.exe` is started). Type only the name of the command and ensure that the command is located in the *`Data_Protector_home`*`\bin` directory on the same system. Do not use double quotes. *DAG environment only:* If you selected the DAG virtual system (host) in the **Application system** option, ensure that the command is located on the currently active node. To find out which Microsoft Exchange Server node is currently active, see "Tip" (page 154).

Perform consistency check
`[-exch_check`
`[-exch_throttle` *`Value`*`] |`
`-exch_checklogs]`
	If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed.
	The check is performed on the backup media after the backup data is created. If the data is found corrupt, it is discarded and the database backup fails.
	Default: selected
	If the **Check log files only** option is selected, only the backup data of the log files is checked, which is enough for Microsoft Exchange Server to guarantee data consistency.
	Default: selected
	By default, the consistency check is I/O intensive, which can negatively affect disk performance. The **Throttle check for 1 second** option throttles down the consistency check of the database file `.edb` to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second.
	This option is not available if only the log files are checked.
	Default: not selected

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context.

In the Microsoft Exchange Server 2013 environment, in the Source page, you can change the Exchange domain user credentials for executing the Exchange Management cmdlet operations remotely by right-clicking the selected backup object and clicking **Configure**. You can also validate your configuration by clicking **Check configuration**.

Click other desired tabs, and apply the changes.

NOTE: To see all databases in the source page, not just those you selected, select **All** in the **Show** option. In a DAG environment, this not only shows all databases, but also updates the current status of databases (active or passive).

Scheduling backup sessions

You can schedule a backup session to start automatically at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: "scheduled backups".

Scheduling example

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

1. In the **Schedule** property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See "Scheduling backup sessions" (page 136). Under **Session options**, select **Differential** from the **Backup type** drop-down list.

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

Figure 63 Scheduling backup sessions

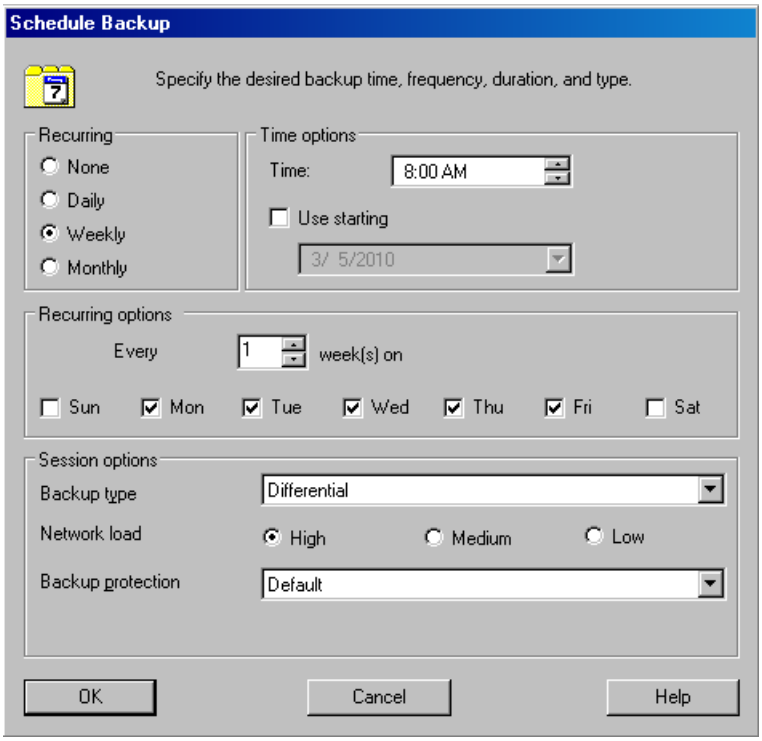

Previewing backup sessions

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 2010+ Server**. 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

1. Log in to the Cell Manager or to any client with the Data Protector `User Interface` component installed, under a user account that is configured as described in "Configuring user accounts" (page 127).
2. Execute the following command:

   ```
   omnib -e2010_list BackupSpecificationName -test_bar
   ```

What happens during the preview?

The following are tested:

- Communication between the Microsoft Exchange Server system on which the backup session is started and the Cell Manager
- If each selected database has at least one copy available for backup after the **Backup policy** options and **Client filtering** options have been applied (this applies to backup specifications that contain backup policy options)
- If the selected databases are ready to be backed up (that is, they should not be dismounted, suspended, or in a failed state)

Starting backup sessions

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

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

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS Exchange 2010+ Server**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

1. Log in to the Cell Manager or to any client that has the Data Protector `User Interface` component installed under a user account that is configured as described in "Configuring user accounts" (page 127).

2. Execute the following command:

   ```
   omnib -e2010_list BackupSpecificationName [-barmode E2010Mode]
   [LIST_OPTIONS]
   ```

   where *E2010Mode* is one of the following:

   ```
   full|copy|incr|diff
   ```

   The default is `full`.

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

   ### Examples

   To start a Full backup using the backup specification `MyDatabases`, execute:

   ```
   omnib -e2010_list MyDatabases -barmode full
   ```

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

   ```
   omnib -e2010_list MyDatabases -barmode diff
   ```

Backup objects

For each database (copy), Data Protector creates the following backup objects:

- *Database file object*
  - *ClientName*:/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/*DBID*/File [MSVSSW-APP]

    (standalone database or active copy)
  - *ClientName*:/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/*DBID*/File [MSVSSW-APP]

    (passive copy)
- *Log file object*
  - *ClientName*:/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/*DBID*/Logs [MSVSSW-APP]

    (standalone database or active copy)
  - *ClientName*:/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/*DBID*/Logs [MSVSSW-APP]

    (passive copy)
- *Database object*

  ClientName:/*DBID*/*DBName* [E2010]

  The database object contains information needed to construct the restore chain. For details on restore chains, see "Restore chain" (page 141).
- *VSS metadata object*

  /BackupSession/Metadata [MSVSSW-APP]

Information on whether the objects were successfully backed up or not is saved in the Data Protector IDB. On how to retrieve the information from the IDB, see "Finding information for restore" (page 142).

Restore

You can restore Microsoft Exchange Server data by performing a standard restore session. For details, see "Restore procedure" (page 143).

IMPORTANT: After you have restored a database, start a Full backup session for the database. Otherwise, the subsequent Incremental and Differential backup sessions will fail.

Considerations

- A Microsoft Exchange Server database that was backed up using the Data Protector Microsoft Volume Shadow Copy Service integration cannot be restored using the Data Protector Microsoft Exchange Server 2010 integration, nor the other way round.

Restore methods

There are various reasons for restoring a Microsoft Exchange Server database. Here are some examples:

- The database has become corrupt.
- The synchronization between an active and passive database copy is broken, but you want to avoid reseeding the passive copy, or simply the resume operation does not work.
- The database needs to be restored to a different point in time.
- The database's backup data needs to be restored for investigation purposes.
- The database's backup data needs to be restored to a recovery database in order to extract individual mailboxes or mailbox files.
- The database's backup data needs to be restored to a dial tone database.

To suit your needs, Data Protector offers different restore methods. You can choose among the following:

- **Repair all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**
- **Restore to a new mailbox database**
- **Restore to a temporary location**

You can specify different restore methods for different databases in the same session.

NOTE: The first three methods restore backup data to the original database and are therefore only available if the original database still exists. The last two methods restore backup data to a new location.

Repair all passive copies with failed status

This method is available only for databases that are part of a DAG. It is useful if some of a database's passive copies become corrupt, acquiring the status `Failed` or `FailedAndSuspended`. The method automatically restores all the corrupt passive copies from the backup created in the last backup session (and the corresponding restore chain). After the data is restored, the copies are synchronized with the active copy, provided that the **Resume database replication** option is selected.

Restore to the latest state

This method is used to restore a corrupt database to the latest possible point in time. Data Protector restores the database from the backup created in the last backup session (and the corresponding restore chain). For details, see "Restore chain" (page 141).

Once the files are restored, all the logs (not only those restored from the backup, but also any existing logs) are replayed to the database file.

NOTE: ***DAG environments:***

When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `ReplayLagTime` parameter setting.

Restore to a point in time

This method is used to restore a database to a specific point in time.

NOTE: When you restore a standalone database or active copy, the existing `.log` and `.chk` files are renamed (a `.keep` extension is added to their names). This feature is useful when you restore files without performing database recovery. It enables you to apply additional logs to the database file; just delete the `.keep` extension of the log files that you also want to be applied and start a database recovery manually. In this way, you can fine-tune the point in time the database is restored to.

When you restore a passive copy, the existing files are deleted.

Once the files are restored, the logs are replayed to the database file (`.edb`) if the **Perform database recovery** option is selected.

NOTE: ***DAG environments:***

- When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `ReplayLagTime` parameter setting.
- For passive copies that are not restored, a full reseed is required once the restore session completes.

Restore to a new mailbox database

This method is used to restore data to a different database, either because the original database no longer exists or in order to move the data elsewhere.

Using it, you can restore data also to a Microsoft Exchange Server recovery database.

Restore files to a temporary location

Using this method, you can restore database files to a location of your choice.

- When you restore from a Differential or Incremental backup session, you can restore the complete restore chain or only the files (`.log`) backed up in the selected session.
- When you restore data from a Full backup session, you have an option to restore only the database file (`.edb`).

Restore destination

Backup data can be restored:

- to an existing database (standalone database, active copy, passive copy),
- to a new database,
- to a temporary location.

Restoring to a standalone database

Restore to the original standalone database (standalone environment) progresses as follows:

1. The database is dismounted.
2. Backup data is restored.

3. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file `.edb` and the database is mounted.

To restore to the original standalone database, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

Restoring to an active copy

Restore to the active copy (DAG environment) progresses as follows:

1. The database is dismounted.
2. All replications are suspended.
3. Backup data is restored.
4. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file `.edb` and the database is mounted.

To restore to the active copy, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

Restoring to a passive copy

Restore to a passive copy (DAG environment) progresses as follows:

1. The replication is suspended.
2. Backup data is restored.
3. Optionally, the replication with the active copy is resumed.

To restore to a passive copy, use one of the following restore methods:

- **Restore all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**

Restoring data to a new database

Restore to a new database progresses as follows:

1. A new mailbox database is created.
2. Backup data is restored to the new database.

NOTE: If you restore to a recovery database, first the backup data is restored and then a recovery database is created.

To restore data to a new mailbox database or recovery database, use the **Restore to a new mailbox database** restore method.

Restoring data to a temporary location

You can restore the database file (`.edb` and/or `.log` and/or `.chk`) to a client and directory of your choice. Select the **Restore files to a temporary location** restore method.

Restore chain

By default, when you select a Differential or Incremental backup session for restore, Data Protector restores not only the logs (`.log`) backed up in the selected session but also files backed up in preceding sessions (**restore chain**):

- If a Differential backup session is selected, Data Protector restores:
  1. The `.edb` file and `.log` files backed up in the most recent Full or Copy backup session.

2. The `.log` files backed up in the selected Differential backup session.

- If an Incremental backup session is selected, Data Protector restores:
  1. The `.edb` file and `.log` files backed up in the most recent Full or Copy backup session.
  2. The `.log` files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session.
- If a Full or Copy backup session is selected, Data Protector restores the `.edb` file and `.log` files backed up in the selected session.

NOTE:

- If the **Restore to the latest state** method is used, `.log` files from the Full or Copy backup session are not restored.
- The only method that enables you to restore only `.log` files backed up in the selected Incremental or Differential session is **Restore to a temporary location**.

Restore parallelism

If device concurrency allows, database copies are restored in parallel, except in the following cases:

- If database copies were backed up from the same client, but are now restored to different clients.
- If backup data of the same database copy is used as a restore source for multiple database copies.

Finding information for restore

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

To retrieve information, use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Internal Database**.
2. In the Scoping pane, expand **Objects** or **Sessions**.

   If you expand **Objects**, backup objects are sorted according to the Microsoft Exchange Server databases for which they were created.

   NOTE: The backup object name contains the database GUID. To find out which GUID belongs to which database, see the *database object* `/DB_GUID/DB_Name`.

   For example, the *database object* for the database `DB1` with the GUID `08bca794-c544-4e27-87e8-533fb81fd517` is:

   `/08bca794-c544-4e27-87e8-533fb81fd517/DB1`

If you expand **Sessions**, backup objects are sorted according to the sessions in which they were created. For example, backup objects created in the session `2013/02/7–7` are listed under `2013/02/7–7`.

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

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

Using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the Data Protector `User Interface` component installed under a user account that is configured as described in "Configuring user accounts" (page 127).
2. Get a list of Microsoft Exchange Server backup objects created in a backup session:

   ```
   omnidb -session SessionID
   ```
3. Get details on a backup object:

   ```
   omnidb -e2010 BackupObjectName -session SessionID -catalog
   ```

   Here is one example of a backup object name:

   ```
   devy.company.com:/08bca794-c544-4e27-87e8-533fb81fd517/DB1
   ```

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

Restore procedure

You can restore multiple Microsoft Exchange Server databases in the same session, specifying a different restore method for each database. For details, see "Restore methods" (page 139).

To restore databases, use the Data Protector GUI or CLI.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Exchange 2010+ Server**, expand the DAG virtual system or standalone Microsoft Exchange Server system and click **MS Exchange 2010+ Server**.
3. In the Source page, Data Protector displays all Microsoft Exchange Server databases backed up from the selected DAG or standalone environment.

   Select which Microsoft Exchange Server databases to restore.

   When you select a database, the Properties for Database dialog box is displayed automatically. Specify a restore method and click **OK**. For databases that are part of a DAG, the default restore method is **Repair all passive copies with failed status**. For standalone databases, the default is **Restore to the latest state**.

Figure 64 Selecting databases for restore

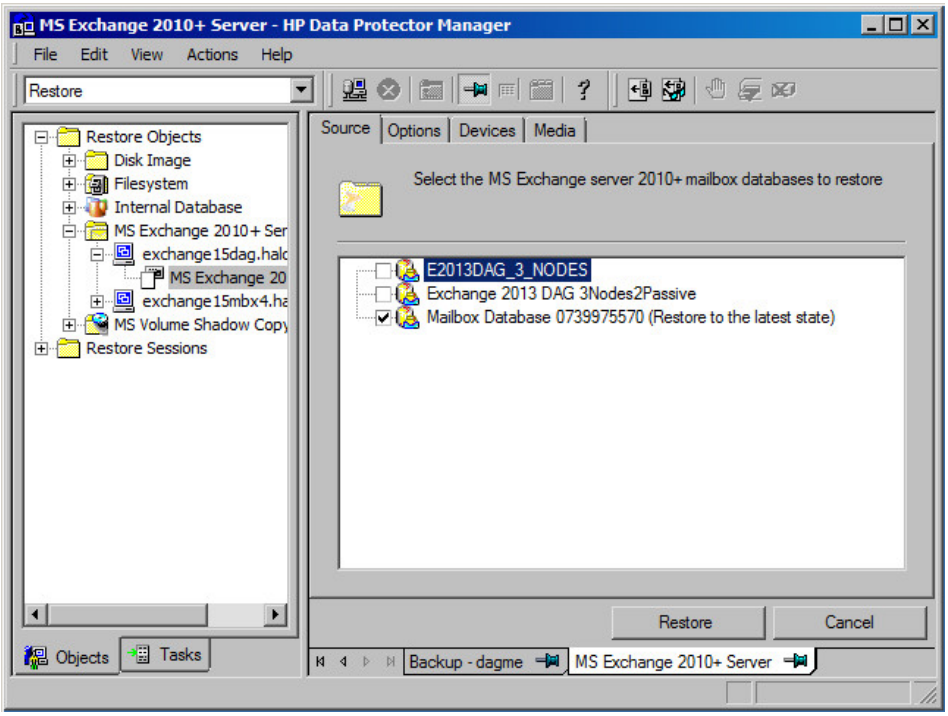

To change the restore method, right-click the database and click **Properties**.

Figure 65 Repair all passive copies with failed status

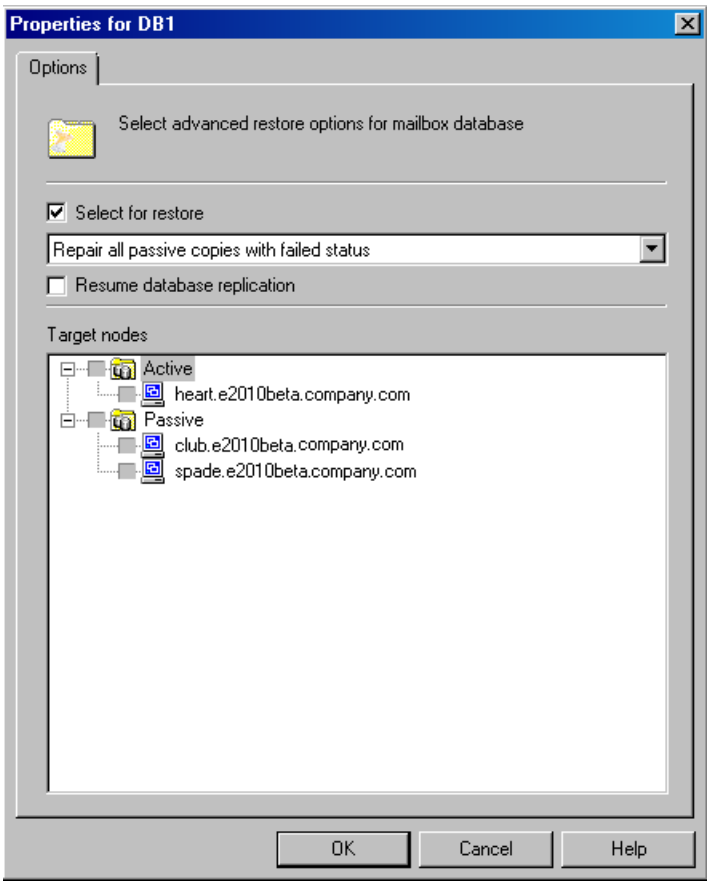

For details, see "Repair all passive copies with failed status" (page 151).

Figure 66 Restore to the latest state

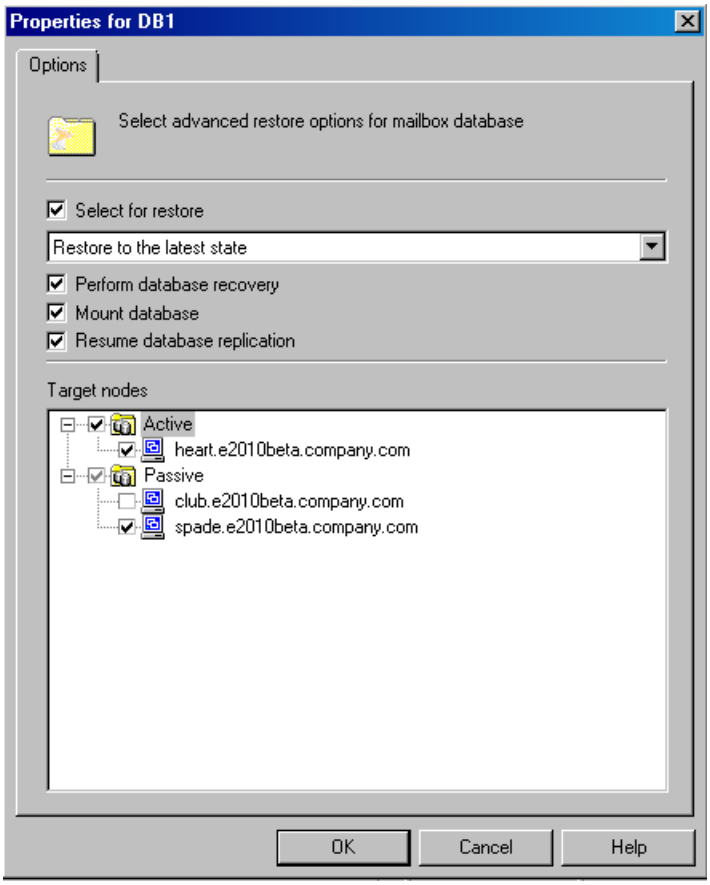

For details, see “Restore to the latest state” (page 139).

Figure 67 Restore to a point in time

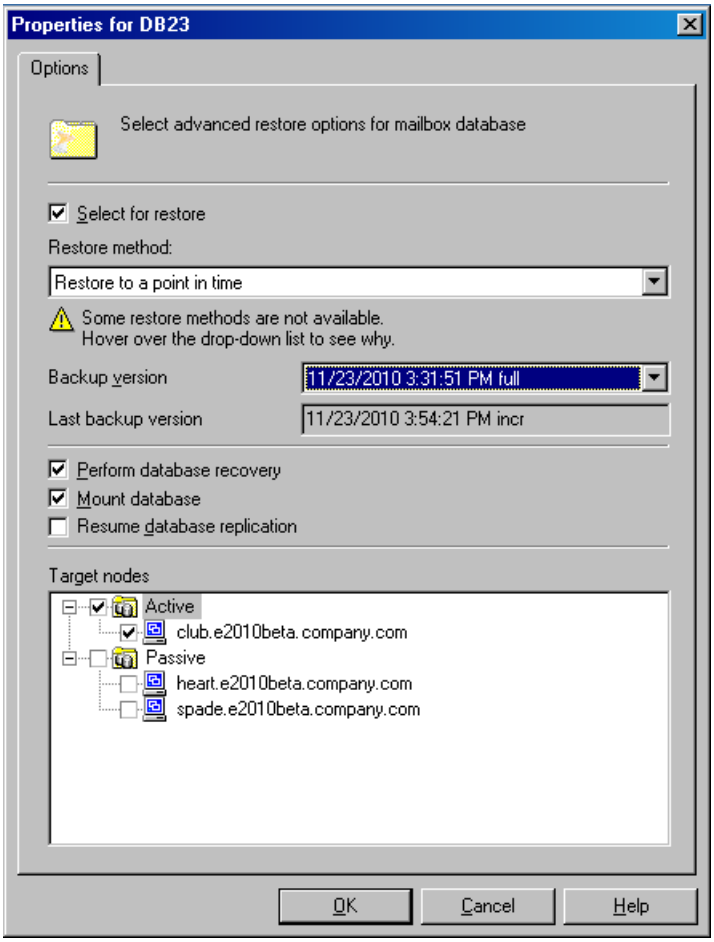

For details, see "Restore to a point in time" (page 140).

Figure 68 Restore to a recovery database

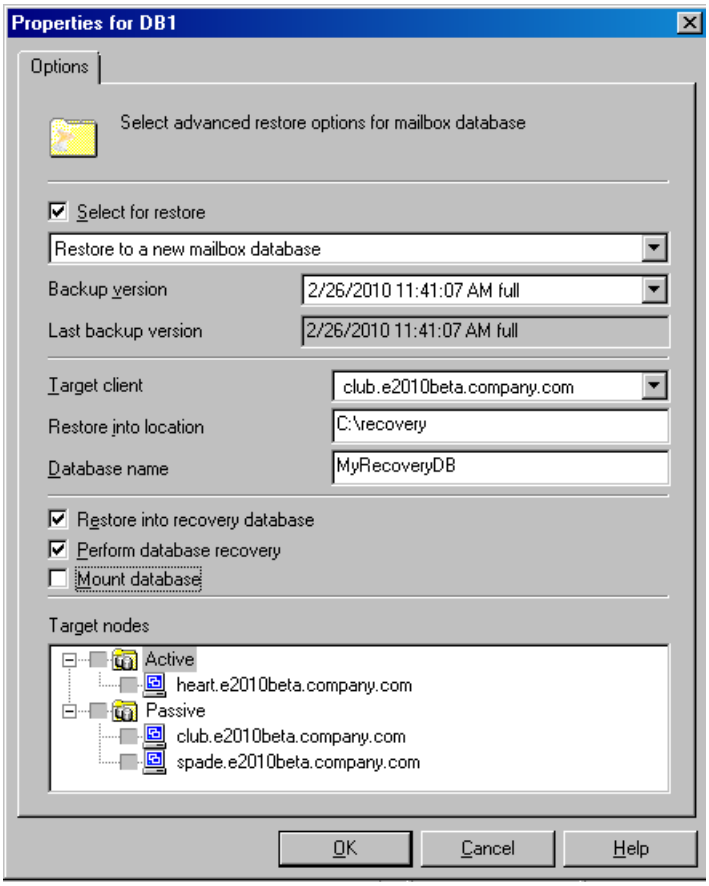

For details, see "Restore to a new mailbox database" (page 140).

Figure 69 Restore files to a temporary location

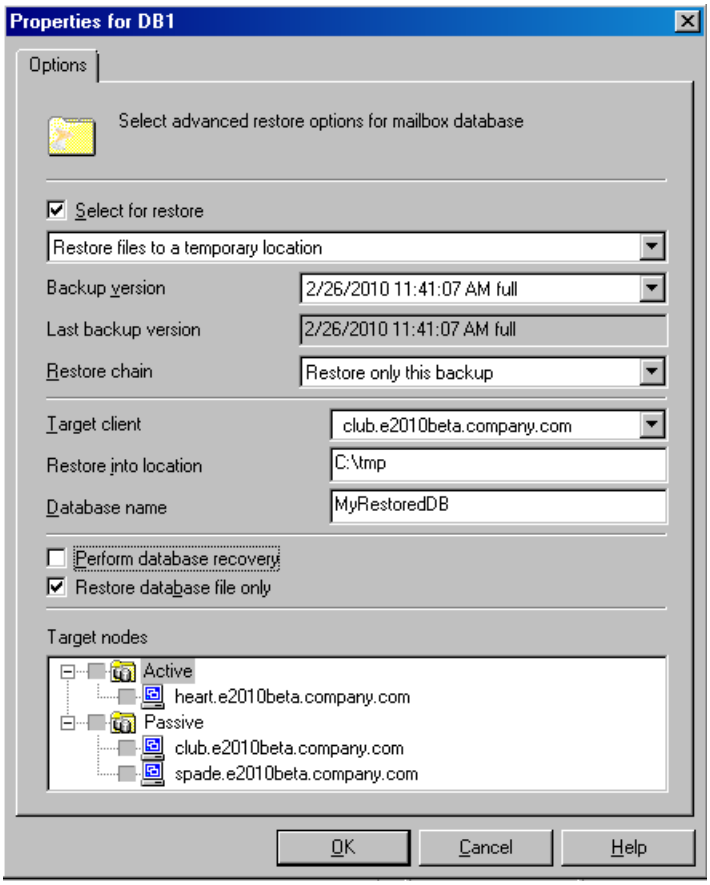

For details, see "Restore files to a temporary location" (page 140).

4. In the **Options** page, specify the Data Protector Microsoft Exchange Server 2010 integration restore options. For details, see Table 37 (page 154).

Figure 70 Restore options

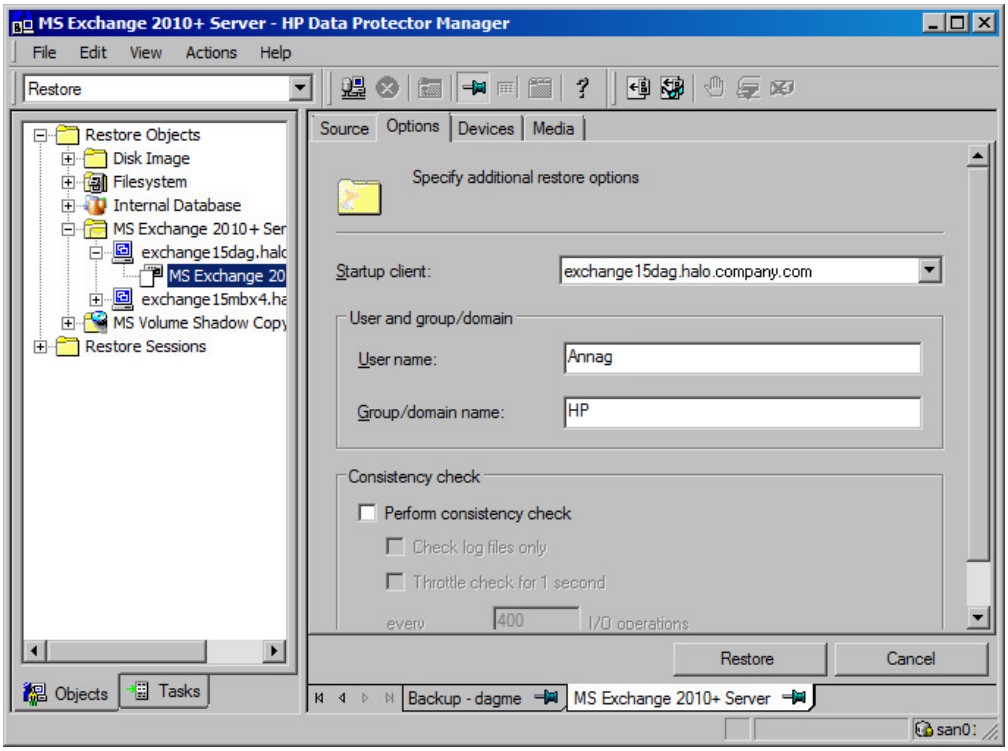

5. In the **Devices** page, select which devices to use for restore.

   For details on how to select devices to be used for restore, see the *HP Data Protector Help* index: "restore, selecting devices for".
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. Specify **Report level** and **Network load**.

   Click **Finish** to start the restore.

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

Restoring using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the `User Interface` component installed under a user account that is configured as described in "Configuring user accounts" (page 127).
2. Execute the following:

```
omnir -e2010
-barhost ClientName
[VSS_EXCHANGE_SPECIFIC_OPTIONS]
Database [Database …]
[-user User:Domain]
[GENERAL_OPTIONS]

Database
{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID }
[-source SourceClientName]
{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS

E2010_REPAIR_METHOD_OPTIONS
[-no_resume_replication]

E2010_LATEST_METHOD_OPTIONS
[-node TargetNode … | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]

E2010_PIT_METHOD_OPTIONS
-session BackupID
[-node TargetNode … | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]
```

```
E2010_NEW_METHOD_OPTIONS
-session BackupID
-client TargetClientName
-location TargetDatabasePath
-name TargetDatabaseName
[-recoverydb]
[-no_recover]
[-no_mount]

E2010_TEMP_METHOD_OPTIONS
-session BackupID
-client TargetClientName
-location TargetDatabasePath
[-no_chain]
[-edb_only]
[-no_recover]
```

For a brief description of the options, see "Restore options" (page 151). For details, see the `omnir` man page or the *HP Data Protector Command Line Interface Reference*.

NOTE: A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.

The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

Example (Restore method – repair)

DAG environment

To restore all corrupt passive copies of the database `DB1`, which was backed up from a DAG whose virtual system name was `dag0.company.com`, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, execute:

```
omnir -e2010 -barhost exchange2.company.com –db_name DB1 –source
dag0.company.com -repair
```

Example (Restore method – latest)

Standalone environment

To restore the corrupt standalone database `DB1`, which resides on the client `exchange1.company.com`, to the latest possible point in time, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, execute:

```
omnir -e2010 -barhost exchange2.company.com –db_name DB1 –source
exchange1.company.com –latest
```

DAG environment

Suppose you want to restore the active copy of the database `DB1`, which resides on the client `exchange1.company.com`, and the passive copies of the database that reside on the clients `exchange2.company.com` and `exchange3.company.com`. Suppose the database `DB1` is part of a DAG whose virtual system name is `dag0.company.com`, and that you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange2.company.com`. Execute the following command:

```
omnir -e2010 -barhost exchange2.company.com –db_name DB1 –source
dag0.company.com –latest –node exchange1.company.com –node
exchange2.company.com –node exchange3.company.com
```

Example (Restore method – pit)

Standalone environment

Suppose you want to restore the corrupt standalone database `DB1`, which resides on the client `exchange1.company.com`, using the backup data created in the session `2013/5/14-1`. Suppose you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange1.company.com`. Execute the following command:

```
omnir -e2010 -barhost exchange1.company.com –db_name DB1 -pit –session
2013/5/14-1
```

NOTE: The `–source` option is not specified, in which case Data Protector assumes that the database was backed up from the client specified with the `–barhost` option.

Example (Restore method – new)

DAG environment

Suppose you want to restore the backup of the database `DB1` to a recovery database that should be created on the client `exchange2.company.com` and named `Recovery1`, with the files in the `C:\Recovery1Folder` directory. Suppose the database `DB1` was backed up in the session `2013/5/14-1` from a DAG whose virtual system name was `dag0.company.com`. To also ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, execute the following command:

```
omnir -e2010 -barhost exchange1.company.com –db_name DB1 –source
dag0.company.com –new –session 2013/5/14-1 –client exchange2.company.com
-location C:\Recovery1Folder –name Recovery1 –recoverydb
```

Example (Restore method – temp)

Standalone environment

Suppose you want to restore the transaction logs of the database `DB1`, which resides on the client `exchange2.company.com`. The logs were backed up in the Incremental backup session `2013/5/14-1`. To restore the logs to the client `exchange2.company.com` to the directory `C:\DB1TransactionLogFolder` without performing database recovery, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, execute the following command:

```
omnir -e2010 -barhost exchange1.company.com –db_name DB1 –source
exchange2.company.com –temp –session 2013/5/14-1 –client
exchange2.company.com –location C:\DB1TransactionLogFolder -no_chain
–no_recover
```

Restoring using another device

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

Restore options

Table 32 Repair all passive copies with failed status

Option GUI / CLI	Description
Resume database replication / `-no_resume_replication`	Available in DAG environments. Resumes the replication between the active and passive copies after the copies are restored.

Table 32 Repair all passive copies with failed status *(continued)*

Option GUI / CLI	Description
	Note that the CLI option `-no_resume_replication` has the opposite meaning. If it is specified, the replication in not resumed.
Target nodes	Not available. The clients (that is, copies) that have the status `Failed` or `FailedAndSuspended` are automatically selected.

Table 33 Restore to the latest state

Option GUI / CLI	Description
Select for restore	Specifies whether the database should be restored.
Perform database recovery / `-no_recover`	Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Applies the logs to the database file (`.edb`) after the restore completes. Note that the CLI option `-no_recover` has the opposite meaning. If it is specified, the database recovery in not performed.
Mount database / `-no_mount`	Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Mounts the database after the database recovery completes. This option is available only if **Perform database recovery** is selected. Note that the CLI option `-no_mount` has the opposite meaning. If it is specified, the database is not mounted.
Resume database replication / `-no_resume_replication`	Available when restoring passive copies (DAG environment). Resumes the replication between the active and passive copies after the copies are restored. Note that the CLI option `-no_resume_replication` has the opposite meaning. If it is specified, the replication is not resumed.
Target nodes `-node | -all`	Available only in DAG environments. Specifies which clients (that is, database copies) to restore.

Table 34 Restore to a point in time

Option GUI/CLI	Description
Select for restore	See the description in " Restore to the latest state" (page 152).
Backup version / `-session`	Specifies from which backup data to restore. Select a backup ID. If a Differential backup session is selected, the `.log` files backed up in the selected Differential backup session are restored. If an Incremental backup session is selected, the `.log` files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session, are restored.
Last backup version	Shows the session in which the database was last backed up.
Perform database recovery / `-no_recover`	See the description in " Restore to the latest state" (page 152).
Mount database / `-no_mount`	
Resume database replication /	

Table 34 Restore to a point in time *(continued)*

Option GUI/CLI	Description
`-no_resume_replication`	
Target nodes / `-node | -all`	See the description in " Restore to the latest state" (page 152). The node (client) hosting the active copy is automatically selected for restore.

Table 35 Restore to a new mailbox database

Option GUI/CLI	Description
Select for restore	See the description in " Restore to the latest state" (page 152).
Target client / `-client`	Specifies the client to restore to.
Restore into location / `-location`	Specifies the directory to restore to.
Database name / `-name`	Specifies the name to be used for the new database. If another database with the same name already exists, the restore fails.
Restore into Recovery database / `-recoverydb`	Restores the data to a Microsoft Exchange Server recovery database. Although multiple recovery databases can exist in parallel, only one recovery database can be mounted to the Microsoft Exchange Server at a time.
Backup version / `-session`	See the description in Table 34 (page 152).
Last backup version	
Perform database recovery / `-no_recover`	See the description in " Restore to the latest state" (page 152).
Mount database / `-no_mount`	
Target nodes	Not available.

Table 36 Restore files to a temporary location

Option GUI/CLI	Description
Select for restore	See the description in " Restore to the latest state" (page 152).
Restore chain	If this option is set to **Restore only this backup**, only files backed up in the selected session are restored. If this option is set to **Full restore (full, incr, diff backups)**, the complete chain is restored.
Target client / `-client`	See the description in " Restore to a new mailbox database" (page 153).
Restore into location / `-location`	
Backup version / `-session`	See the description in " Restore to a point in time" (page 152).
Last backup version	

Table 36 Restore files to a temporary location *(continued)*

Option GUI/CLI	Description
Restore database files only / `-edb_only`	Restores only the database files (`.edb`). The logs (`.log`) and checkpoint files (`.chk`) are not restored.
Perform database recovery / `-no_recover`	See the description in " Restore to the latest state" (page 152).
Target nodes	Not available.

Table 37 General restore options

Option GUI / CLI	Description
Startup client / `-barhost`	Specifies the client on which the integration agent (`e2010_bar.exe`) should be started. If the DAG virtual client (host) is selected, the integration agent is started on the currently active node. To find out which node is currently active, see "Tip" (page 154). Default: The same client that was specified for the backup session. If the DAG virtual client was specified, this client is now selected. However, note that the integration agent may not be started on the same physical node as during the backup session; it depends which node is currently active.
Username **Group/Domain name** / `-user`	Specifies which Windows domain user account to use for the restore session. Ensure that the user is configured as described in "Configuring user accounts" (page 127). If these options are not specified, the restore session is started under the user account under which the `Data Protector Inet` service is running.
Perform consistency check / `[-exch_check [-exch_throttle` *Value*`] | -exch_checklogs]`	If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed. The check is performed at the target location after the backup data is restored. You do not need to perform the consistency check if it was already performed at the time of backup. Default: not selected If the **Check log files only** option is selected, only the log file backup data is checked, which is enough for Microsoft Exchange Server to guarantee data consistency. Default: not selected By default, the consistency check is I/O intensive, which can negatively affect disk performance. The **Throttle check for 1 second** option throttles down the consistency check of the database file `.edb` to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second. This option is not available if only the log files are checked. Default: not selected

TIP: To find out which Microsoft Exchange Server node is currently active, connect to one of the nodes and run:

```
cluster group
```

Example

```
C:\Users\administrator.E2010BETA>cluster group
Listing status for all available resource groups:

Group                Node            Status
```

```
-------------------- --------------- ------
Available Storage    spade           Offline
Cluster Group        club            Online
```

The currently active node has the status `Online`. In the example, this is `club`.

Monitoring sessions

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

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

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

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft Exchange Server 2010 integration.

Because the Data Protector Microsoft Exchange Server 2010 integration is based on the Data Protector Microsoft Volume Shadow Copy Service integration, also see troubleshooting information in the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

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

Before you begin

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

Checks and verifications

If your browsing, backup, or restore failed:

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

Problems

Problem

It takes a long time to display Microsoft Exchange Server topology in the Data Protector GUI

When you open the Data Protector GUI and try to display the source page, either in the Backup or Restore context, you must wait a long time.

This may happen if there is an unresponsive system in the same domain (for example, a system that is shut down). The problem occurs even if the unresponsive system is not part of your backup environment. This is due to Microsoft Exchange Server problems with execution of Microsoft Exchange Server Shell commands.

Action

Remove the system from the domain or fix the problem.

Problem

A database backup cannot be performed

When you start a backup session for a database, the database is not backed up, appearing to be locked by other session, though there are no other backup sessions currently running. A message similar to the following is displayed:

```
[Minor] From: OB2BAR_E2010_BAR@exch03.e2010.company.com "MS Exchange
 Server"  Time: 1/17/2013 3:07:13 PM
[170:313]  One or more copies of database DEMAR are already being
backed up in a different session.
```

This may happen if the integration agent (`e2010_bar.exe`) was terminated by force while a previous backup session was in progress, either because the Microsoft Exchange Server system was restarted or for some other reason, so the lock remains.

Action

Execute the following command:

```
omnidbutil -free_cell_resources
```

NOTE: This command line removes all existing locks, so ensure that none of the existing locks is still needed.

Problem

Restore fails

When you try to restore a database, the session fails.

This may happen if a database has been restored before (probably unsuccessfully), and during that previous restore session, the Microsoft Exchange Server created an `.env` file in the database directory. This file now prevents the database from being restored again.

Action

Delete the `.env` file and start a new restore session.

Problem

Restore from an object copy fails in a DAG environment

When restoring a database from a media set created in an object copy session, as the media set created in the original backup session no longer exists, the session fails with an error similar to the following:

```
[Critical] From: OB2BAR_E2010_BAR@computer1.company.com "MS Exchange
2010 Server" Time: 28/02/2013 16:08:12 No mailbox database copy can be
selected for restore/instant recovery.
```

Action

1. Verify that the media set created in the object copy session still exists.
2. On all Microsoft Exchange Server system nodes, set the environment variable `OB2BARHOSTNAME` to the name of the DAG virtual system and restart the Data Protector Inet service.
3. Start a new restore session.

Problem

Restore to the latest state fails

When you try to restore a database all of whose log files were lost, using the restore method **Restore to the latest state** with the **Perform database recovery** option selected, the database recovery fails.

This may happen if a database is restored from a Full backup (that is, the restore chain consists of only the Full backup session). Since in the **Restore to the latest state** session, only the `.edb` file is restored from the Full backup (see "Restore chain" (page 141)), when the database recovery is started, there are no logs to be applied to the database file, and the database recovery fails.

Action

Restore the database using the restore method **Restore to a point in time**. For details, see "Restore to a point in time" (page 140).

6 Data Protector Microsoft Exchange Single Mailbox integration

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 38 Microsoft Exchange Single Mailbox integration 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 the *HP Data Protector 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. "Microsoft Exchange Single Mailbox integration architecture" (page 159) shows the architecture of the Data Protector Exchange Single Mailbox integration.

Figure 71 Microsoft Exchange Single Mailbox integration architecture

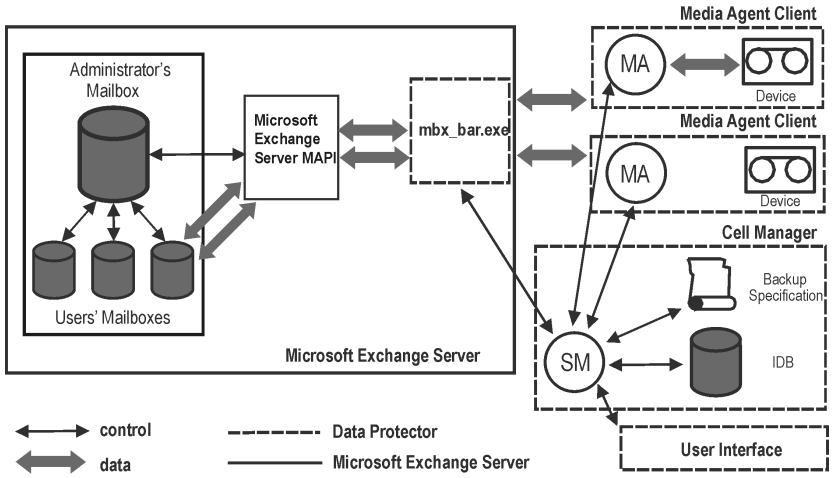

Legend:

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, `mbx_bar.exe` 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 latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
  - For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- On Microsoft Exchange Server 2007 systems, ensure that:
  - Microsoft Exchange Server MAPI Client and Collaboration Data Objects are installed.

    The installation package, which provides both components, can be obtained free of charge from the Microsoft website http://www.microsoft.com/downloads/Search.aspx?displaylang=en.
  - Microsoft Office Outlook is not installed.
- 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 *HP Data Protector 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.

   For information on the **User and group/domain** options, press **F1**.

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

Figure 72 Configuring the Exchange Server

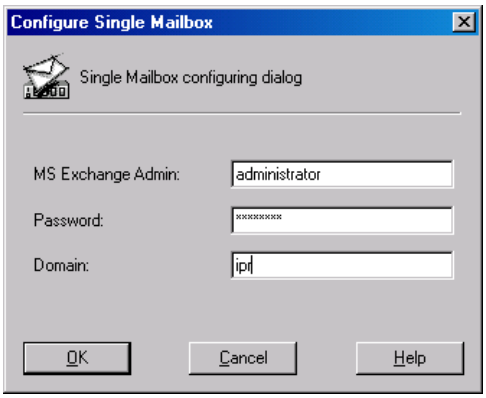

   Click **OK**.
6. The Exchange Server is configured. Exit the GUI or proceed with creating the backup specification at Step 6.

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 39 Microsoft Exchange Single Mailbox integration 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 system that has been struck by a disaster. For information, see "Backup" (page 110).

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 73 Selecting a template

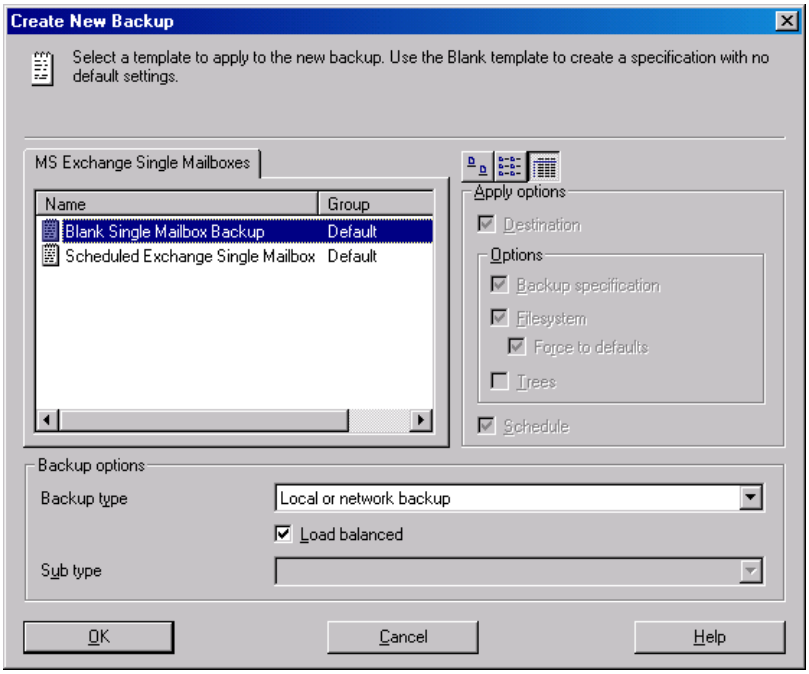

4. In **Client**, select the Exchange Server system. In a cluster environment, select the virtual server.

   For information on the **User and group/domain** options, press **F1**.

   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" (page 160).
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.

NOTE: If some mailboxes have the same display name (for example, `user`), Data Protector appends a user unique string at the end of each mailbox name (for example, `user@@user1`, `user@@user2`, and so on).

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 74 Selecting Exchange Server items for backup

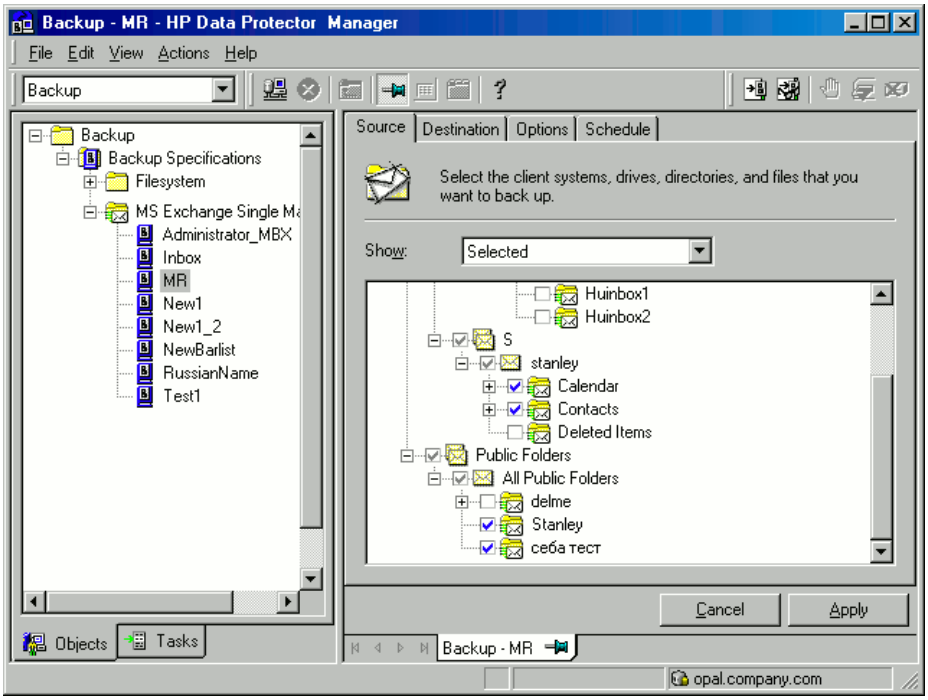

Click **Next**.

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 ("Microsoft Exchange Single Mailbox integration-specific backup options" (page 164)), see "Microsoft Exchange Single Mailbox integration-specific backup options" (page 164).

Figure 75 Microsoft Exchange Single Mailbox integration-specific backup options

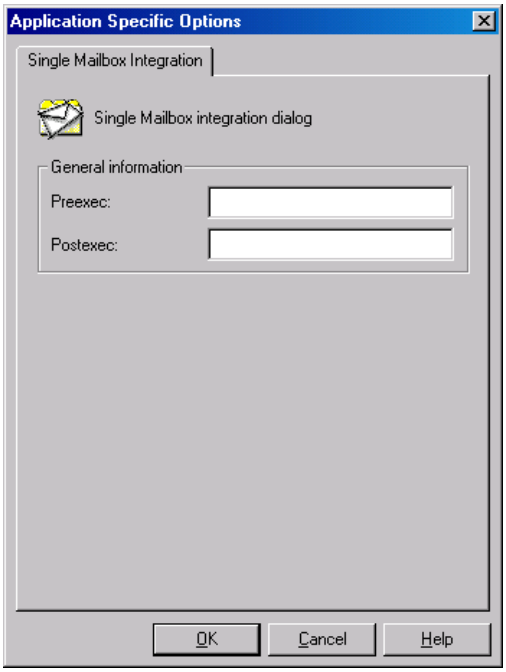

Click **Next**.

9. Optionally, schedule the backup. See "Scheduling backup sessions" (page 164).
   Click **Next**.
10. Save the backup specification, specifying a name and a backup specification group.

TIP: Preview backup session for your backup specification before using it. See "Previewing backup sessions" (page 165).

Table 40 Microsoft Exchange Single Mailbox integration-specific backup options

Option	Description
Pre-exec, **Post-exec**	Specify a command to be executed by `mbx_bar.exe` on the Exchange Server system before the backup (`pre-exec`) or after it (`post-exec`). Do not use double quotes. Type only the name of the command and ensure that the command resides in the *`Data_Protector_home`*`\bin` 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 sessions

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

Scheduling example

To 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 "Scheduling a backup session" (page 165).

   Click **OK**.
3. Repeat Step 1 and Step 2 to schedule backups at 18:00 and 20:00.
4. Click **Apply** to save the changes.

Figure 76 Scheduling a backup session

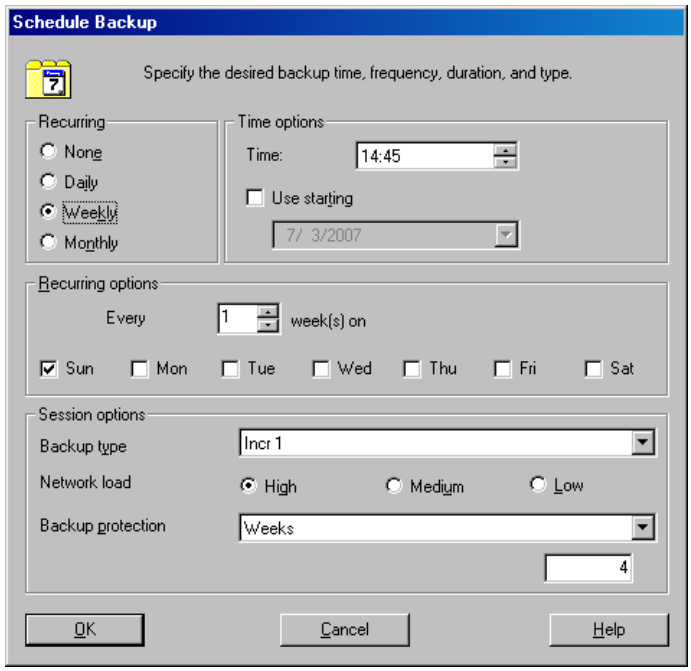

Previewing backup sessions

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

Execute the following command:

```
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 use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

On the Exchange Server system, execute:

`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, execute:

```
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 2007 Integration` component installed and that the Exchange Server is configured for use with Data Protector.

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.

To restore the data from the root mailbox folder, select **Top of Information Store** under the appropriate user mailbox.

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

See "Selecting Exchange Server items for restore" (page 167).

Figure 77 Selecting Exchange Server items for restore

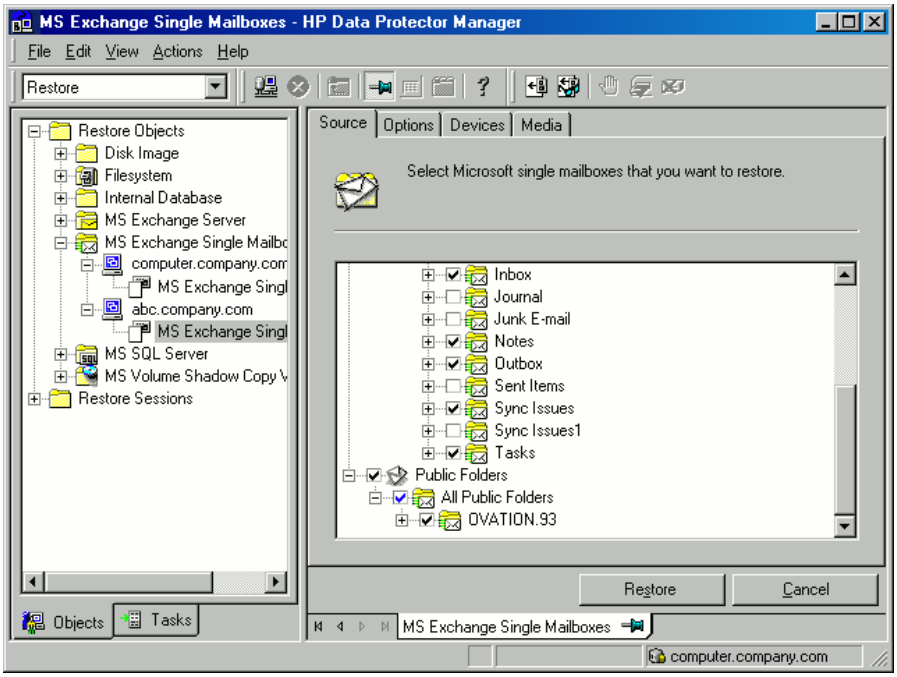

You can specify the backup version, the chain of backups to be used, and the restore destination for each mailbox or Public Folders separately.

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 "Version properties" (page 167).

Figure 78 Version properties

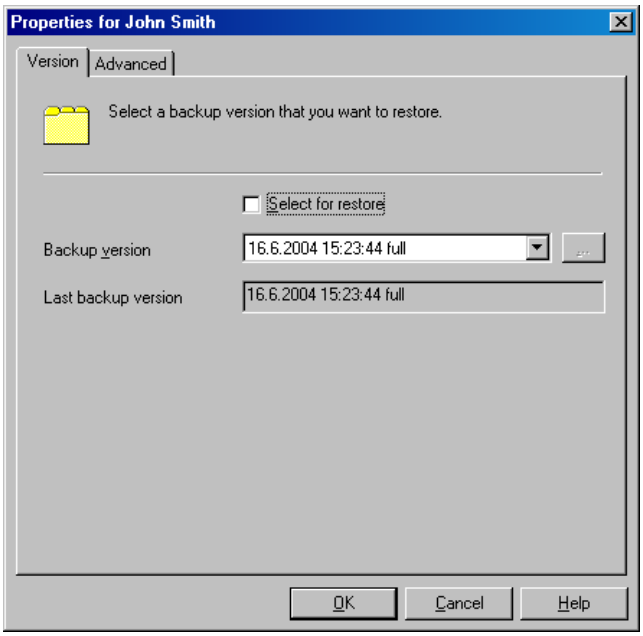

To specify the restore destination and the chain of backup sessions to be used, click the **Advanced** tab. See "Advanced properties" (page 168).

Figure 79 Advanced properties

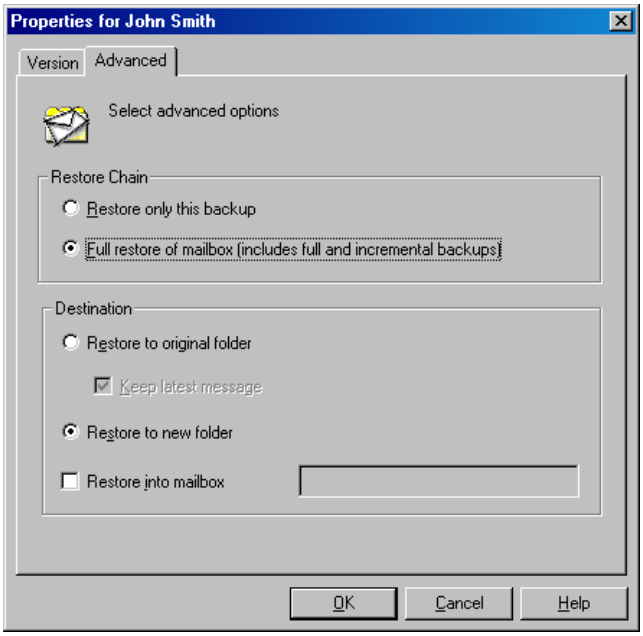

For details on these options, see "Microsoft Exchange Single Mailbox integration restore options" (page 170).

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 80 Selecting the destination Exchange Server system

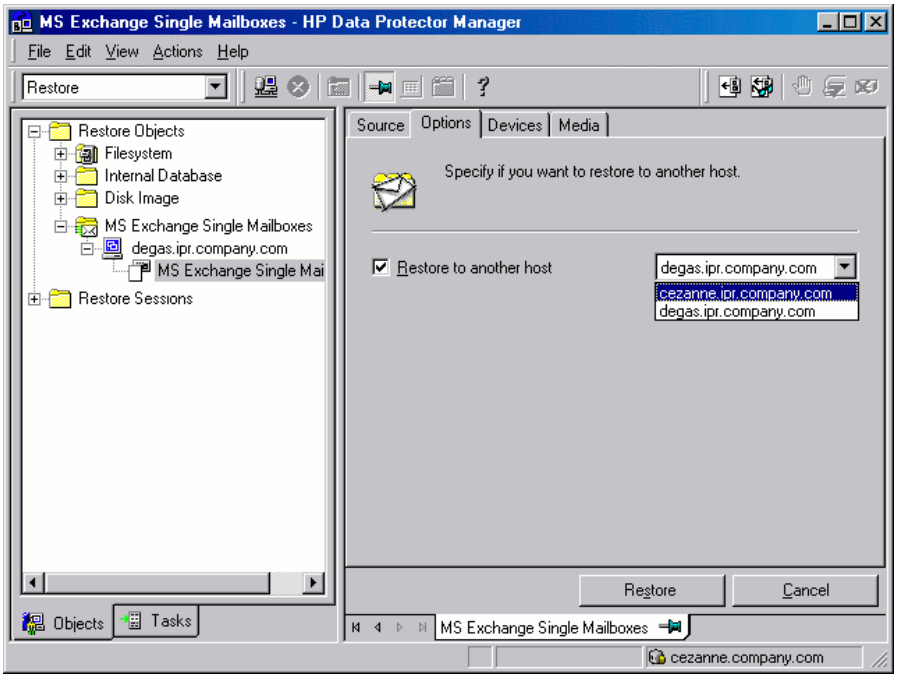

5. In the **Devices** page, select the devices to be used for the restore.

   For more information of how to select devices for a restore, see the *HP Data Protector Help* index: “restore, selecting devices for”.
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. Specify **Report level** and **Network load**.

   Click **Finish** to start the restore.

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

Figure 81 Restored mailbox and public folders content with the restore to new folder option selected.

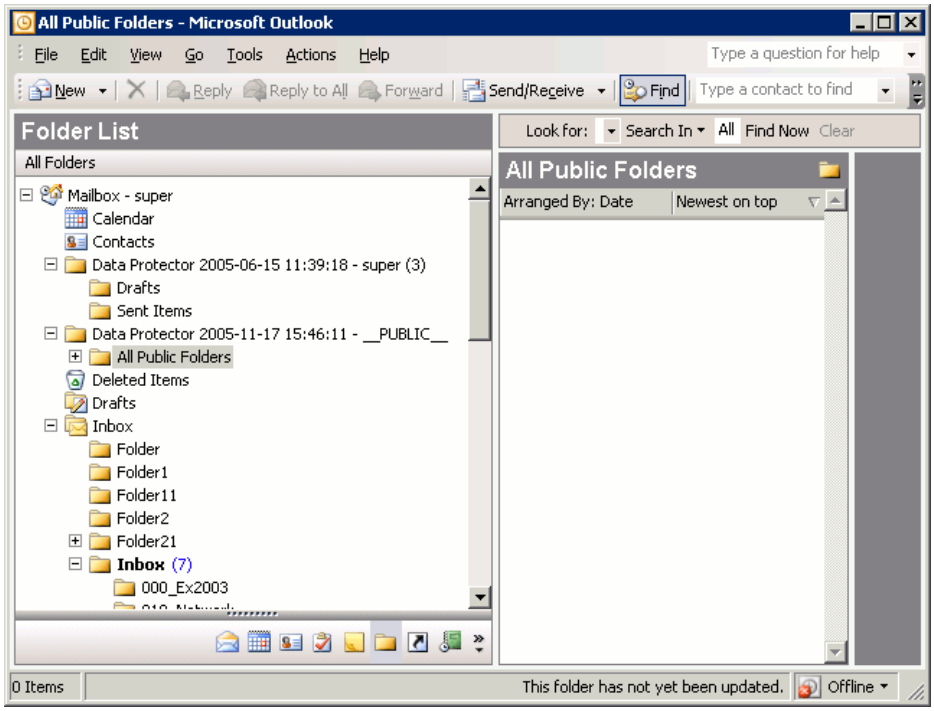

To transfer restored data to `.pst` files:

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 41 Microsoft Exchange Single Mailbox integration restore options

Option	Description
Restore only this backup	Select this option to restore data only from the selected backup session.
Full restore of mailbox	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. Note that any Exchange item that was backed up in any of these sessions is displayed and can be selected for restore.
Restore to original folder	Data Protector restores Exchange items to the same location from which they were backed up. If `Keep latest message` is selected, existing messages in the destination mailbox or Public Folders are not restored even if they differ from their backed up versions. If `Keep latest message` 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. 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. By default, this option is not selected.
Restore to new folder	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 "Restored mailbox and public folders content with the restore to new folder option selected." (page 170).

Table 41 Microsoft Exchange Single Mailbox integration restore options *(continued)*

Option	Description
	When restoring a mailbox, the folder is named `Data Protector backup_date backup_time`. When restoring Public Folders it is named `Data Protector backup_date backup_time - public folder`. 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 `Data Protector backup_date backup_time (1)` is created.
Restore into mailbox	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. For privacy protection, you cannot restore Exchange items from mailboxes to Public Folders.
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, execute:

```
omnir -mbx
-barhost ClientName
[-destination DestClientName]
-mailbox MailboxName -session BackupID [MAILBOX_OPTIONS]
-public -session BackupID [PUBLIC_FOLDERS_OPTIONS]
[GENERAL_OPTIONS]

MAILBOX_OPTIONS
-destmailbox DestMailboxName
-folder Folder
-exclude ExFolder
-originalfolder {-keep_msg | -overwrite_msg}
-chain

PUBLIC_FOLDERS_OPTIONS
-folder Folder
-exclude ExFolder
-originalfolder {-keep_msg | -overwrite_msg}
-chain
```

NOTE: To restore multiple mailboxes, repeat the options –`mailbox MailboxName -session BackupID [MAILBOX_OPTIONS]`.

To restore or exclude from restore multiple folders, repeat the options `-folder Folder` and `-exclude ExFolder`.

Parameter description

ClientName	Original Exchange Server system, from which Exchange items to be restored were backed up.
DestClientName	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).
BackupID	A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session. Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose

	the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the object *copy* session. The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.
MailboxName	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"`.
DestMailboxName	Destination mailbox, to which the Exchange items from the mailbox will be restored (needed only if you are not restoring to the original mailbox).
Folder	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"`.
ExFolder	Subfolder to be excluded from restore of the mailbox or Public Folders.

Option description

`-originalfolder`	This option is equivalent to the Data Protector GUI option `Restore to original folder`. If not specified, the same results occur as if the Data Protector GUI option `Restore to new folder` was selected.
`-chain`	This option is equivalent to the Data Protector GUI option `Full restore of mailbox`. If not specified, the same results occur as if the Data Protector GUI option `Restore only this backup` 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 `2011/01/10-1` from the Exchange Server system `infinity.ipr.company.com`, to a new folder in the mailbox `TEMP` on the same Exchange Server system, execute:

```
omnir -mbx -barhost infinity.ipr.company.com -mailbox FIRST -session
2011/01/10-1 -destmailbox TEMP
```

Example 2

To restore the folder `Inbox` from the mailbox `User 1`, backed up in the session `2010/03/10-18` from the Exchange Server system `exchange.hp.com`, to the original folder without overwriting the messages in the original folder, execute:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 1" -session
2010/03/10-18 -folder Inbox -originalfolder -keep_msg
```

Example 3

To restore the mailbox `User 2`, backed up in the session `2010/03/10-19` 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`, execute:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 2" -session
2010/03/10-19 -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 `2010/06/10-19` from the Exchange Server system `exchange.hp.com`, to a new folder in Public Folders and to the original folders in the mailbox respectively, execute:

```
omnir -mbx -barhost exchange.hp.com -public -session 2010/06/10-19
-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 *HP Data Protector 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 82 Example of backup with concurrency 1

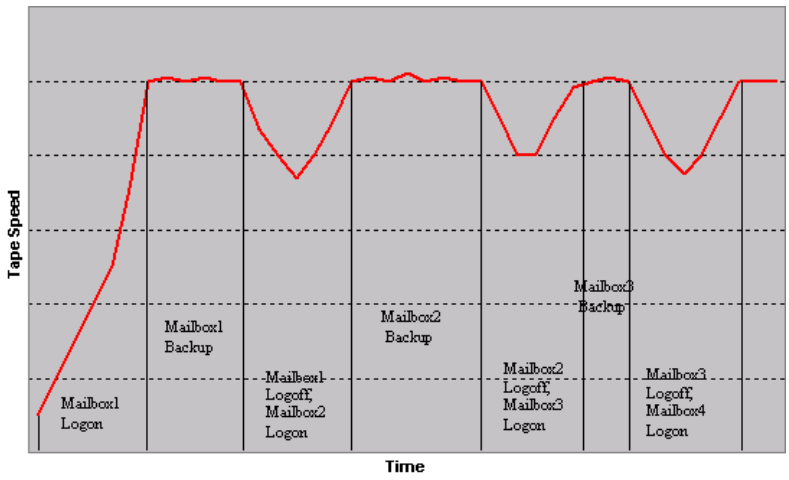

Tests have shown that best performance is achieved when backing up mailboxes and Public Folders using two concurrent data streams, either by specifying one device with concurrency 2 or two devices with concurrency 1.

Figure 83 Example of backup with concurrency 2

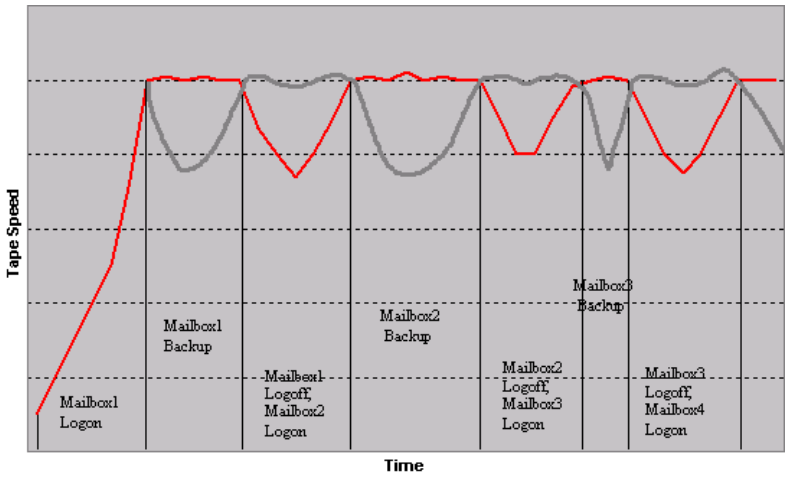

NOTE: Data Protector cannot create more than one backup object out of Exchange items from a single mailbox or Public Folders.

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" (page 175). 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. On how to verify this, see the *HP Data Protector Help* index: "patches".
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see http://support.openview.hp.com/selfsolve/manuals.

Checks and verifications

If your configuration, backup, or restore failed:

- Ensure that the following directories exist on the Data Protector Cell Manager:

  Data_Protector_program_data\config\server\barlists\Mailbox

  Data_Protector_program_data\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 CellServer 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" (page 160).
  - 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 the *HP Data Protector 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. Execute:

      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 "Specifying the Exchange Server administrator's mailbox" (page 177).

Figure 84 Specifying the Exchange Server administrator's mailbox

Problem

Restore to another client fails

Action

Ensure that Exchange Server and the Data Protector **MS Exchange 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 system.

Glossary

A

access rights *See* user rights.

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

Active Directory *(Windows specific term)* The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AES 256-bit encryption Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

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

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

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

application system *(ZDB specific term)* A system the application or database runs on. The application or database data is located on source volumes.
See also backup system and source volume.

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

archived log files *(Data Protector specific term)* Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online and offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.

archived redo log *(Oracle specific term)* Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also online redo log.

ASR set A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory *Data_Protector_program_data*`\Config\Server\dr\asr` (Windows systems) or `/etc/opt/omni/server/dr/asr` (UNIX systems) as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs Data files to which auditing information is stored.

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

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

autochanger *See* library.

autoloader *See* library.

Automatic Storage Management (ASM) *(Oracle specific term)* A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

auxiliary disk A bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

B

BACKINT *(SAP R/3 specific term)* SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain *See* restore chain.

backup device A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

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

backup ID An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image.

A backup object is defined by:

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

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

backup session A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type. The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also backup specification, full backup, and incremental backup.

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

backup set *(Oracle specific term)* A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)* A system connected to a disk array together with one or multiple application systems. The backup system is typically connected to a disk array to create target volumes (a replica) and is used for mounting the target volumes (the replica).
See also application system, target volume, and replica.

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

backup view Data Protector provides different views for backup specifications:

By Type - according to the type of data available for backups/templates. Default view.

By Group - according to the group to which backup specifications/templates belong.

By Name - according to the name of backup specifications/templates.

By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

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

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

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splitable 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 Data Protector Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also BRBACKUP and BRRESTORE.

BRBACKUP *(SAP R/3 specific term)* An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also BRARCHIVE and BRRESTORE.

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

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

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also BRBACKUP and BRARCHIVE.

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

C

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

Catalog Database (CDB)
A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database.
See also MMDB.

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

CDB
See Catalog Database (CDB).

CDF file
(UNIX systems specific term) A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN or SAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also MoM.

Centralized Media Management Database (CMMDB)
See CMMDB.

Certificate Server
A Windows Certificate Server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal
(Windows specific term) A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

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

channel
(Oracle specific term) An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type 'disk'
- type 'sbt_tape'

If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging
(Microsoft Exchange Server and Lotus Domino Server specific term) Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup
A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification:

- If you select the check box next to the client system name, a single backup object of the `Client System` type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, `CONFIGURATION` is also backed up.
- If you individually select all volumes that are mounted on the client system, a separate backup object of the `Filesystem` type is created for each volume. As a result, at the time of the

backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the backup specification was created are not backed up.

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

cluster continuous replication *(Microsoft Exchange Server specific term)* Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also Exchange Replication Service and local continuous replication.

cluster-aware application It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

CMD script for Informix Server *(Informix Server specific term)* A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Class Registration Database *(Windows specific term)* The COM+ Class Registration Database and the Windows Registry store application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command device *(HP P9000 XP Disk Array Family specific term)* A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

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

concurrency *See* Disk Agent concurrency.

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

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

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

CRS The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

CSM The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D

data file *(Oracle and SAP R/3 specific term)* A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also catalog protection.

data replication (DR) group *(HP P6000 EVA Disk Array Family specific term)* A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log.
See also copy set.

data stream Sequence of data transferred over the communication channel.

Data_Protector_home A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is `%ProgramFiles%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_program_data.

Data_Protector_program_data A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012. Its default path is `%ProgramData%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_home.

database library A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

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

database server A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

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

DC directory A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database.
See also Detail Catalog Binary Files (DBCF) and Internal Database (IDB).

DCBF *See* Detail Catalog Binary Files (DCBF).

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.

Detail Catalog Binary Files (DCBF) A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files.
See also DC directory and Internal Database (IDB).

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

device chain A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)* A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written

to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

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

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

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

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

directory junction *(Windows specific term)* Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

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

disaster recovery operating system *See* DR OS.

Disk Agent A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

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

disk group *(Veritas Volume Manager specific term)* The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image backup A high-speed backup where Data Protector backs up files as bitmap images. A disk image backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup.
See also virtual full backup.

Distributed File System (DFS) A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

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

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

DR OS An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.

drive A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

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

E

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

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

encrypted control communication Data Protector secure communication between the clients in the Data Protector cell is based on Secure Socket Layer (SSL) that uses SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.

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

enhanced incremental backup Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

enterprise backup environment Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also MoM.

Event Log (Data Protector Event Log) A central repository of all Data Protector-related notifications. By default, all notifications are sent to the Event Log. The events are logged on the Cell Manager into the file *`Data_Protector_program_data`*`\log\server\Ob2EventLog.txt` (Windows systems),

or `/var/opt/omni/server/log/Ob2EventLog.txt` (UNIX systems). The Event Log is accessible only to users of the Data Protector `Admin` user group and to users who are granted the Data Protector `Reporting and notifications` user rights. You can view or delete all events in the Event Log.

Event Logs *(Windows specific term)* Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology.
See also cluster continuous replication and local continuous replication.

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

exporting media A process that removes all data about backup sessions, such as systems, objects, and filenames, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also importing media.

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

F

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

failover *(HP P6000 EVA Disk Array Family specific term)* An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations.
See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

FC bridge *See* Fibre Channel bridge.

Fibre Channel An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

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

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

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

File Replication Service (FRS) A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

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

file version The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror *(HP P9000 XP Disk Array Family specific term)* A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level

mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used.
See also primary volume and mirror unit (MU) number.

flash recovery area *(Oracle specific term)* A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files).
See also recovery files.

formatting A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

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

full database backup A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

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

full ZDB A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.
See also incremental ZDB.

G

global options A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager

group *(Microsoft Cluster Server specific term)* A collection of resources (for example disk volumes, application services, IP names, and addresses) that are needed to run a specific cluster-aware applications.

GUI A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.

H

hard recovery *(Microsoft Exchange Server specific term)* A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM) A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

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

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

HP Business Copy (BC) P6000 EVA *(HP P6000 EVA Disk Array Family specific term)* A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware.

See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

HP Business Copy (BC) P9000 XP
(HP P9000 XP Disk Array Family specific term) An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system.
See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit, application system, and backup system.

HP Command View (CV) EVA
(HP P6000 EVA Disk Array Family specific term) The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, and mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.
See also HP P6000 / HP 3PAR SMI-S Agent and HP SMI-S P6000 EVA Array provider.

HP Continuous Access (CA) P9000 XP
(HP P9000 XP Disk Array Family specific term) An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs).
See also HP Business Copy (BC) P9000 XP, Main Control Unit, and LDEV.

HP Continuous Access + Business Copy (CA+BC) P6000 EVA
(HP P6000 EVA Disk Array Family specific term) An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array.
See also HP Business Copy (BC) P6000 EVA, replica, and source volume.

HP P6000 / HP 3PAR SMI-S Agent
A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the HP P6000 / HP 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface.
See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP P9000 XP Agent
A Data Protector component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It uses RAID Manager Library for communication with a P9000 XP Array storage system.
See also RAID Manager Library.

HP SMI-S P6000 EVA Array provider
An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses.
See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

I

ICDA
(EMC Symmetrix specific term) EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB
See Internal Database (IDB).

IDB recovery file
A file that maintains information about completed IDB backup sessions and the backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is `obdrindex.dat`.

importing media A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also exporting media.

incremental (re)-establish *(EMC Symmetrix specific term)* A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental backup A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length.
See also backup types.

incremental backup *(Microsoft Exchange Server specific term)* A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also backup types.

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

incremental restore *(EMC Symmetrix specific term)* A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape.
See also full ZDB.

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

Inet A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also Key Management Service and Site Replication Service.

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

initializing *See* formatting.

Installation Server A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery *(ZDB specific term)* A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

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

Internal Database (IDB) An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It stores its data in an embedded database and a collection of proprietary data files which reside on the Cell Manager.
See also DC directory and Detail Catalog Binary Files (DBCF).

Internet Information Services (IIS) *(Windows specific term)* Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

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

J

jukebox *See* library.

jukebox device A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the "file jukebox device".

K

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

keychain A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

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

KMS Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

L

LBO *(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

LDEV *(HP P9000 XP Disk Array Family specific term)* A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array.
See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or unattended operation A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

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

load balancing By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used

for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication

(Microsoft Exchange Server specific term) Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.

An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal.

A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.

See also cluster continuous replication and Exchange Replication Service.

lock name

You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script

(Informix Server UNIX systems specific term) A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server `ALARMPROGRAM` configuration parameter defaults to the *`INFORMIXDIR`*`/etc/log_full.sh`, where *`INFORMIXDIR`* is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the `ALARMPROGRAM` configuration parameter to *`INFORMIXDIR`*`/etc/no_log.sh`.

logging level

An option that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID

(Microsoft SQL Server specific term) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database

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

- *`user_name`* is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- *`password`* must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.
- *`service`* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *`user_name`*`/`*`password`*`@`*`service`*, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the

Oracle target database. In this case, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API *(Lotus Domino Server specific term)* An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

M

Magic Packet *See* Wake ONLAN.

mailbox *(Microsoft Exchange Server specific term)* The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store *(Microsoft Exchange Server specific term)* A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

Main Control Unit (MCU) *(HP P9000 XP Disk Array Family specific term)* An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP or HP CA+BC P9000 XP configuration and acts as a master device.
See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

maintenance mode An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the Data Protector installation.

make_net_recovery `make_net_recovery` is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX `bootsys` command or interactively specified on the boot console.

make_tape_recovery `make_tape_recovery` is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

Manager-of-Managers (MoM) *See* MoM.

MAPI *(Microsoft Exchange Server specific term)* The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU *See* Main Control Unit (MCU).

Media Agent A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

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

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

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

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

media pool A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

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

media usage policy The media usage policy controls how new backups are added to the already used media. It can be `Appendable`, `Non-Appendable`, or `Appendable for incrementals only`.

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

merging This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored.
See also overwrite.

Microsoft Exchange Server A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)* An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

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

Microsoft Volume Shadow Copy Service (VSS) A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also shadow copy, shadow copy provider, replica, and writer.

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

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

mirror unit (MU) number *(HP P9000 XP Disk Array Family specific term)* A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family.
See also first-level mirror.

mirrorclone *(HP P6000 EVA Disk Array Family specific term)* A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also CMMDB and Catalog Database (CDB).

MoM Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

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

mount request A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

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

multisnapping *(HP P6000 EVA Disk Array Family specific term)* Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot.
See also snapshot.

O

OBDR capable device A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

obdrindex.dat *See* IDB recovery file.

object *See* backup object.

object consolidation The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

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

object copy session A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

object ID *(Windows specific term)* The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

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

object mirroring The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

object verification The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

offline backup A backup during which an application database cannot be used by the application. In an offline backup session, the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the time period of the data replication process. For instance, for backup to tape, until streaming of data to the tape is finished. Normal database operation is resumed before potential post-backup operations are started.
See also zero downtime backup (ZDB) and online backup.

offline recovery Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Cell Manager can only be recovered offline.

offline redo log *See* archived redo log.

ON-Bar *(Informix Server specific term)* A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

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

ONCONFIG *(Informix Server specific term)* An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file in the directory *`INFORMIXDIR`*`\etc` (on Windows systems or *`INFORMIXDIR`*`/etc/` (on UNIX systems).

online backup A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period of the data replication process. For instance, for backup to tape, until streaming of data to tape is finished. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. Normal database operation is resumed before potential post-backup operations are started.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.

See also zero downtime backup (ZDB) and offline backup.

online recovery A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.

online redo log *(Oracle specific term)* Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also archived redo log.

Oracle Data Guard *(Oracle specific term)* Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

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

ORACLE_SID *(Oracle specific term)* A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *`ORACLE_SID`*. The *`ORACLE_SID`* is included in the CONNECT DATA

parts of the connect descriptor in a `TNSNAMES.ORA` file and in the definition of the TNS listener in the `LISTENER.ORA` file.

original system The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also merging.

ownership Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.

If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

P1S file P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into the directory *`Data_Protector_program_data`*`\Config\Server\dr\p1s` (Windows systems), or `/etc/opt/omni/server/dr/p1s` (UNIX systems) with the filename `recovery.p1s`.

package *(MC/ServiceGuard and Veritas Cluster specific term)* A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP P9000 XP Disk Array Family specific term)* The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP P9000 XP Agent:

- PAIR – The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.
- SUSPENDED – The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.
- COPY – The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

parallel restore Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism The concept of reading multiple data streams from an online database.

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

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

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

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

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

post-exec A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems.
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 systems and as shell scripts on UNIX systems.

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

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

primary volume (P-VOL) *(HP P9000 XP Disk Array Family specific term)* An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU).
See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection *See* data protection and also catalog protection.

public folder store *(Microsoft Exchange Server specific term)* The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

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

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

R

RAID Redundant Array of Independent Disks.

RAID Manager Library *(HP P9000 XP Disk Array Family specific term)* A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands.
See also HP P9000 XP Agent.

RAID Manager P9000 XP *(HP P9000 XP Disk Array Family specific term)* A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.

rawdisk backup *See* disk image backup.

RCU *See* Remote Control Unit (RCU).

RDBMS Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)* A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

Recovery Catalog *(Oracle specific term)* A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

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

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

recovery files *(Oracle specific term)* Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces.
See also flash recovery area.

Recovery Manager (RMAN) *(Oracle specific term)* An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

RecoveryInfo When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

recycle or unprotect A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)* Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit (RCU) *(HP P9000 XP Disk Array Family specific term)* An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.

Removable Storage Management Database *(Windows specific term)* A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)* A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)* An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on a UNIX system, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition is replicated.

See also snapshot, snapshot creation, split mirror, and split mirror creation.

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 Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.

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

resync mode *(HP P9000 XP Disk Array Family VSS provider specific term)* One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.
See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), mirror unit (MU) number, and replica set rotation.

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

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

RSM *(Windows specific term)* Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

S

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

scanning A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

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

secondary volume (S-VOL) *(HP P9000 XP Disk Array Family specific term)* An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration.
See also primary volume (P-VOL) and Main Control Unit (MCU).

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

session ID An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.

session key This environment variable for the pre-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the `omnimnt`, `omnistat`, and `omniabort` commands.

shadow copy *(Microsoft VSS specific term)* A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original

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

shadow copy provider *(Microsoft VSS specific term)* An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also shadow copy.

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

shared disks A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

Site Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
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.

SMBF The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA) *See* HP P6000 / HP 3PAR SMI-S Agent.

snapshot *(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)* A type of target volumes created using a specific replication technology. Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume.
See also replica and snapshot creation.

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

snapshot creation *(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)* A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation.
See also snapshot.

source (R1) device *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also target (R2) device.

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

sparse file A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix Disk Array and HP P9000 XP Disk Array Family specific term)* A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes.
See also replica and split mirror creation.

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

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

split mirror creation *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also split mirror.

split mirror restore *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)* A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method.
See also ZDB to tape, ZDB to disk+tape, and replica.

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

SRD file *(disaster recovery specific term)* A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows or Linux system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster.
See also target system.

SRDF *(EMC Symmetrix specific term)* The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

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

sst.conf file The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

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

Storage Group *(Microsoft Exchange Server specific term)* A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

storage volume *(ZDB specific term)* An object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, filesystems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

StorageTek ACS library *(StorageTek specific term)* Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

switchover *See* failover.

Sybase Backup Server API *(Sybase specific term)* An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)* The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

SYMA *See* EMC Symmetrix Agent.

synthetic backup A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

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

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

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

System Recovery Data file *See* SRD file.

System State *(Windows specific term)* The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

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

T

tablespace A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

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

target (R2) device *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also source (R1) device.

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

target system *(disaster recovery specific term)* A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.

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

Terminal Services *(Windows specific term)* Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(Microsoft SQL Server specific term)* An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)* A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

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

transaction log backup Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

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

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

transportable snapshot *(Microsoft VSS specific term)* A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed.
See also Microsoft Volume Shadow Copy Service (VSS).

U

unattended operation *See* lights-out operation.

user account (Data Protector user account) You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

User Account Control (UAC) A security component in Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)* Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

user_restrictions file A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than *admin* and *operator*.

V

vaulting media The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

verify A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

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

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

virtual disk *(HP P6000 EVA Disk Array Family specific term)* A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array.
See also source volume and target volume.

virtual full backup An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

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

virtual server A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape *(VLS specific term)* An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs.
See also Virtual Library System (VLS) and virtual tape library (VTL).

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

VMware management client *(VMware (Legacy) integration specific term)* The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).

volser *(ADIC and STK specific term)* A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint *(Windows specific term)* An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

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

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

VSS compliant mode *(HP P9000 XP Disk Array Family VSS provider specific term)* One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching the disks.
See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS Veritas Journal Filesystem.

VxVM (Veritas Volume Manager) A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

W

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

Web reporting The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows configuration backup Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

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

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer *(Microsoft VSS specific term)* A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

X

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

Z

ZDB *See* zero downtime backup (ZDB).

ZDB database *(ZDB specific term)* A part of the IDB, storing ZDB-related information such as source volumes, replicas, and security information. The ZDB database is used in zero downtime backup, instant recovery, and split mirror restore sessions.
See also zero downtime backup (ZDB).

ZDB to disk *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape *(ZDB specific term)* A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB) A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index