

HP OpenView Storage Data Protector Integration Guide

**for Microsoft Applications:
SQL Server
Exchange Server
Volume Shadow Copy Service**

Manual Edition: July 2006



Manufacturing Part Number: B6960-96007

Release A.06.00

© Copyright 2006 Hewlett-Packard Development Company, L.P.

Legal Notices

©Copyright 2006 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

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

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

UNIX® is a registered trademark of The Open Group

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

1. Integrating Microsoft SQL Server and Data Protector

In This Chapter	2
Introduction	3
Prerequisites and Limitations.....	5
Integration Concept	6
Advanced Concept - Parallelism	9
Data Protector Microsoft SQL Server Configuration File	12
Configuring the Integration	14
Configuring Microsoft SQL Server	14
Configuring a Microsoft SQL Server Backup	18
Backing Up Microsoft SQL Server Databases	25
Scheduling a Backup	25
Running an Interactive Backup.....	26
Restoring a Microsoft SQL Server Database	27
Restoring Using the Data Protector GUI	27
Restore Options.....	31
Restoring Using the Data Protector CLI.....	34
Disaster Recovery	35
Performance Tuning	39
Monitoring a Microsoft SQL Server Backup and Restore	44
Monitoring Current Sessions	44
Viewing Previous Sessions.....	45
Troubleshooting	47
Before You Begin	47
Configuration Problems	47
Backup Problems	50
Restore Problems	51
Before You Call Support.....	53

2. Integrating Microsoft Exchange Server and Data Protector

In This Chapter	56
Introduction	57
Prerequisites and Limitations.....	60
Integration Concept	61
Configuring the Integration	64
Configuring a Microsoft Exchange Backup.....	65
Backing Up Microsoft Exchange Server.....	72
Scheduling a Backup	72

Contents

Running an Interactive Backup	73
Restoring a Microsoft Exchange Server Database	75
Restore Using the GUI	76
Restore Using the CLI	83
Troubleshooting	85
Before You Begin	85
Checks and Verifications	85
Backup Problems	86
Restore Problems	86

3. Integrating Microsoft Exchange Single Mailbox and Data Protector

Introduction	90
Integration Concepts	91
Configuring the Integration	93
Prerequisites	93
Limitations	93
Before You Begin	93
Cluster-Aware Clients	94
Configuring Exchange Server Users	94
Configuring Exchange Servers	94
Checking the Configuration	95
Backup	96
Creating Backup Specifications	96
Modifying Backup Specifications	100
Scheduling Backup Specifications	100
Previewing Backup Sessions	101
Starting Backup Sessions	102
Restore	104
Before You Begin	104
Restoring Using the Data Protector GUI	104
Restoring Using the Data Protector CLI	110
Monitoring Sessions	114
Performance Tuning	115
Troubleshooting	117
Before You Begin	117
Checks and Verifications	117
Problems	118

4. Integrating Microsoft Volume Shadow Copy Service with Data Protector	
In This Chapter	124
Introduction	125
Prerequisites and Limitations.	129
Integration Concepts	131
Backup.	132
Restore.	133
Configuring the Integration	136
Configuring the Data Protector VSS Cluster-Aware Integration	136
Writers Specifics	139
Backing Up Writers Data	144
Creating Backup Specifications Using GUI	144
Scheduling the Backup.	151
Running an Interactive Backup.	152
Restoring Writers Data	154
Restore Procedure.	154
MSDE Writer Restore Specifics	158
Microsoft Exchange Server 2003 Writer Restore Specifics	159
Microsoft Data Protection Manager 2006 Writer Restore Specifics	162
Monitoring a VSS Backup and Restore	165
Monitoring Current Sessions	165
Viewing Previous Sessions.	166
Troubleshooting	168
Before You Begin	168
Backup Problems	168
Restore Problems	168

Glossary

Index

Contents

Printing History

Manual 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	Manual Edition	Product
B6960-90108	October 2004	Data Protector Release A.05.50
B6960-96007	July 2006	Data Protector Release A.06.00

Conventions

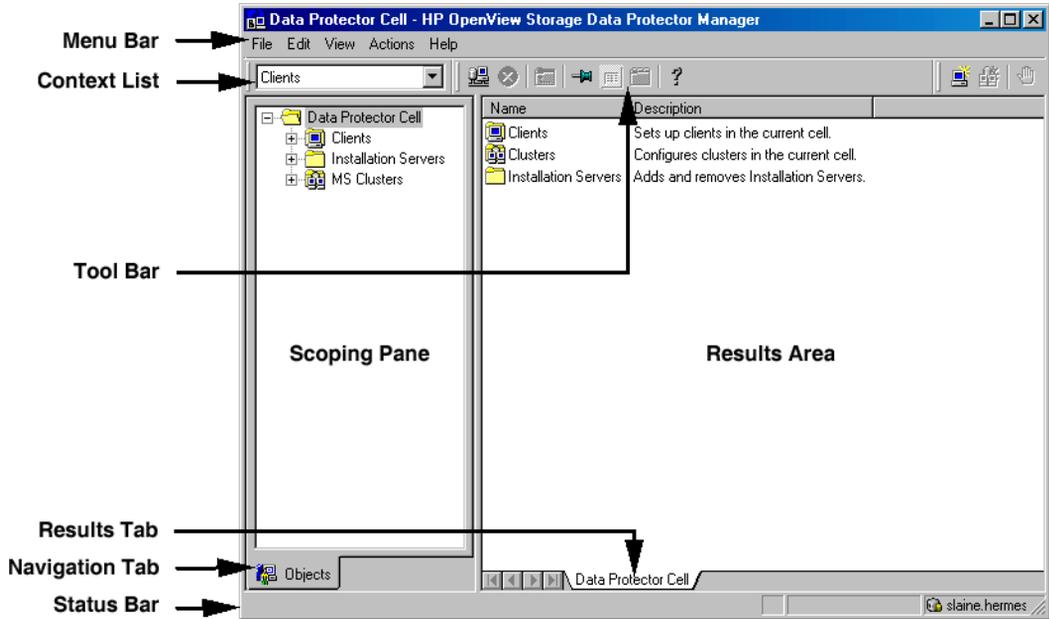
The following typographical conventions are used in this manual.

Table 2

Convention	Meaning	Example
<i>Italic</i>	Book or manual titles, and manual page names	Refer to the <i>HP OpenView Storage Data Protector Integration Guide</i> for more information.
	Provides emphasis	You <i>must</i> follow these steps.
	Specifies a variable that you must supply when entering a command	At the prompt type: rlogin <i>your_name</i> where you supply your login name.
Bold	New terms	The Data Protector Cell Manager is the main ...
Computer	Text and items on the computer screen	The system replies: Press Enter
	Command names	Use the grep command ...
	File and directory names	/usr/bin/X11
	Process names	Check to see if Data Protector Inet is running.
	Window/dialog box names	In the Backup Options dialog box...
	Text that you must enter	At the prompt, type: ls -l
Keycap	Keyboard keys	Press Return .

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. Refer to the online Help for information about the Data Protector graphical user interface.

Figure 1 Data Protector Graphical User Interface



Contact Information

General Information

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

Technical Support

Technical support information can be found at the HP Electronic Support Centers at

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

Information about the latest Data Protector patches can be found at

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

HP does not support third-party hardware and software. Contact the respective vendor for support.

Documentation Feedback

Your comments on the documentation help us to understand and meet your needs. You can provide feedback at

storagedocs.feedback@hp.com

Training Information

For information on currently available HP OpenView training, see the HP OpenView World Wide Web site at

<http://www.openview.hp.com/training/>

Follow the links to obtain information about scheduled classes, training at customer sites, and class registration.

Data Protector Documentation

Data Protector documentation comes in the form of manuals and online Help.

Manuals

Data Protector manuals are available in printed format and in PDF format. Install the PDF files during the Data Protector setup procedure by selecting the `User Interface` component on Windows or the `OB2-DOCS` component on UNIX. Once installed, the manuals reside in the `<Data_Protector_home>\docs` directory on Windows and in the `/opt/omni/doc/C/` directory on UNIX. You can also find the manuals in PDF format at <http://www.hp.com/support/manuals>

HP OpenView Storage Data Protector Concepts Guide

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

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Troubleshooting Guide

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

HP OpenView Storage Data Protector Disaster Recovery Guide

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

HP OpenView Storage Data Protector Integration Guide

This manual describes how to configure and use Data Protector to back up and restore various databases and applications. It is intended for backup administrators or operators. There are four versions of this manual:

- *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server, Exchange Server, and Volume Shadow Copy Service*

This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft SQL Server 7/2000/2005, and Volume Shadow Copy Service.

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

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

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

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

- *HP OpenView Storage Data Protector Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol*

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for UNIX

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on UNIX.

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for Windows

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on Windows.

There are two versions of the manual:

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

This manual 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 OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide

This manual describes how to configure and use the integration of Data Protector with HP StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array, EMC Symmetrix Remote Data Facility and TimeFinder, and HP StorageWorks Disk Array XP. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server 2000/2003, and Microsoft

SQL Server 2000 databases. The manual also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

HP OpenView Storage Data Protector MPE/iX System User Guide

This manual describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.

HP OpenView Storage Data Protector Media Operations User's Guide

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

HP OpenView Storage Data Protector Product Announcements, Software Notes, and References

This manual gives a description of new features of HP OpenView Storage Data Protector A.06.00. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at <http://www.hp.com/support/manuals>

There are also four other *Product Announcements, Software Notes and References*, which serve a similar purpose for the following:

- OVO UNIX integration
- OVO 7.1x/7.2x Windows integration
- OVO 7.5 Windows integration
- Media Operations

Online Help

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

Documentation Map

Abbreviations

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

Abbreviation	Manual
CLI	Command Line Interface Reference Guide
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
Help	Online Help
IG-IBM	Integration Guide—IBM Applications
IG-MS	Integration Guide—Microsoft Applications
IG-O/S	Integration Guide—Oracle, SAP R/3, and SAP DB/MaxDB
IG-OV	Integration Guide—HP OpenView Service Information Portal/OpenView Reporter
IG-OVOU	Integration Guide—HP OpenView Operations, UNIX
IG-OVOW	Integration Guide—HP OpenView Operations 7.1x, 7.2x, Windows
IG-OVOW	Integration Guide—HP OpenView Operations 7.5, Windows
IG-Var	Integration Guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and Licensing Guide
MO GS	Media Operations Getting Started Guide
MO RN	Media Operations Product Announcements, Software Notes, and References
MO UG	Media Operations User Guide
MPE/iX	MPE/iX System User Guide

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

Map

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

	Help	GS	Concepts	Install	Trouble	DR	PA	Integration Guides							ZDB			MO						
								MS	O/S	IBM	Var	OV	OVOU	OVOW	Concpt	Admin	IG	GS	User	PA	MPE/iX	CLI		
Backup	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	
Disaster Recovery	X		X			X																		
Installation/Upgrade	X	X		X			X					X	X	X				X	X			X		
Instant Recovery	X		X												X	X	X							
Licensing	X			X			X												X					
Limitations	X				X		X	X	X	X	X			X			X					X		
New features	X						X															X		
Planning strategy	X		X									X												
Procedures/Tasks	X			X	X	X		X	X	X	X	X	X	X		X	X		X					
Recommendations			X				X							X								X		
Requirements				X			X	X	X	X	X		X					X	X	X				
Restore	X	X	X					X	X	X	X					X	X						X	
Support matrices							X																	
Supported configurations														X										
Troubleshooting	X			X	X			X	X	X	X	X				X	X							

Integrations

Look in these manuals for details of the following integrations:

Integration	Guide
HP OpenView Operations (OVO)	IG-OVOU, IG-OVOW
HP OpenView Reporter (OVR)	IG-OV
HP OpenView Reporter Light	IG-OVOW
HP OpenView Service Information Portal (OVSIP)	IG-OV
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX System	MPE/iX
Microsoft Exchange Servers	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Servers	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG
Sybase	IG-Var
Symmetrix (EMC)	all ZDB
VMware	IG-Var

In This Book

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

Audience

It 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 OpenView Storage Data Protector Concepts Guide*, which is recommended to fully understand the fundamentals and the model of Data Protector.

Organization

The manual is organized as follows:

- Chapter 1** “Integrating Microsoft SQL Server and Data Protector” on page 1.
- Chapter 2** “Integrating Microsoft Exchange Server and Data Protector” on page 55.
- Chapter 3** “Integrating Microsoft Exchange Single Mailbox and Data Protector” on page 89.
- Chapter 4** “Integrating Microsoft Volume Shadow Copy Service with Data Protector” on page 123.
- Glossary** Definition of terms used in this manual.

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft SQL Server integration.

The chapter is organized into the following sections:

“Introduction” on page 3

“Prerequisites and Limitations” on page 5

“Integration Concept” on page 6

“Data Protector Microsoft SQL Server Configuration File” on page 12

“Configuring the Integration” on page 14

“Backing Up Microsoft SQL Server Databases” on page 25

“Restoring a Microsoft SQL Server Database” on page 27

“Performance Tuning” on page 39

“Monitoring a Microsoft SQL Server Backup and Restore” on page 44

“Troubleshooting” on page 47

Introduction

The Data Protector integration with Microsoft SQL Server allows you to perform online backups.

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for information about platforms and devices that are supported by the Data Protector Microsoft SQL Server integration.

The online backup concept is widely accepted. It addresses the business requirements for high application availability better than the offline backup concept.

Data Protector Microsoft SQL integration supports the following backup types:

- Full online database backups
- Transaction log online backups
- Differential online database backups

Data Protector Microsoft SQL integration supports the following restore options:

- Point-in-time restore
- Restore database to another SQL Server
- Recovery completion state
- Force restore over existing database

Advantages

Using Data Protector together with Microsoft SQL Server offers several advantages over using Microsoft SQL Server alone:

- Central Management for all backup operations

The administrator can manage backup operations from a central point.

- Media Management

Data Protector has an advanced media management system that allows users to monitor media usage, set the protection for stored data, as well as organize and manage devices in media pools.

Introduction

- Backup Management

Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.

- Scheduling

Data Protector has a built-in scheduler, which allows the administrator to automate backups to run periodically. With the Data Protector Scheduler, the backups you configure run unattended at specified times, provided the devices and media are properly set.

- Device Support

Data Protector supports a wide range of devices: files, standalone drives, very large multiple drive libraries, etc.

- Reporting

Data Protector has reporting capabilities that allow you to receive information about your backup environment. You can schedule reports to be issued at a specific time or attached to a predefined set of events, such as the end of a backup session or a mount request.

- Monitoring

Data Protector has a feature that allows the administrator to monitor currently running sessions and view finished sessions from any system that has the Data Protector GUI installed.

All backup sessions are logged in the IDB, which provides the administrator with a history of activities that can be queried at a later time.

Prerequisites and Limitations

Prerequisites

- You need a license to use the Data Protector Microsoft SQL Server integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Before you begin, make sure that you have correctly installed and configured the Microsoft SQL Server and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector Microsoft SQL Server integration.
 - *SQL Server Books Online* for online information on Microsoft SQL Server.
- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

Limitations

Preview is not possible for SQL backup and restore sessions.

It is assumed that you are familiar with the Microsoft SQL Server database administration and the basic Data Protector functionality.

Integration Concept

Virtual Device Interface

The Microsoft SQL Server introduces a new backup interface called Virtual Device Interface (VDI). VDI allows much faster backups and restores than the backup interface used in previous versions of the Microsoft SQL Server.

The central component of the integration is the Data Protector `sql_bar.exe` executable, which is installed on the Microsoft SQL Server system. From the perspective of the Microsoft SQL Server, Data Protector is seen as media management software. The `sql_bar.exe` executable implements multiple virtual devices used for backup and restore, and transforms VDI commands from the Microsoft SQL Server into Data Protector backup or restore streams.

Fast Direct Mode

The VDI architecture allows the Data Protector General Media Agent to access data directly in the Microsoft SQL Server's memory, provided that the devices are attached directly to the Microsoft SQL Server system. Therefore, high backup and restore speeds on large databases can be achieved.

The high performance Data Protector mode is called **Fast direct mode**.

Backup Types

There are three online backup types of the Microsoft SQL Server system that can be performed using the Data Protector Microsoft SQL integration:

Full database backup

Full database backup includes all data in a database regardless of whether the database has changed after the last backup was created. This means that the entire database backup does not depend on any other backup media.

Differential database backup

A differential database backup records only the data changes made to the database since the last full database backup. A differential database backup takes less time to complete than a full database backup. By creating differential database backups more frequently than full database backups, you can conserve the media used for a 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.

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

Backup Objects

When selecting objects for a backup, you can choose to back up the whole server system or may select particular databases, which are listed below:

Database	Description
user databases	Contain user data. For example, there is the pubs database providing learning tools, which is the basis for most of the examples in Microsoft SQL Server manuals.
master	Controls the user databases and the operation of Microsoft SQL Server as a whole. It keeps track of such information as user accounts, configurable environment variables, and system error messages.
model	Provides a template or prototype for new user databases.
distribution	The distribution database is one of the system databases used by the replication components of the Microsoft SQL Server, such as Distribution Agent, to store data, including transactions, snapshot jobs, synchronization status, and replication history information. This database does not necessarily reside on a server unless the system is used for remote distribution or as a combined Publisher/Distributor.
msdb	Provides a storage area for scheduling information and information about backups.

IMPORTANT

Table backup is not supported by the Data Protector Microsoft SQL Server integration.

Integration Concept

For a complete description of system databases, refer to the *Microsoft SQL Server Books Online*.

Full and differential backups, together with regular transaction log backups, prevent a user from data loss in the event of a disk failure. Furthermore, transaction log backups are needed to perform a point-in-time restore. Data Protector restores the 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 point-in-time restore option.

The recovery itself is performed by the Microsoft SQL Server.

Backup Flow

A backup session is started by the Data Protector Backup Session Manager (BSM), which reads the Data Protector backup specification and invokes `sql_bar.exe`. The BSM also starts the General Media Agents. The `sql_bar.exe` executable connects to the Microsoft SQL Server and receives data from it via VDI. It then passes the instructions on to Data Protector General Media Agents, which write the data to the backup devices.

Messages from the backup session are sent to the Backup Session Manager, which then writes messages and information regarding the respective session to the IDB.

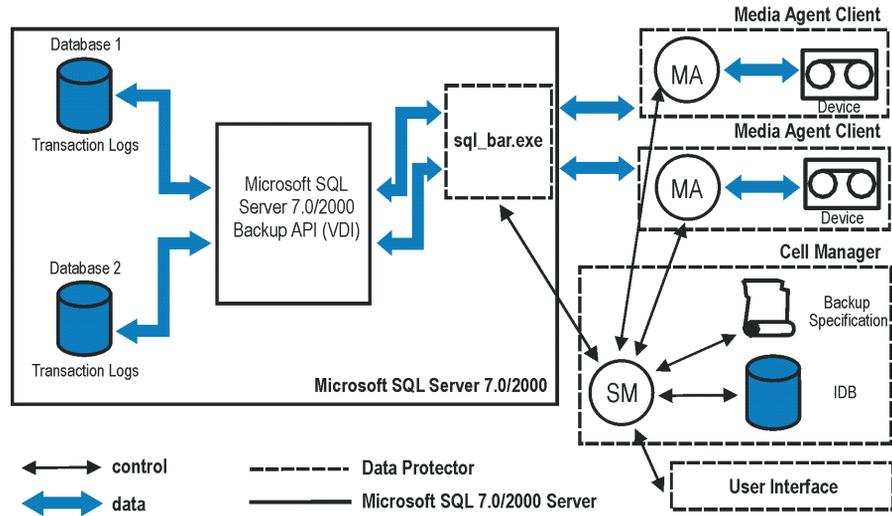
Restore Flow

Using the Data Protector User Interface, the objects and object versions which are to be restored are defined by the user. A restore session is started by the Restore Session Manager (RSM), which starts `sql_bar.exe` and the Data Protector General Media Agents. The `sql_bar.exe` connects to Microsoft SQL Server and receives data from the General Media Agents. The Microsoft SQL Server then writes the data restored by Data Protector to the disks.

Messages from the restore session are sent to the Data Protector RSM, which writes messages and information regarding the respective session to the IDB.

The concept of a backup and a restore session is shown in the Figure 1-1 on page 9.

Figure 1-1 Microsoft SQL Server Integration Concept



Legend:

- SM A Data Protector Session Manager, which is the Data Protector Backup Session Manager during backup, or the Data Protector Restore Session Manager during restore.
- BackupAPI or VDI The Microsoft SQL Server Virtual Device Interface, the Microsoft backup interface introduced with the Microsoft SQL Server 7.0.
- MA The Data Protector General Media Agent.

Advanced Concept - Parallelism

Data Protector can back up more than one Microsoft SQL Server database at a time. It can even back up a single database using multiple streams.

The two parallelism types that are used with Microsoft SQL Server are **database parallelism** and the number of **concurrent** streams.

Integration Concept

- 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.

The allocation of streams to devices is done automatically.

Data Protector tries to use all available devices to run the backup in parallel.

- Number of concurrent streams

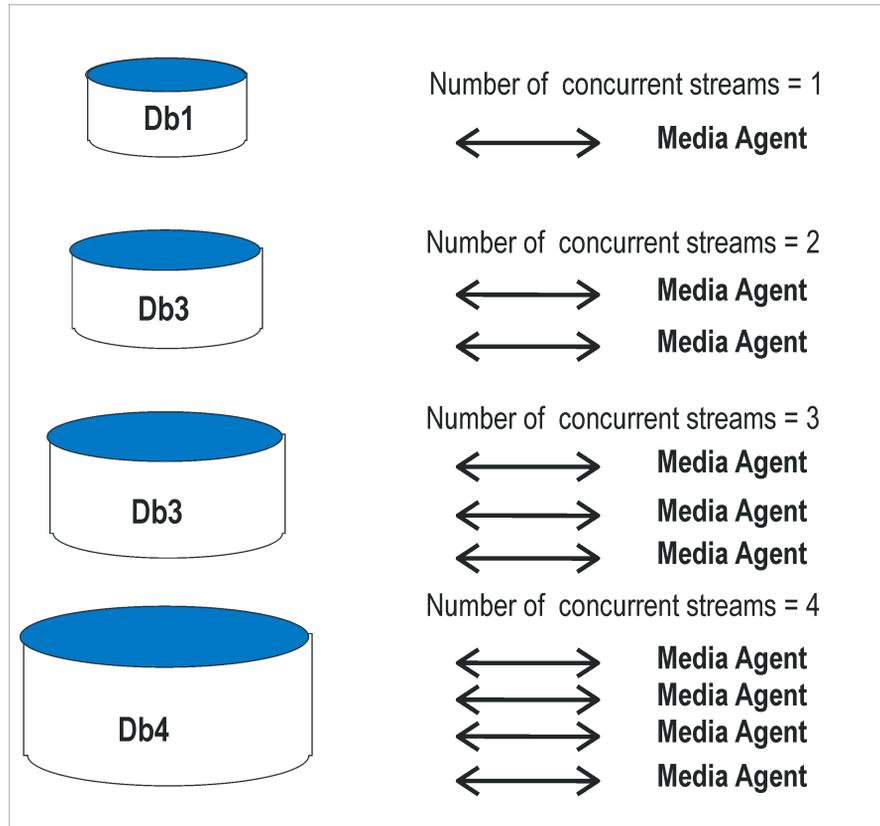
The number of concurrent streams is defined as the number of devices used to back up a particular database or Microsoft SQL Server. It can be specified by the user or calculated automatically.

NOTE

Microsoft SQL Server does not support the backup of multiple streams to one device.

Figure 1-2 on page 11 gives an example of a Data Protector backup session where four Microsoft SQL Server databases are being backed up using a different number of concurrent streams per database.

Figure 1-2 Database Parallelism = 4, Overall Concurrency = 10



Data Protector Microsoft SQL Server Configuration File

Data Protector stores the Microsoft SQL Server integration parameters for every configured Microsoft SQL Server in the `/etc/opt/omni/server/integ/config/MSSQL/<client_name>%<instance_name>` file (HP-UX and Solaris systems), or in the `<Data_Protector_home>\Config\Server\Integ\Config\MSSQL\<client_name>%<instance_name>` file (Windows systems) on the Cell Manager. The parameters stored are the user name and password for the Microsoft SQL Server user, who must have permissions to run backup and restore within Microsoft SQL Server (assuming the standard security is used during the configuration of the integration).

The configuration parameters are written to the Data Protector Microsoft SQL Server configuration file:

- during the configuration of the integration
- during the creation of a backup specification

Syntax

The syntax of the Data Protector Microsoft SQL Server configuration file is as follows:

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

IMPORTANT

To avoid problems with your backups, take extra care to ensure that the syntax of your configuration file matches the examples.

Example

This is an example of the file:

- if standard security is used:

```
Login='TROLL\Administrator';  
Password='dsjf08m80fh43kdf';
```

- if integrated security is used:
Login='';
Password='dsjf08m80fh43kdf';

Configuring the Integration

The configuration of the Data Protector Microsoft SQL Server integration consists of the following:

1. “Configuring Microsoft SQL Server” on page 14
2. “Configuring a Microsoft SQL Server Backup” on page 18

Before You Begin

It is recommended that you configure and run some test filesystem backups using Data Protector.

This includes installing the Disk Agent on the Microsoft SQL Server system. Any device can be used for this test. Configure a standard filesystem backup, which can include one directory only.

Thus, you can check whether the Microsoft SQL Server client system and the Data Protector Cell Manager are communicating properly.

In case of problems, this type of backup is much easier to troubleshoot than the integration of Microsoft SQL Server with Data Protector.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions. In case of any difficulties with the filesystem backup, do not continue configuring the integration until you have solved the problems.

Configuring Microsoft SQL Server

The configuration is performed during the creation of the first backup specification for Microsoft SQL Server databases. For creating a backup specification, see “Configuring a Microsoft SQL Server Backup” on page 18.

However, you can change the configuration any time after you have created at least one backup specification. For information on changing the configuration, see “Changing and Checking the Microsoft SQL Server Configuration” on page 17.

Prerequisites

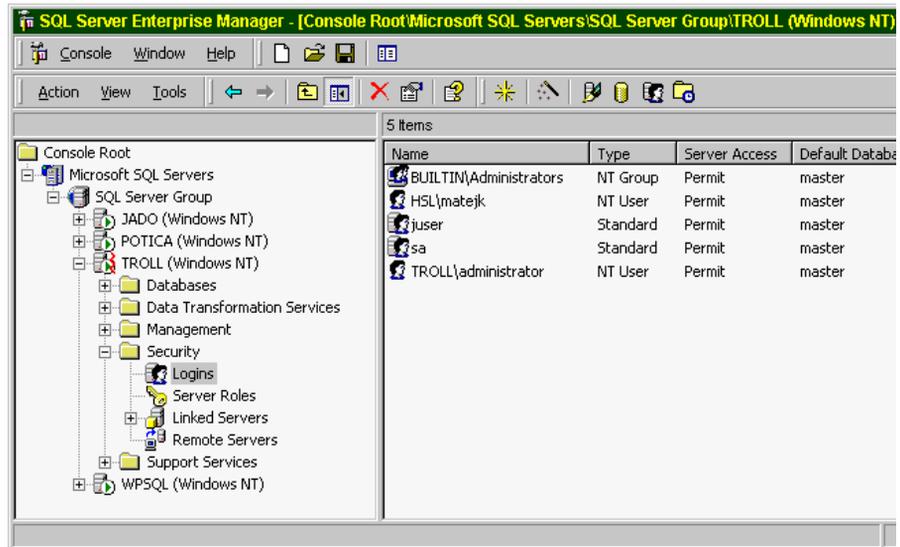
- Microsoft SQL Server must be online during the configuration procedure.
- Configuration must be performed for every single Microsoft SQL Server system.

The configuration consists of setting the user name and password for the Data Protector services. Thereafter, the services are able to connect to the Microsoft SQL Server and operate under the specified account.

The user must have appropriate permissions to run backup and restore on the respective Microsoft SQL Server.

You can check this using the Microsoft SQL Server Enterprise Manager.

Figure 1-3 Microsoft SQL Server Users



You need to define the way in which the Data Protector sessions will run on the Microsoft SQL Server system, using either Data Protector Inet account (in most cases the system account) or a specified user account (preferred option).

Configuration Procedure

To configure the Microsoft SQL Server while creating the first backup specification or while changing the configuration, proceed as follows:

In the Configure MS SQL dialog box, select either Integrated Security or Standard Security. See Figure 1-4.

Figure 1-4 Configuring the Microsoft SQL Server



NOTE

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

- If you use Standard Security, provide a user name in the format <DOMAIN>\<user_name> and a password for a Microsoft SQL Server user, who must have permissions to run backup and restore of the Microsoft SQL Server.
- If you use Integrated Security, the Data Protector SQL Server integration will use the Data Protector Inet account to connect to the Microsoft SQL Server.

See Microsoft SQL Server documentation for more detailed information about security and for a description of the two connection types.

Click OK to confirm the configuration.

What Happens?

The login information is written to the Data Protector Microsoft SQL Server configuration file on the Cell Manager:

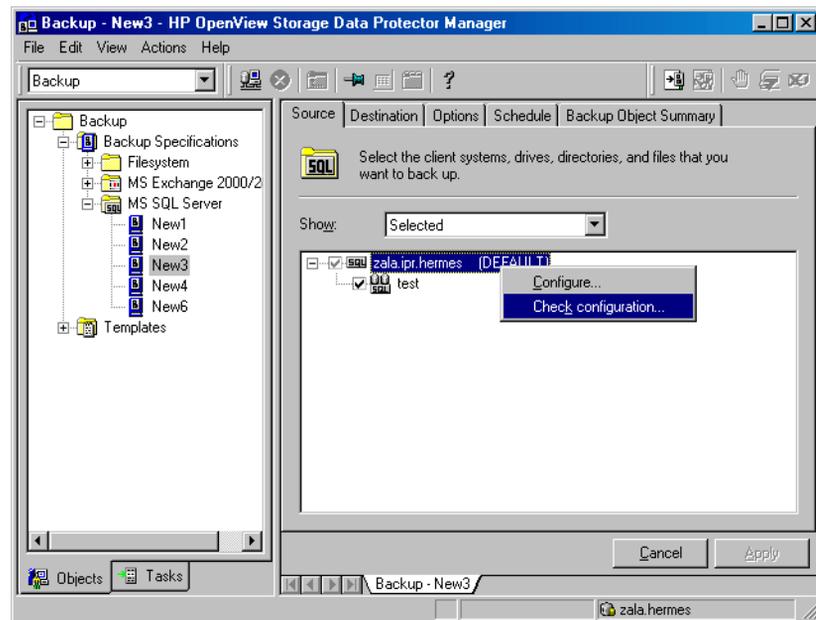
```
<Data_Protector_home>\Config\server\Integ\Config\MSSQL\<hostname>%<instance name> (Windows Cell Manager) or  
/etc/opt/omni/server/integ/config/MSSQL/<hostname>%<instance name> (UNIX Cell Manager).
```

Changing and Checking the Microsoft SQL Server Configuration

You can change the configuration of a specific Microsoft SQL Server and its instance any time after you have created at least one backup specification for this Microsoft SQL Server. To change the configuration, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then MS SQL Server. Click an existing backup specification for the Microsoft SQL Server for which you want to change the configuration.
3. In the Source property page, right-click the name of the Microsoft SQL Server and select Configure.
4. Configure the Microsoft SQL Server as described in “Configuring Microsoft SQL Server” on page 14.
5. Right-click the name of the Microsoft SQL Server and select Check Configuration. See Figure 1-5.

Figure 1-5 Checking Configuration



Configuring the Integration

Once you start checking the configuration procedure, the Data Protector service reads the login information from the configuration file.

Configuring a Microsoft SQL Server Backup

To configure a Microsoft SQL Server backup, perform the following steps:

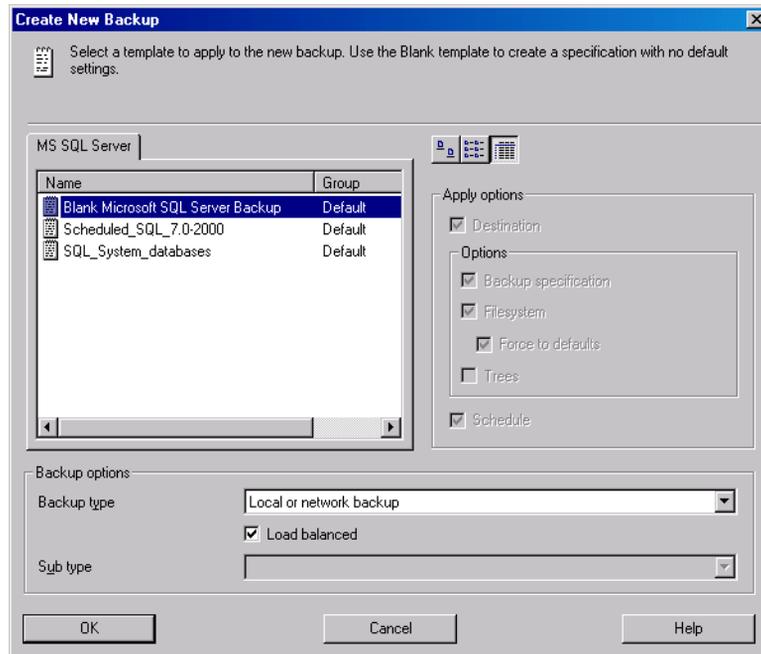
1. Configure the devices you plan to use for a backup. Refer to the online Help index keyword “configuring devices” for instructions. See also “Performance Tuning” on page 39 for advanced options.
2. Configure media pools and media for a backup. Refer to the online Help index keyword “creating media pools” for instructions.
3. Create a Data Protector Microsoft SQL Server backup specification.

Creating a Backup Specification

To create a Microsoft SQL Server backup specification, perform the following steps:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click MS SQL Server and click Add Backup.
3. In the Create New Backup dialog box, select the Blank Microsoft SQL Server Backup template. See Figure 1-6.

Figure 1-6 **Selecting a Blank Template**



Click OK.

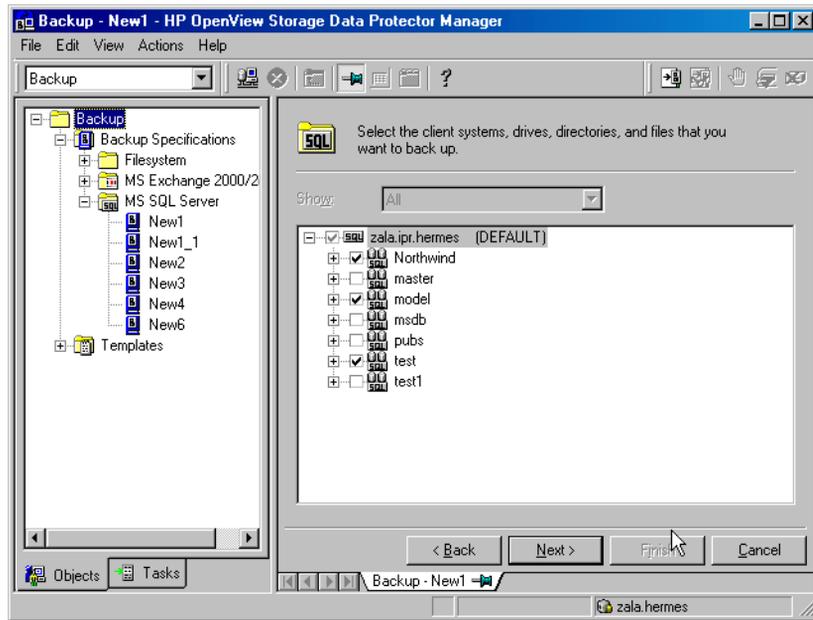
4. In the Client drop-down list, select the Microsoft SQL Server system. If the application is cluster-aware, select the virtual server of the Microsoft SQL Server resource group.

In the Application database drop-down list, leave the instance name for Microsoft SQL Server.

Click Next.

5. If the client has not been configured yet, the Configure Microsoft SQL dialog box appears. See “Configuring Microsoft SQL Server” on page 14 for detailed steps.
6. Select the Microsoft SQL Server databases you want to back up.

Figure 1-7 Selecting Backup Objects



Click Next.

7. Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**.

You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror. The minimum number of devices required for mirroring Microsoft SQL Server integration objects equals the number of devices used for backup.

For detailed information on the object mirror functionality, see *HP OpenView Storage Data Protector Administrator's Guide*.

Click Next.

8. Select the backup options.

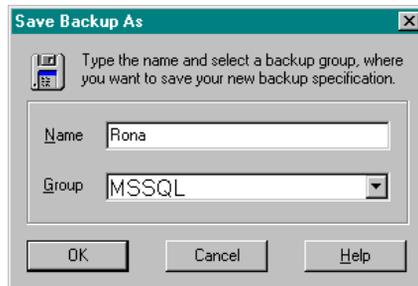
For information on **Backup Specification Options** and **Common Application Options**, refer to the online **Help**.

For information on Application Specific Option, see “Microsoft SQL Server Specific Backup Options” on page 21 or online Help.

Click Next.

9. Optionally, schedule the backup. For information on scheduler, press **F1**.
10. Save the backup specification. It is recommended that you save all Microsoft SQL Server backup specifications in the MSSQL group. See the following figure.

Figure 1-8 Saving a Backup Specification



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

See “Backing Up Microsoft SQL Server Databases” on page 25 for information on starting a backup.

Microsoft SQL Server Specific Backup Options

The Microsoft SQL Server specific backup options are specified using the Data Protector GUI by clicking the Advanced tab in the Application Specific Options group box.

The following are the Microsoft SQL Server application specific backup options:

- Concurrent streams** This option is available only if the whole Microsoft SQL Server is backed up. It sets the number of user-specified concurrent streams (devices) used for a backup. To set the number of concurrent streams for a particular Microsoft SQL

Server database (if one or more Microsoft SQL Server databases were selected for a backup), see “Object Specific Options” on page 23.

Fast Direct Mode This option can only be used with locally connected devices in order to optimize performance. This operational mode must be combined with special device settings. See “Performance Tuning” on page 39 for details.

NOTE

It is recommended that the local devices with special block-size settings be dedicated to the Microsoft SQL Server high performance backup only.

Check Database Integrity By selecting this option the structure of the MS SQL data is verified before the backup. In other words, the MS SQL data integrity validation is performed.

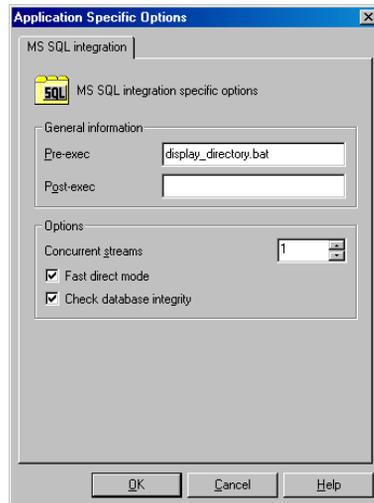
NOTE

If the check fails, the backup session is still completed (with warnings).

Pre-exec Specifies a command with arguments or a script that will be started on the Microsoft SQL Server before the backup starts. The command/script is started by the Data Protector `sql_bar.exe` and must reside in the `<Data_Protector_home>\bin` directory of the Microsoft SQL Server system. Only the filename must be provided in the backup specification.

Post-exec Specifies a command with arguments or a script that will be started on the Microsoft SQL Server after the backup. The command/script is started by the Data Protector `sql_bar.exe` and must reside in the `<Data_Protector_home>\bin` directory of the Microsoft SQL Server system. Only the filename must be provided in the backup specification.

Figure 1-9 **Application Specific Options**



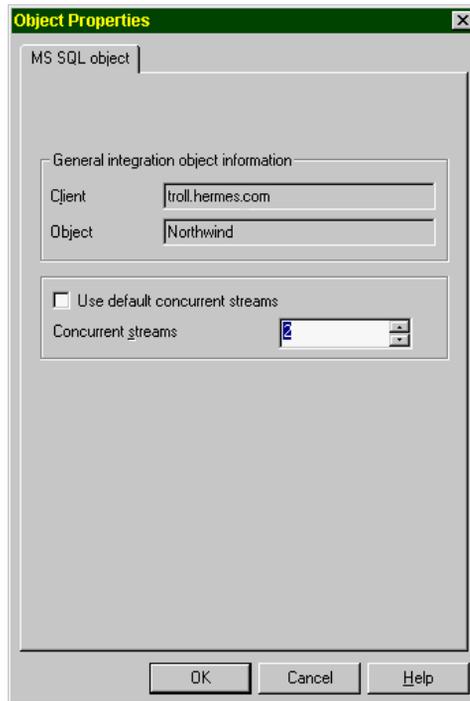
Object Specific Options

If you have specified one or several databases for a backup, so that the backup specification is started as a database backup (as opposed to a whole server backup), you can set the backup options on a single database level.

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

In the Backup Specification Summary property page, double-click an object to open the Object Properties window, where you can customize object properties for the respective database.

Figure 1-10 **Object Properties**



The following backup options can be selected per backup object:

Use default concurrent streams

You may keep this box checked so that the number of concurrent streams is defined by Data Protector using all available devices.

Concurrent streams

The number of user-specified concurrent streams (devices) used for backup. VDI supports up to 32 virtual devices per database.

Backing Up Microsoft SQL Server Databases

To run an online backup of a Microsoft SQL Server database, use any of the following methods:

Backup Methods

- Schedule a backup of an existing Microsoft SQL Server backup specification using the Data Protector Scheduler.
- Start an interactive backup of an existing Microsoft SQL Server backup specification using the Data Protector GUI or the Data Protector CLI.

For information on starting an interactive backup using the Data Protector CLI, refer to the `omnib` man page.

Scheduling a Backup

Scheduling a backup specification means setting time, date, and type of a backup that starts unattended once the scheduling options are defined and saved in the backup specification.

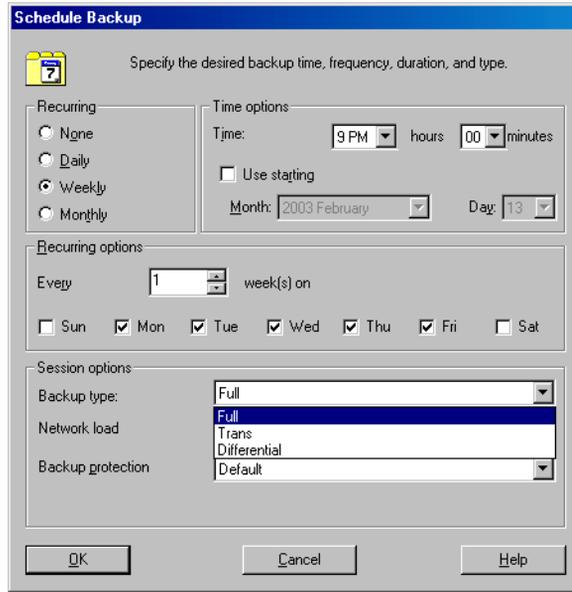
For more information on scheduling, refer to the online Help index keyword “scheduled backups”.

To schedule a Microsoft SQL Server backup, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then MS SQL Server.
3. Double-click the backup specification you want to schedule and click the Schedule tab.
4. In the Schedule page, select a date in the calendar and click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options. See Figure 1-11.

You can select one of the following backup types: full, differential, or transaction log backup. See “Integration Concept” on page 6 for a detailed description of backup types.

Figure 1-11 **Scheduling Backups**



Click OK and then Apply to save the changes.

Running an Interactive Backup

An interactive backup can be performed any time after a backup specification has been created and saved.

To start an interactive backup of a Microsoft SQL Server database, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then MS SQL Server.
3. Right-click the backup specification, and then select Start Backup.

In the Start Backup dialog box, select the Backup type and Network load options. For information on these options, click Help.

Click OK.

Restoring a Microsoft SQL Server Database

You can restore a Microsoft SQL Server object using the Data Protector GUI or using the Data Protector CLI.

Prerequisite

Before you start a restore session, verify that the database is not being used by any user.

Restoring Using the Data Protector GUI

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

If the database already exists and has a different structure, the restore will fail unless you select the Force restore over existing database option.

See “Restore Options” on page 31 for a detailed description.

General restore options that apply to all objects within a restore session, such as Restore database to another Microsoft SQL Server and Restore using a different device can be combined with the object-specific restore options, which are the following:

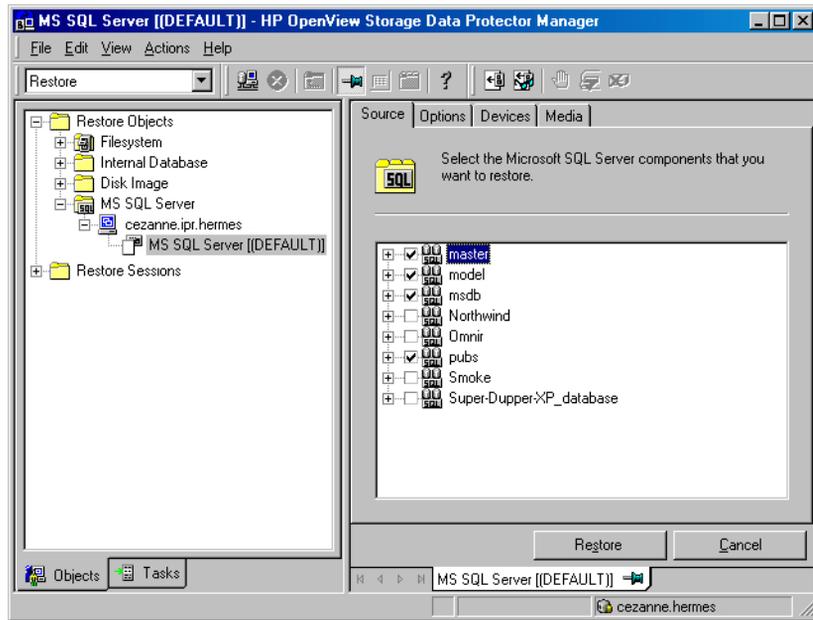
- Point-in-time restore
- Recovery completion state
- Force restore over existing database

In this way, you can choose among several restore scenarios.

Restore Procedure To restore the Microsoft SQL Server databases using the Data Protector GUI, proceed as follows:

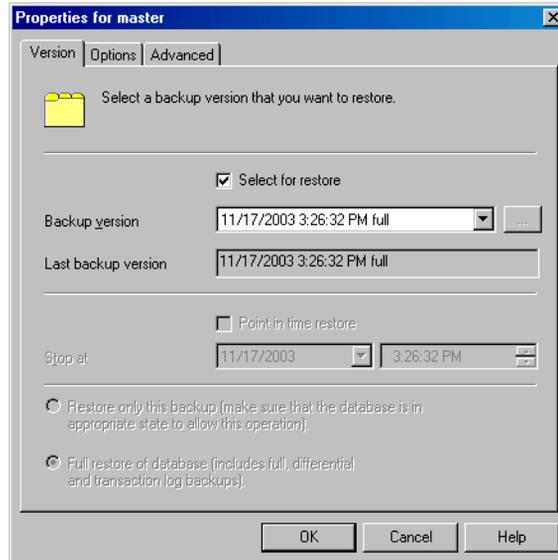
1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand Restore Objects, MS SQL Server, and then select the MS SQL Server from which you want to restore. A list of backed up objects is displayed in the Results Area.
3. Select the backed up Microsoft SQL Server databases you want to restore. See Figure 1-12.

Figure 1-12 Restore Objects



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

Figure 1-13 Selecting the Object Specific Options



Here you can select the version (date of a backup) from which you want to restore and choose the Microsoft SQL Server specific restore options. See “Restore Options” on page 31 for details about these options. Click OK.

4. In the Options property page, specify whether you want to restore your data to another client or instance. In this case, you have to specify new locations for the databases you want to restore. See “Restore Options” on page 31.

NOTE

When you click the Options tab, Data Protector browses the cell for the running Microsoft SQL Server instances that can be selected as target instances for restore. If no instances are found, the Restore to another instance option is automatically disabled and the message There are no instances on this client system is displayed.

Select one of the following Restore actions:

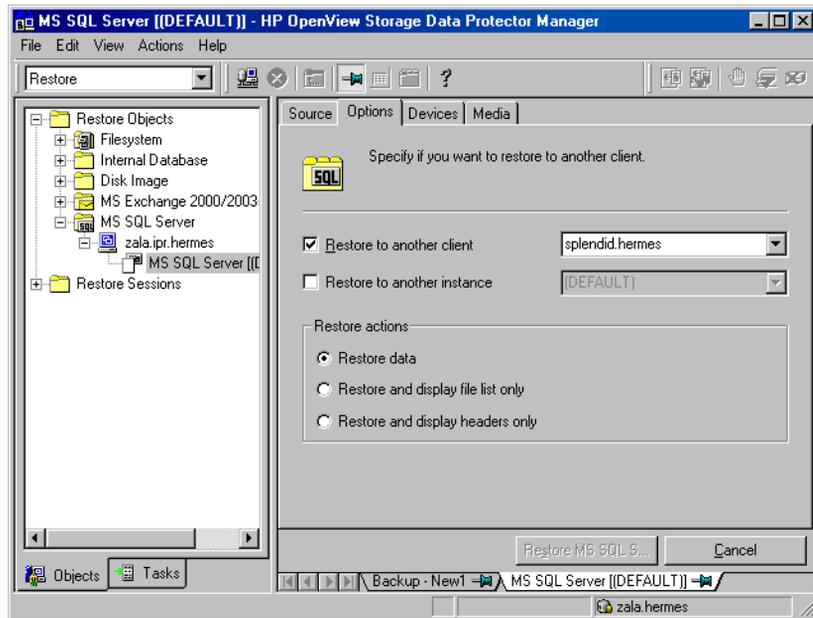
- Restore data (default). Select this action to restore the whole database.

Integrating Microsoft SQL Server and Data Protector

Restoring a Microsoft SQL Server Database

- Restore and display file list only. Use this action if you do not know the original file names. In this case, the list of files backed up in a particular backup session is displayed.
- Restore and display headers only. Select this action if you need specific details about the database backup. The SQL Server's header information is displayed.

Figure 1-14 Restore Options



5. Click Devices and then Media to select the devices and media to be used for the restore.

Note that you can use a different device for the restore than the one used for the backup. Refer to the “Restoring Under Another Device” section in the *HP OpenView Storage Data Protector Administrator’s Guide* for more information on how to perform a restore using another device.

IMPORTANT

If the devices used for the restore are not those used for the backup, select the same number of devices in the `Devices` property page as you used when you backed up the databases.

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

Click `Finish` to start the restore session.

Restore Options

See also “Before You Call Support” on page 53 for restore options available from the command line only.

Backup Version

Specify the backup session from which the selected objects will be restored.

Point-in-Time Restore

Point-in-time restore means that a user can specify a point in time to which the database state must be restored. After recovery, the database is recovered in the state it was at the specified date and time.

Only transaction log records written before the specified date and time are applied to the database.

Point-in-time restore is specified by selecting a backup version and by setting the `Stop at` option.

Stop at

The `Stop at` option specifies the exact time when the rollforward of transactions will be stopped. Therefore, the backup you restore from must include transaction log backups so that the Microsoft SQL Server can recover the database to a particular point in time.

This option cannot be used with `NORECOVERY` or `STANDBY`. If you specify a `Stop at` time that is after the end of the `RESTORE LOG` operation, the database is left in a non-recovered state, just as if `RESTORE LOG` had been run with `NORECOVERY`.

Restore only this backup

If you have restored a version of the database and left it in a non-operational or standby state, then you can subsequently restore differential or transaction log backups one by one, leaving each version non-operational to restore additional backups.

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

The existing database residing on the target restore server system will be overwritten.

If a database with the same name as the one that you want to restore already exists on the server, and it has a different internal structure, then the Microsoft SQL Server does not let you rewrite the database without the Force Restore over existing database option turned on.

Recovery Completion State

These options let you select the state of the database after the recovery. You may select among the following:

- Leave the database operational. Once the last transaction log has been restored and the recovery has completed, the database is already operational.
- Leave the database non-operational after the last transaction log has been restored. You may further restore additional transaction logs one by one.
- Leave the database as read-only. You may further restore transaction logs before the database is set to read-write mode.

Restore database with a new name

This option lets you restore your database under a different name. You have to specify the database's logical file name and the destination file name (suboptions of the **Restore files to new locations** option) when selecting this option.

Restore files to new locations

This option allows you to restore files to a new location. You need to specify the database's logical file name, and a destination target file name for the specified logical file name. Use this option if you are restoring data to another server, instance or making a copy of the database on the same server.

Restoring to Another Microsoft SQL Server Instance or (and) to Another Microsoft SQL Server

To restore databases to a different Microsoft SQL Server system or (and) to a different Microsoft SQL Server instance, check the prerequisites below:

Limitations

- Restore to another instance is supported only on Microsoft SQL Server 2000 and higher.

Prerequisites

- Both Microsoft SQL Servers must have the same local settings, such as code page and sort order. This information is displayed in the session monitor for each backup.
- The Target Microsoft SQL Server must be in the same Data Protector cell as the original Microsoft SQL Server and it must be configured.

Proceed as follows:

1. If the target Microsoft SQL Server is not yet configured, create a backup specification and configure the server.
See "Configuring a Microsoft SQL Server Backup" on page 18.
2. Select the databases you want to restore and their versions.

Restoring a Microsoft SQL Server Database

3. Select whether you want to restore the data to another Microsoft SQL Server client or (and) to another Microsoft SQL Server instance:

- To restore the data to another Microsoft SQL Server client, select the `Restore to another client` option and then select the target client from the drop-down list.
- To restore the data to another Microsoft SQL Server instance, select the `Restore to another instance` option. If you do not see the list of instances in the drop-down list, enter the instance name by yourself.
- To restore the data to another Microsoft SQL Server client and to another Microsoft SQL Server instance, make sure you entered the name of the instance that exists on the target client. Otherwise, the restore will fail.

Also, specify the new locations for the databases you want to restore.

4. Start restore.

See “Restoring a Microsoft SQL Server Database” on page 27

Restoring Using the Data Protector CLI

A restore session can also be started from the command line. Switch to the `<Data_Protector_home>\bin` directory on any client system within the Data Protector cell that has the Data Protector User Interface installed, and run the following command:

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

Provide the `Session_ID` of the backup session. In case of object copies, do not use the copy session ID, but the object's backup ID, which equals the object's backup session ID.

Example

To start a restore of the database RONA that was running on the Microsoft SQL Server called Alma, execute the following command to restore the latest backup session to the same destination:

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

Disaster Recovery

Disaster recovery is a very complex process that involves products from several vendors. As such, a successful disaster recovery depends on all the vendors involved. The information provided here is intended to be used as a guideline.

Check the instructions from the database/application vendor on how to prepare for a disaster recovery. See also the “Disaster Recovery” chapter in the *HP OpenView Storage Data Protector Administrator’s Guide* for instructions on how to approach system disaster recovery using Data Protector.

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.
2. Install, configure, and initialize the database/application so that data on the Data Protector media can be loaded back to the system. Consult the documentation from the database/application vendor for a detailed procedure and the steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and in the troubleshooting section.
4. Start the restore. When the restore is complete, follow the instructions from the database/application vendor for any additional steps required to bring the database back online.

In case of a disk failure, you need to recover the operating system prior to any other recovery tasks. Data Protector disaster recovery is used to bring the operating system back up on the damaged system.

Refer to the *HP OpenView Storage Data Protector Administrator’s Guide* for more information about the Data Protector disaster recovery.

The next step is restoring the Microsoft SQL Server database.

IMPORTANT

If you need to reinstall SQL Server, ensure that you use the original local settings. Before you perform a restore to another client, ensure that the local settings on the restore target system match the original.

You need to recover the master database first. See the following section for the procedure on how to do this.

Recovering the Master Database

The master database holds the vital information about the Microsoft SQL Server as a whole. If you lose the master database, you cannot access any other databases and all of them are therefore unavailable.

If the master database gets corrupted or lost, recover the master database first to make the Microsoft SQL Server operational.

Next, restore all other databases or reattach them, as described in the section “Recovering User Databases” on page 38.

To recover the master database, proceed as follows:

1. Rebuild the master database

Create the basic master database in order to make the Microsoft SQL Server operational, since some databases might be corrupted or might contain inconsistent data:

- a. Shut down the Microsoft 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, refer to the *Microsoft SQL Server Books Online*.

2. Set user rights or reconfigure the integration

At this stage, either reconfigure the integration or set user rights.

After the master database is rebuilt, proceed as follows to set user rights using the Microsoft SQL Server Enterprise Manager:

- a. From the server’s desktop, click *Start*, *Programs*, *Microsoft SQL Server 7/Microsoft SQL Server*, *Enterprise Manager* to start the utility.

Setting User Rights

- b. Right-click the server in question, and then select Register Server. Configure the Microsoft SQL Server to use trusted connections.
- c. Close the dialog box and go to Security, Logins.
Select the user rights you want to use (such as sa, <password>).
- d. Then return to the server in question, right-click its name, and then select Register Server.
Enter the account you have just selected in Manage, Logins.
Perform any additional administration tasks that are required to run the SQL Server at this point.

Reconfiguring the Microsoft SQL Server Integration

Reconfigure the Microsoft SQL Server Integration.

See “Configuring a Microsoft SQL Server Backup” on page 18 for instructions.

3. Stop all Microsoft SQL Server services

You can do this using the Windows desktop:

Start, Programs, Microsoft SQL Server 7.

Start the SQL Service Manager and stop the services.

4. Start the Microsoft SQL Server service in single-user mode

This can be done from the Windows desktop:

- a. In the Control Panel, go to Administrative Tools, Services.
- b. Select the MSSQL Server Service.
- c. Enter -m as a start-up parameter and start the services.

5. Restore the master database using the Data Protector Manager

6. Restart the Microsoft SQL Server services in normal mode

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

Remember that you must restore *all* other databases if you perform disaster recovery.

Restoring a Microsoft SQL Server Database

If you restored selected databases only, you need to reattach databases (if they exist on disks) to the newly-rebuilt master database. See the next section for details.

Recovering User Databases

To restore a user database, proceed as described in “Restoring a Microsoft SQL Server Database” on page 27.

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

Suppose you have the following backup sequence:

F D T T D T T T T T

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

Restore is performed and automated by `sql_bar.exe`.

TIP

It is also possible to 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 about restoring the master database, see “Recovering the Master Database” on page 36.

For more information on disaster recovery, refer to the *HP OpenView Storage Data Protector Administrator’s Guide* and *Microsoft SQL Server Books Online*.

Performance Tuning

Performance tuning means customizing your Microsoft SQL Server and Data Protector in a way that enables them to achieve better backup and restore results. You can improve the backup or restore performance of your Microsoft SQL Server by following these guidelines:

1. Ensure that your Microsoft SQL Server database files are on separate disks.
2. Calculate the number of devices to be used in parallel. The main consideration is to select a number of devices to match the bandwidth of the incoming data stream and to identify the bottleneck. This can be either the network, if devices are connected to remote systems, or the Microsoft SQL Server itself, if the devices are connected locally.

As the bandwidths of networks are most often either ~1 MB/s (10 Mbit Ethernet), or ~10 MB/s (100 Mbit Ethernet), though the actual throughput is usually lower, you will not need more than one fast device, such as a DLT 7000 for a remote backup.

There are two possibilities when you have devices connected locally:

- a. Devices are dedicated to local backups of the Microsoft SQL Server and it is very likely that backup/restore performance is important. Fast direct mode should be used.
- b. Devices are shared within the Data Protector cell and backup/restore performance is not very important. Fast direct mode should be disabled.

Determine the maximum backup speed by performing a backup to a few null file devices on a local server, and select the number of devices that fit 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 the block size for local backup devices.

- Enable/Disable the `Fast direct mode` option.

`Fast direct mode` is an application-specific option per backup specification, which enables Data Protector to read data directly from the SQL Server's shared memory, and can therefore increase the backup/restore speed to/from local devices. See Figure 1-9 on page 23.

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

You should disable the `Fast direct mode` option (as well as the special local device settings) if backup performance is not very critical and/or other data is also backed up to devices connected to the Microsoft SQL Server system.

The `Fast direct mode` option can only be used for local devices. This mode is ignored for remote devices.

- Set the block size.

Before you activate the `Fast direct mode` option, a special block size has to be set for the devices, as referenced in the backup specification. The adjusted block sizes can be calculated as follows:

$$\text{block size (kB)} = 64 * N + 4 \quad (N=1, \dots, 64)$$

$$\text{block size (kB)} = 68, 132, \dots, 4100 \text{ kB}$$

All selected devices must have the same block size.

Some performance improvement can be gained by specifying a block size larger than 68 KB. You can increase the block size step by step and compare the performance achieved for each step.

Note that it is recommended to set the block size to 68 KB.

If you keep the `Fast direct mode` inactive, you do not need to adjust the local device settings; otherwise, adjust the block size for all selected local devices, where the block size must meet the above requirements.

You can adjust the block size during the initial device definition for local devices by checking the attached check box and selecting the block size. See Figure 1-15 on page 41.

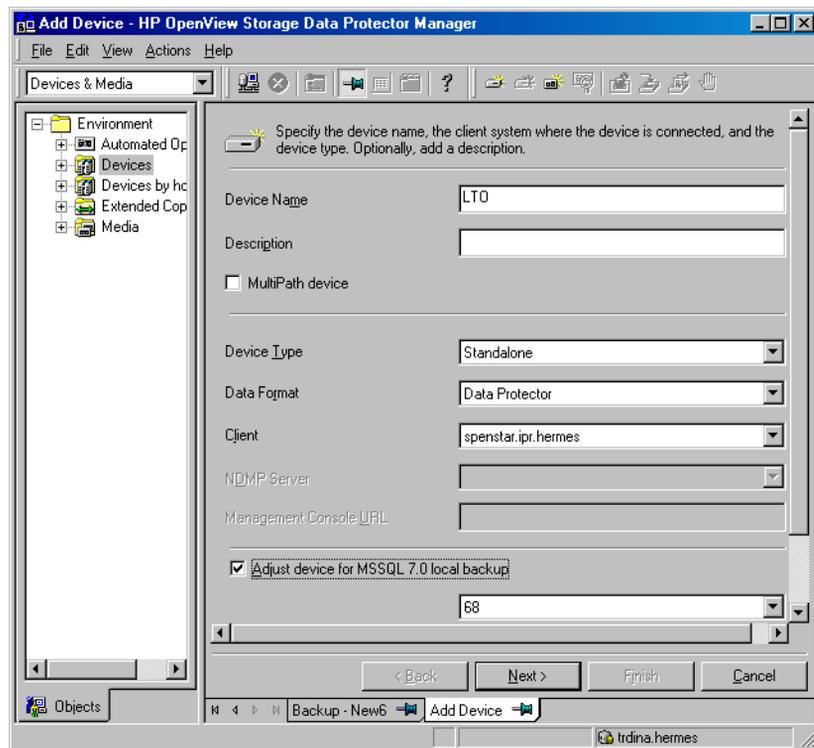
You can also modify the block size later, however you will have to first calculate the block size by using the formula above and then insert the value as shown in Figure 1-16 on page 42.

- Modify the registry.

To use a 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.

Refer to the online Help index keyword "changing block size" for information on how to adjust registry values for block sizes larger than 56 KB.

Figure 1-15 **Creating an Adjusted Local Device**



To modify the block sizes of an existing device, proceed as follows:

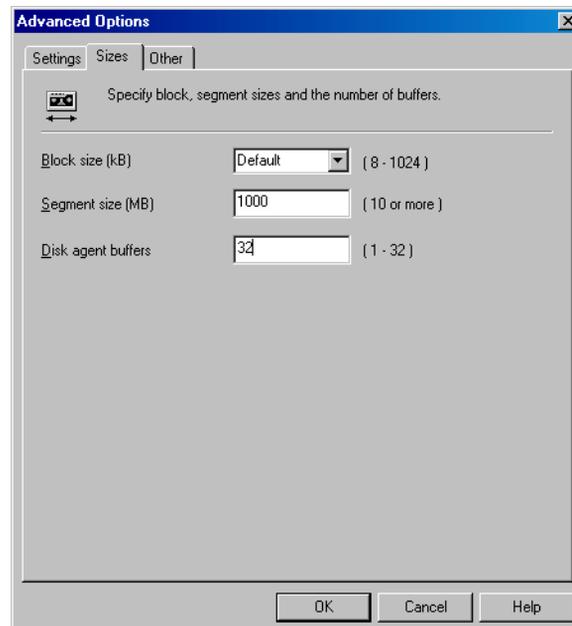
- a. In the HP OpenView Storage Data Protector Manager, switch to the Devices & Media context.

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

- b. In the Advanced Options window, click Sizes, where you can modify block sizes.

Figure 1-16

Advanced Options



If the `Fast direct mode` option has been activated and not all selected local devices in a backup specification are adjusted accordingly, you will get the following warning message when saving the backup specification:

Figure 1-17 **Device Block Sizes Are Not Adjusted**



4. Scheduling

The backup schedule depends on how many transactions are done on the server. Generally, it is not wise to let the transaction log files grow over a certain limit, which depends on the certain production database and the size of its transaction log files. These are some general rules on how to schedule backups for production databases:

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

You should schedule full and differential backups when the server is not heavily loaded (nights and weekends), while transaction log backups should be done several times during the day.

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

For more information, refer to the *Microsoft SQL Server Books Online* and *HP OpenView Storage Data Protector Administrator's Guide*.

Monitoring a Microsoft SQL Server Backup and Restore

The Data Protector GUI enables you to monitor current or view previous backup and restore sessions.

Monitoring is automatically activated when you start a backup or restore interactively.

See “Running an Interactive Backup” on page 26 and “Before You Call Support” on page 53 for information about how to use the command line.

Monitoring Current Sessions

To monitor a currently running session using the Data Protector GUI, proceed as follows:

1. In the Context List, click Monitor.
In the Results Area, all currently running sessions are listed. See Figure 1-18.
2. Double-click the session you want to monitor.

Figure 1-18

Monitoring a Current Session



Clearing Sessions To remove all completed or aborted sessions from the Results Area of the Monitor context, proceed as follows:

1. In the Scoping Pane, click Current Sessions.
2. In the Actions menu, select Clear Sessions. Or click the Clear Sessions icon on the toolbar.

To remove a particular completed or aborted session from the current sessions list, right-click the session and select `Remove From List`.

NOTE

All completed or aborted sessions are automatically removed from the Results Area of the `Monitor` context if you restart the Data Protector GUI.

For detailed information on a completed or aborted session, see “Viewing Previous Sessions”.

Viewing Previous Sessions

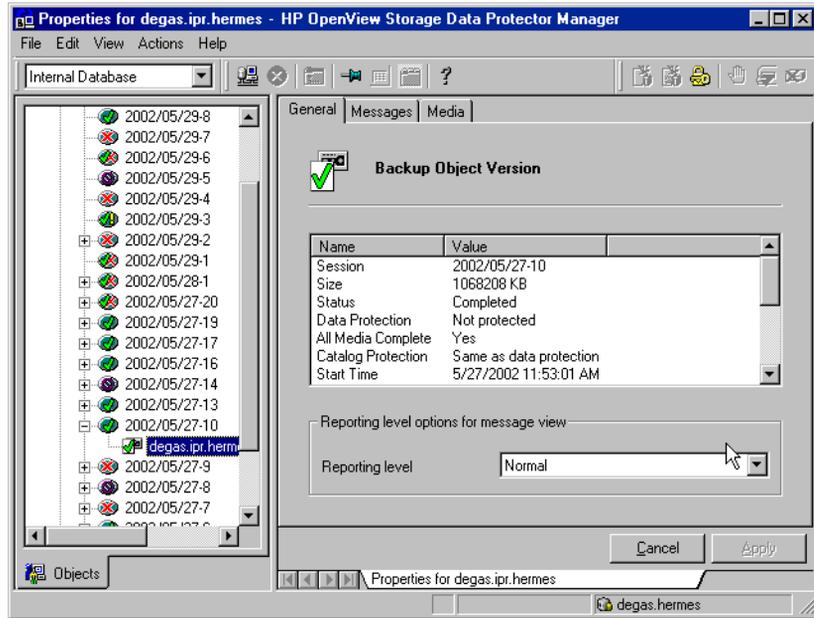
To view a previous session using the Data Protector GUI, proceed as follows:

1. In the `Context List`, click `Internal Database`.
2. In the `Scoping Pane`, expand `Sessions` to display all the sessions stored in the IDB.

The sessions are sorted by date. Each session is identified by a session ID consisting of a date in the `YY/MM/DD` format and a unique number.

3. Right-click the session and select `Properties` to view details on the session.
4. Click the `General`, `Messages` or `Media` tab to display general information on the session, session messages, or information on the media used for this session, respectively. See Figure 1-19.

Figure 1-19 Viewing Previous Sessions



Troubleshooting

This section lists general checks and verifications, and problems you might encounter when using the Data Protector Microsoft SQL integration.

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

Before You Begin

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

Configuration Problems

If a configuration procedure does not work:

- Check that the Microsoft SQL Server services are running.
- Examine system errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server, functioning as a Data Protector client.
- Connect to the Microsoft SQL Server via Microsoft SQL Server Enterprise Manager using the same login ID as the one specified in the Data Protector Configuration dialog box.
- Perform a backup of Microsoft SQL Server databases using the Microsoft SQL Server Enterprise Manager.

If the backup fails, fix any Microsoft SQL Server problems, and then perform a backup using Data Protector.

Using the Data Protector Command Line to Check Configuration

The configuration can also be performed from the command line. Enter the following string at the `<Data_Protector_home>\bin` directory on the Microsoft SQL Server system:

```
sql_bar config -dbuser:<dbuser> -password:<password>  
-appsrv:<appsrv> [-instance:<instance name>]
```

Enter the same information as if using the Data Protector GUI:

- The username and password of the SQL Server user who has permissions to back up and restore the SQL Server backup objects.
- The name of the SQL Server system.

To check configuration using the command line, enter the following string from the `<Data_Protector_home>\bin` directory on the Microsoft SQL Server computer:

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

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

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

```
*RETVAl*8523
```

If you want the information about the existing configuration, enter the following string:

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

If the optional parameter `-instance:<instance_name>` is not specified, the configuration for the default instance is returned.

What Happens?

Once you start the configuration from the command line, the login information will be written in the `<hostname>%<instance name>` file in the `<Data_Protector_home>\Config\Server\Integ\Config\MSSQL` directory.

Once you start checking the configuration procedure, the Data Protector service reads login information from the `<hostname>%<instance name>` file on the Data Protector Cell Manager and tries to connect to the server using this account.

Miscellaneous Problems

Problems

- The integration is properly configured and the backup of all databases fails after a timeout, with an error message similar to the following:

```
[Warning] From: OB2BAR@paradajz.hermes.com "MSSQL70" Time:
3/14/2000 8:19:22 PM
Error has occurred while executing SQL statement.
```

```
Error message: '<Microsoft SQL-DMO (ODBC SQLState: 42000)> Error
number: bc5'
```

```
[Microsoft] [ODBC SQL Server Driver] [SQL Server]Backup or restore
operation terminating abnormally.'
```

```
[Critical] From: OB2BAR@paradajz.hermes.com "MSSQL70" Time:
3/14/00 8:19:24 PM
```

```
Received ABORT request from SM => aborting.
```

- The SQL Server's error log contains an entry similar to the following:

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

- The SQL Server's VDI.LOG file contains an entry similar to the following:

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

Cause

The SQL Server service and the Data Protector Inet service are running under different accounts. The SQL Server integration cannot access the SQL Server's data for backup due to security problems.

Solution

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

Backup Problems

1. If a backup does not work:
 - Verify the configuration file to check if the Cell Manager is correctly set on the Microsoft SQL Server that is functioning as a Data Protector client.
 - Check that the Microsoft SQL Server services are running.
 - Check that `sql_bar.exe` is installed on the system.
 - Examine system errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server system.

Check also the `errorlog` and `VDI.log` files in the `<MSSQL>\log` directory on the server system.
 - Perform a backup of Microsoft SQL Server databases using the Microsoft SQL Server Enterprise Manager.

If the backup fails, fix the Microsoft SQL Server problems and perform a backup using Data Protector.
2. If during the creation of a backup specification you do not see the instance of the Microsoft SQL Server as the application database, enter the instance name by yourself. When the “not-named instance” is not displayed, the `DEFAULT` string must be inserted as an application database.
3. When performing a backup, Microsoft SQL Server reports that the database backup cannot take place because of inappropriate user rights.

If the Data Protector Manager (and `sql_bar.exe`) reports that the integration is properly configured, verify that the Microsoft SQL Server user has appropriate rights to access the databases that cannot be backed up.

It is recommended that the Microsoft SQL Server system administrator (`sa`) configure the Data Protector Microsoft SQL Server integration.

Backup Hangs if Concurrency Is Set to More Than One and One of the Devices Fails

Problem	Backup of a Microsoft SQL Server can hang if Disk Agent concurrency is set to more than one and one of the devices fails during backup or is not started at all, for example because of a medium error.
Action	Set the device concurrency to one or replace the invalid media.

Restore Problems

- If a restore does not work:
 - Check whether a filesystem backup of the problematic client works. It is much easier to troubleshoot a filesystem backup.
 - Check that the Microsoft SQL Server services are running.
 - Check that `sql_bar.exe` is installed on the system.
 - Examine the system errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server that is functioning as a Data Protector client.

Check also the `errorlog` and `VDI.log` files in the `<MSSQL>\log` directory on the same system.

- The following error has occurred when executing an SQL statement:

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

It is an expected behavior when the master database is restored in a single user mode, so this message should not be treated as an error.

A Restore from an Object Copy Hangs

Problem	A restore from an object copy hangs.
Action	Before restarting the restore: <ul style="list-style-type: none">• Increase the number of Disk Agent buffers for the device used for the restore.

- If all objects of the backup are recorded in the IDB, perform the following steps:
 1. In the Internal Database context of the Data Protector GUI, search for all objects belonging to the same backup. The objects are identified by the same backup ID.
 2. Copy each object in a separate object copy session to a separate device, for example a file library. For each object, use a separate medium with the non-appendable media policy.
 3. Set the highest media location priority for the newly created copies.

Database is Left in Unrecovered State After the “Invalid value specified for STOPAT parameter” Message is Reported in the Data Protector Monitor

Problem If the Invalid value specified for STOPAT parameter Message is Reported in the Data Protector Monitor, the database remains in the unrecovered state as if the RESTORE LOG operation was run with the Leave the database non-operational option.

Action The database can be recovered to the latest point in time by using the Microsoft SQL Query Analyzer. To recover the database, run the following T-SQL command:

```
RESTORE DATABASE <database_name> WITH RECOVERY
```

After the database is recovered, additional transaction logs cannot be applied.

Restoring to Another Client

Problem You want to perform a restore of the Microsoft SQL Server database to another client in the Data Protector cell not configured to use with the Microsoft SQL Server, but the restore does not work.

Action Create the configuration file by configuring the Microsoft SQL integration on this client. See “Configuring the Integration” on page 14.

Database Left in Unrecovered State After the Restore Session Completed Successfully

Problem

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

Action

The database can be recovered to the latest point in time by using the Microsoft SQL Query Analyzer. To recover the database, run the following T-SQL command:

```
RESTORE DATABASE <database_name> WITH RECOVERY
```

After the database is recovered, additional transaction logs cannot be applied.

Before You Call Support

If you have performed all the troubleshooting procedures without solving your problem, you should gather the following information for Data Protector support before you make a call:

1. Provide details about your hardware and software configuration, including the official patches you use, the Microsoft SQL Server version, the SP, the Windows version and the SP.
2. Provide a detailed description of the action you failed to perform. If you had backup problems, attach the backup specification.
3. Provide the information from the following files:
 - `<Data_Protector_home>\log\debug.log`
 - `<MSSQL>\log\errorlog`
 - `<MSSQL>\log\vdi.log`Copy the session output into a file.

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft Exchange integration.

The chapter is organized into the following sections:

“Introduction” on page 57

“Prerequisites and Limitations” on page 60

“Integration Concept” on page 61

“Configuring the Integration” on page 64

“Backing Up Microsoft Exchange Server” on page 72

“Restoring a Microsoft Exchange Server Database” on page 75

“Troubleshooting” on page 85

Introduction

The Data Protector integration with Microsoft Exchange enables you to perform online backup of Microsoft Exchange Server.

Extensible Storage Engine (ESE 98)

Microsoft Exchange Server uses a database technology called **Extensible Storage Engine (ESE 98)** as a storage system for information exchange.

Microsoft Exchange Server uses a common **Application Programming Interface (API)** that provides a unified interface to back up and restore all data that is written to any instance of the ESE database present on the system.

Using new functionality of Exchange Server, you can configure multiple databases for each server. Databases are grouped into **storage groups**. For each storage group a database server instance is running. Up to 4 storage groups and 5 databases per storage group are supported, so that up to 20 databases for each server or cluster can be configured. Each storage group can only run one backup at a time, so that databases within a storage group can only be backed up sequentially. Storage groups can be backed up in parallel.

There are two kinds of databases: mailbox stores and public folder stores. The service that is responsible for storage management is called the **Information Store**. The database that permits compatibility with Exchange 5.5 by emulating an Exchange 5.5 directory service is called the **Site Replication Service**. The database that provides security encryption services is called the **Key Management Service**.

For other information about Microsoft Exchange Server, refer to the *Microsoft Exchange Server Books Online*.

Using the Data Protector Microsoft Exchange Server integration, you can perform online backups and restores of single mailboxes located on a Microsoft Exchange Server system. For more information on configuring, backing up, and restoring single mailboxes, see Chapter 3, “Integrating Microsoft Exchange Single Mailbox and Data Protector,” on page 89.

Advantages

Integrating Data Protector with Microsoft Exchange Server offers several advantages over using the Windows backup utility with support for Exchange Server:

- Central Management for all backup operations

The administrator can manage backup operations from a central point.

- Backup Management

Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.

- Media Management

Data Protector has an advanced media management system, which allows users to monitor media usage and set protection for stored data, as well as organize and manage devices in media pools.

- Scheduling

Data Protector has a scheduler that allows the administrator to automate backups to run periodically. Using the Data Protector Scheduler, you can configure the backups to run unattended, at specified times, if the devices and media are set properly.

- Device Support

Data Protector supports a wide range of devices: files, standalone drives, very large multiple drive libraries, etc.

- Reporting

Data Protector has reporting capabilities that allow you to receive information about your backup environment. You can schedule reports to be issued at a specific time or attached to a predefined set of events, such as the end of a backup session or a mount request.

- Monitoring

Data Protector has a feature that allows the administrator to monitor currently running sessions and view finished sessions from any system that has the Data Protector User Interface installed.

All backup sessions are logged in the IDB, which provides the administrator with a history of activities that can be queried later.

Refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for more information about supported platforms and devices.

Prerequisites and Limitations

- You need a license to use the Data Protector Microsoft Exchange integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Before you begin, make sure that you have correctly installed and configured the Microsoft Exchange and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector Microsoft Exchange Server integration.
 - *Microsoft Exchange Server Books Online* for online information on Microsoft Exchange Server.
- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.
- Preview is not possible for Exchange Server backup and restore sessions.

It is assumed that you are familiar with the Microsoft Exchange database administration and the basic Data Protector functionality.

Integration Concept

The central component of the Data Protector Microsoft Exchange integration is the Data Protector `ese_bar.exe` executable, which is installed on the Microsoft Exchange Server system and which controls the activities between Microsoft Exchange Server and Data Protector backup and restore processes.

From the perspective of Microsoft Exchange Server, Data Protector is seen as media management software. On the other hand, Microsoft Exchange Server is a Data Protector client from the Data Protector Cell Manager's point of view.

Backup Flow

A Data Protector backup session can be started only from the Data Protector GUI.

The Data Protector Backup Session Manager reads the backup specification and starts the `ese_bar.exe` command on the Microsoft Exchange Server system.

The `ese_bar.exe` command reads the data from Microsoft Exchange Server and passes it to the Data Protector General Media Agents.

Multiple storage groups are backed up in parallel. Multiple databases within a storage group are backed up sequentially. The maximum number of devices used in a backup session equals the number of storage groups you want to back up.

The two types of backup supported by the Data Protector Microsoft Exchange integration are **Full** and **Incremental**.

A full backup selects for a backup the whole database and all log files of a storage group regardless of whether they have been changed since the last backup. Incremental backup selects only log files. After a full or incremental backup, the log files are deleted.

There is only one level of incremental backup, which refers to the previous full or incremental backup, whichever was performed last.

Restore Flow

Using the Data Protector User Interface, you define which objects and object versions to restore. The Data Protector Restore Session Manager is invoked, which then starts `ese_bar.exe` and passes the information

Integrating Microsoft Exchange Server and Data Protector
Integration Concept

about the objects and backup versions on to the backup API. General Media Agents are started by `ese_bar.exe`, and the data flows from the media to the target Microsoft Exchange Server. See Figure 2-1.

Messages from the restore session are sent to the Data Protector Restore Session Manager, which writes the messages and the information regarding the respective session to the IDB.

Figure 2-1 Data Protector Microsoft Exchange Integration Concept

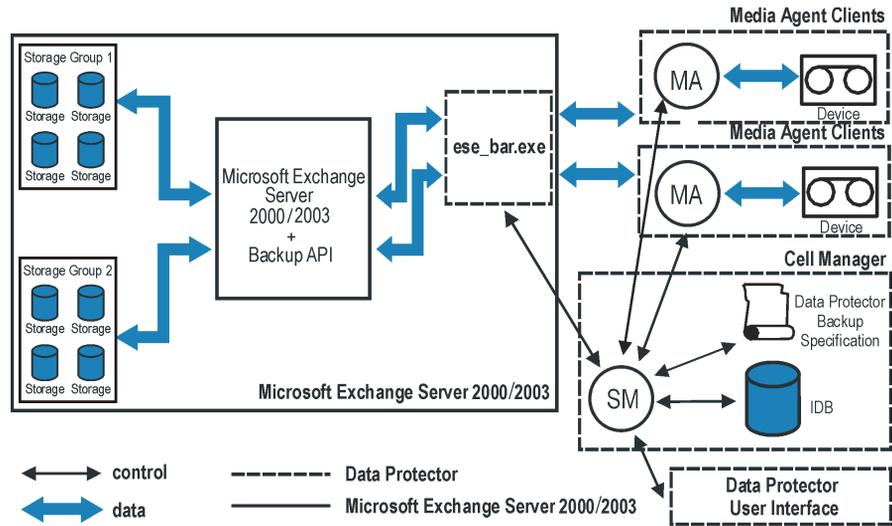


Table 2-1

Legend:

SM	Data Protector Session Manager, which is a Data Protector Backup Session Manager during a backup, or Data Protector Restore Session Manager during a restore.
MA	Data Protector General Media Agent
Backup API	The Microsoft defined interface that enables the data transfer between Data Protector and Microsoft Exchange Server.

Table 2-1

Legend:

Storage group	A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange manages each storage group using a separate server process.
---------------	---

Configuring the Integration

Switching Off the Circular Logging

Before performing an incremental backup of Microsoft Exchange Server, make sure that **circular logging** for storage groups is switched off.

Circular logging is a Microsoft Exchange mode where transaction log files are automatically overwritten as soon as the data they contain is transferred to the database(s).

If turned on, this option reduces disk storage space requirements, but does not allow you to perform incremental backups.

Cluster-Aware Clients

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

Extending the Path Environment Variable

The `<Exchange_home>\bin` directory must be added to the Windows Path environment variable before any operation is performed.

To add this directory, proceed as follows:

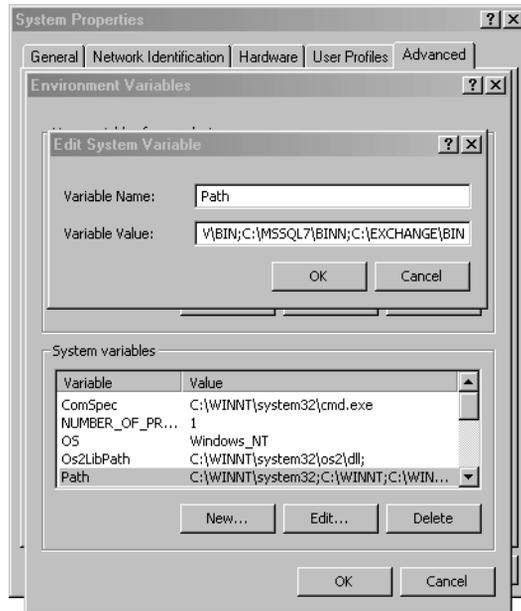
1. In the Microsoft Windows Explorer, right-click My Computer and click Properties.
2. In the Properties dialog box, click Advanced and then Environment Variables.
3. In the Environment Variables dialog box, select Path in the System Variables list and click Edit.
4. Add `<Exchange_home>\bin` in the Variable Value text box and click OK.

See Figure 2-2 on page 65.

Cluster-Aware Clients

If the Microsoft Exchange integration is cluster-aware, add this directory to the Windows Path environment variable on all cluster nodes.

Figure 2-2 Path System Variable



Configuring a Microsoft Exchange Backup

To configure a Microsoft Exchange backup, perform the following three steps:

1. Configure the devices which you plan to use for a backup. Refer to the online Help index keyword “configuring devices” for instructions.
2. Configure media pools and media for a backup. Refer to the online Help index keyword “creating media pools” for instructions.
3. Create a Data Protector Microsoft Exchange Server backup specification.

See the following section for the procedure on creating a backup specification.

See the following section for the procedure on how to create a backup specification.

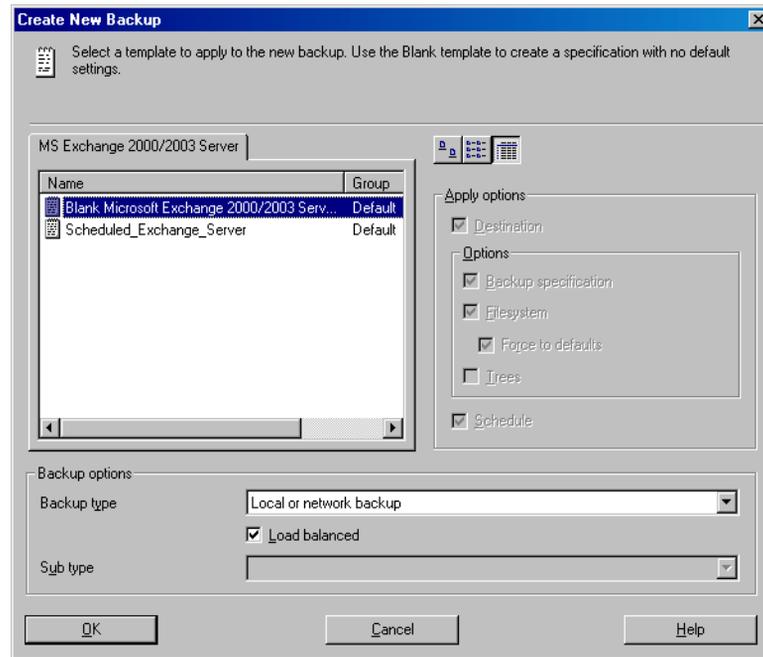
Creating a Backup Specification

To create a Microsoft Exchange backup specification, perform the following steps:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications.
3. Right-click MS Exchange 2000/2003 Server and click Add Backup.
4. In the Create New Backup dialog box, select the Blank Microsoft Exchange 2000/2003 Server Backup template, and click OK.

Figure 2-3

Selecting a Blank Template



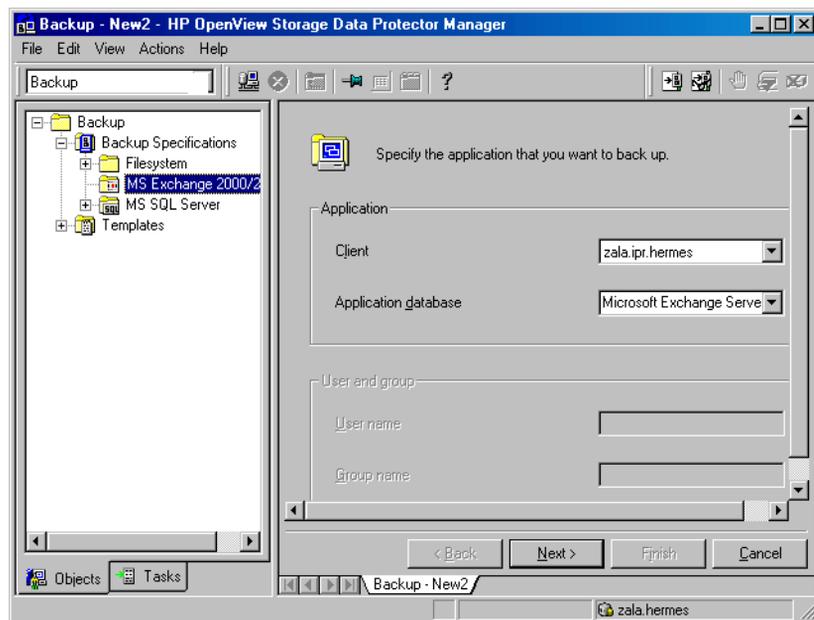
5. In the Client drop-down list, select the Microsoft Exchange Server system. If the application is cluster-aware, select the virtual server of the Microsoft Exchange Server resource group.

In the Application database drop-down list, select one of the following:

- Microsoft Exchange Server (Microsoft Information Store)
Select this item to back up the Information Store.
- Microsoft Exchange Server (Microsoft Key Management Service) (if installed)
Select this item to back up the Key Management Service.
- Microsoft Exchange Server (Microsoft Site Replication Service) (if installed)
Select this item to back up the Site Replication Service.

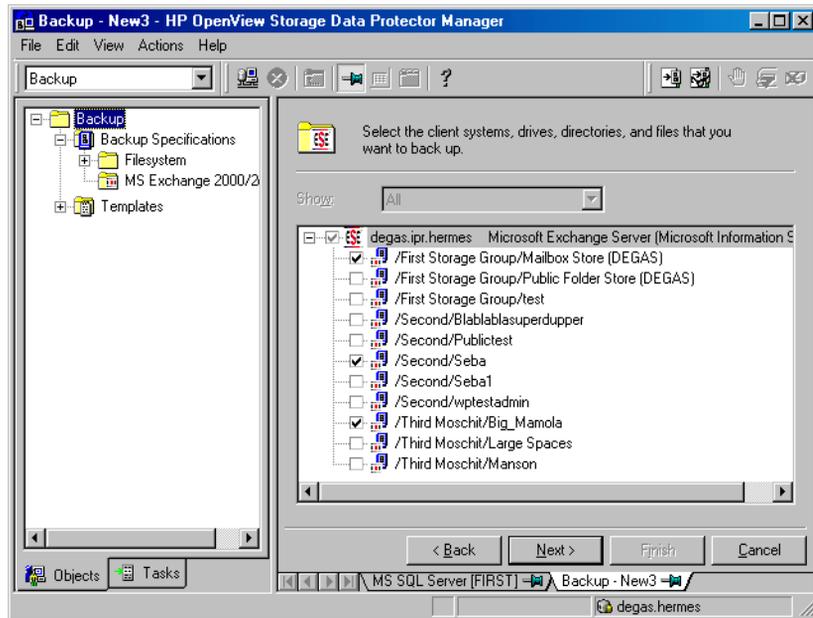
Click Next.

Figure 2-4 Specifying a Client Name and Selecting an Application Database



6. Select the Microsoft Exchange Server databases you want to back up.

Figure 2-5 Selecting Backup Objects



Click Next.

7. Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**.

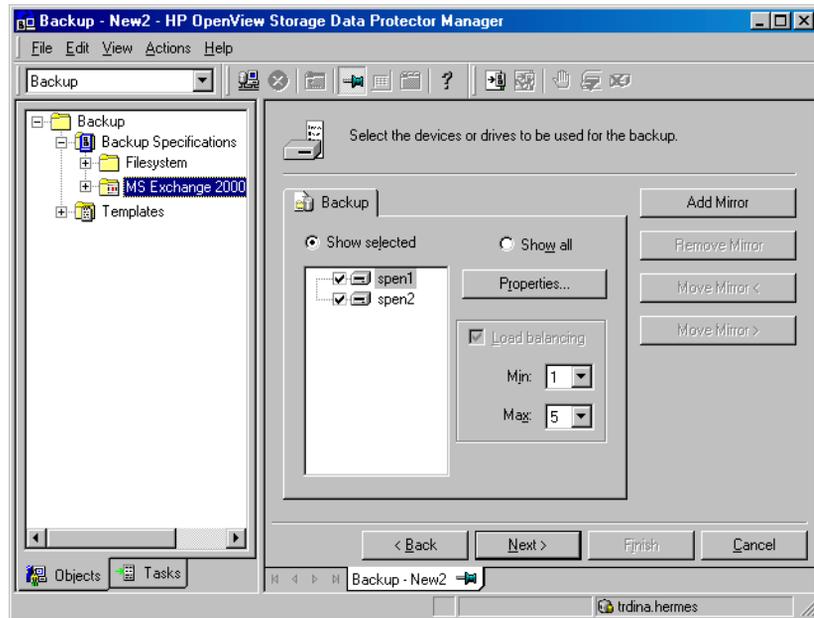
You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror.

For detailed information on the object mirror functionality, see online **Help** index: "object mirroring".

NOTE

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

Figure 2-6 **Selecting Backup Devices**



Click Next to proceed.

8. Select the backup options.

For information on Backup Specification Options and Common Application Options, refer to the online Help.

For information on Application Specific Option, see “Microsoft Exchange Specific Backup Options” on page 70 or online Help.

Click Next.

9. Optionally, schedule the backup. For information on scheduler, press **F1**.
10. Save the backup specification.

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

See “Backing Up Microsoft Exchange Server” on page 72 for information on starting a backup.

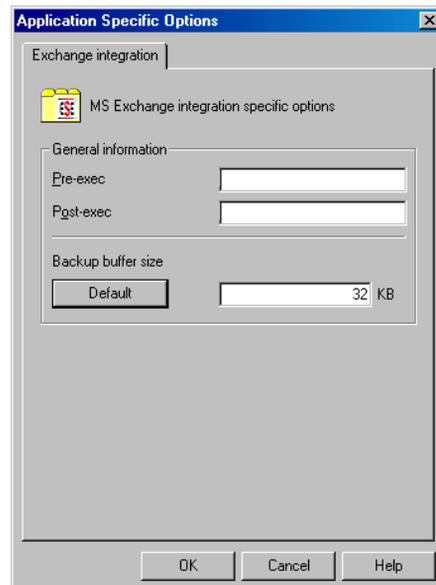
Microsoft Exchange Specific Backup Options

This section describes backup options specific to the Data Protector Microsoft Exchange integration.

You can access these options from the Options property page of a backup specification. Click the Advanced button next to the Application Specific Options. See Figure 2-7.

Figure 2-7

Application Specific Options



The following options can be selected from this window:

Pre-exec

Specifies a command with arguments or a script that will be started on the Microsoft Exchange client before the backup starts. The command/script is started by Data Protector `ese_bar.exe` and must reside in `<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 that will be started on the Microsoft Exchange client after the backup. The command/script is started by Data Protector <code>ese_bar.exe</code> and must reside in <code><Data_Protector_home>\bin</code> directory. Only the filename must be provided in the backup specification.
Backup buffer size	This is the size of the Microsoft Exchange buffer, which is used for transferring data to Data Protector.

NOTE

Note that the pre-exec and post-exec commands/scripts must reside in the `<Data_Protector_home>\bin` directory on the Microsoft Exchange Server system.

Backing Up Microsoft Exchange Server

To run an online backup of a Microsoft Exchange database, use any of the following methods:

Backup Methods

- Schedule the backup of an existing Microsoft Exchange backup specification using the Data Protector Scheduler.
- Start an interactive backup of an existing Microsoft Exchange backup specification using the Data Protector GUI or the Data Protector command-line interface.

For information on starting an interactive backup using the Data Protector command-line interface, refer to the `omnib` man page.

Scheduling a Backup

Scheduling a backup specification means setting time, date, and type of a backup that starts unattended once the scheduling options are defined and saved in the backup specification.

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

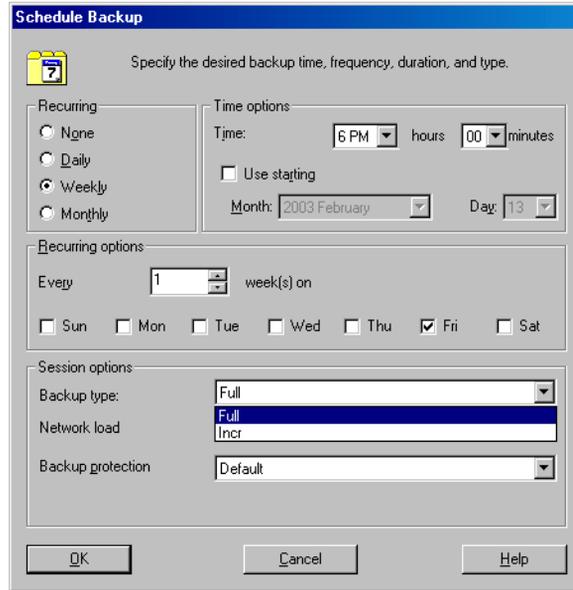
To schedule a Microsoft Exchange backup, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then Filesystem.
3. Double-click the backup specification you want to schedule and click the Schedule tab.
4. In the Schedule page, select a date in the calendar and click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options. See Figure 2-8.

The incremental backup performs a backup of the transaction log files that record changes to the database.

Microsoft Exchange Server automatically deletes transaction log files after they have been backed up.

Figure 2-8 **Scheduling Backups**



Click OK and then Apply to save the changes.

Running an Interactive Backup

An interactive backup can be performed any time after a backup specification has been created and saved.

To start an interactive backup of a Microsoft Exchange database, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then Filesystem.

Backing Up Microsoft Exchange Server

3. Right-click the backup specification, and then select `Start Backup`.

In the `Start Backup` dialog box, select the `Backup type` and `Network load`. For information on these options, click `Help`.

Click `OK`.

Restoring a Microsoft Exchange Server Database

You can restore a Microsoft Exchange Server database using the Data Protector GUI or using the Data Protector CLI.

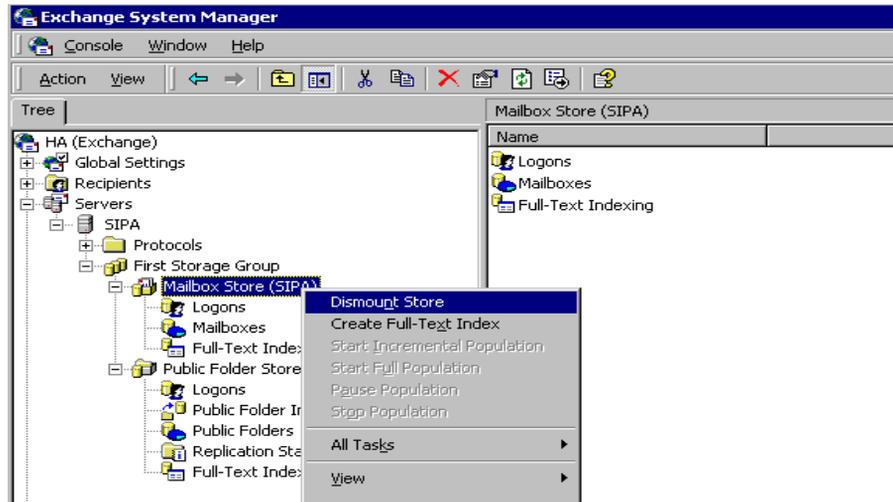
IMPORTANT

The database (store) must be dismounted to perform a restore.

To unmount the database (store), perform the following steps using the Exchange Administration GUI:

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

Figure 2-9 Unmounting the Database (Store)



2. A warning message appears. Click Yes to continue unmounting. When the dismounting is completed, you may start a restore session.

After a hard recovery, databases can be mounted automatically. See Table 2-2 on page 79 for details on restore options.

NOTE

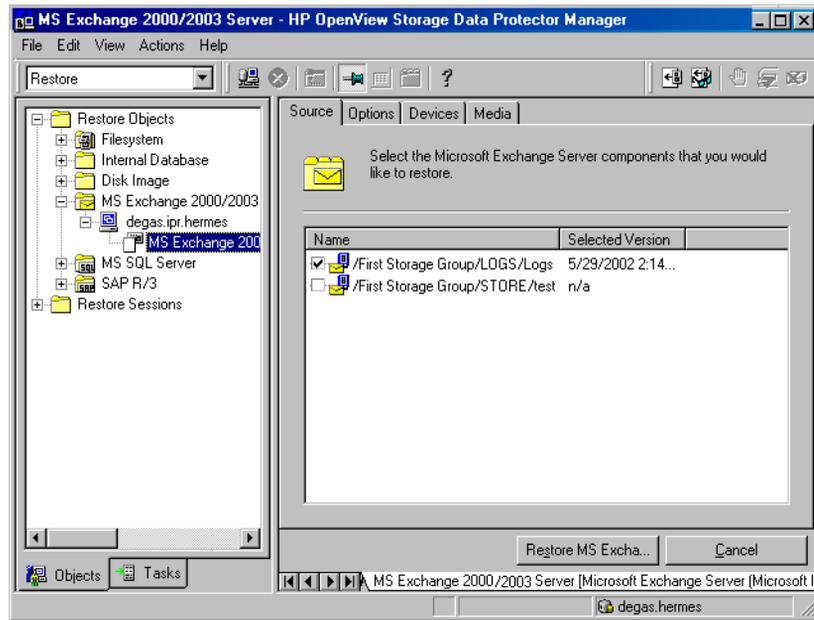
Log files for storage groups are saved in the subdirectory of the specified log directory. See Table 2-2 on page 79 for details on restore options.

Restore Using the GUI

Restore Procedure Use the following procedure to restore a Microsoft Exchange Server database:

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand Restore Objects, MS Exchange 2000/2003 Server, and then the name of the client to which you want to restore.
3. Select the backed up Microsoft Exchange Server databases you want to restore.

Figure 2-10 Restore Objects

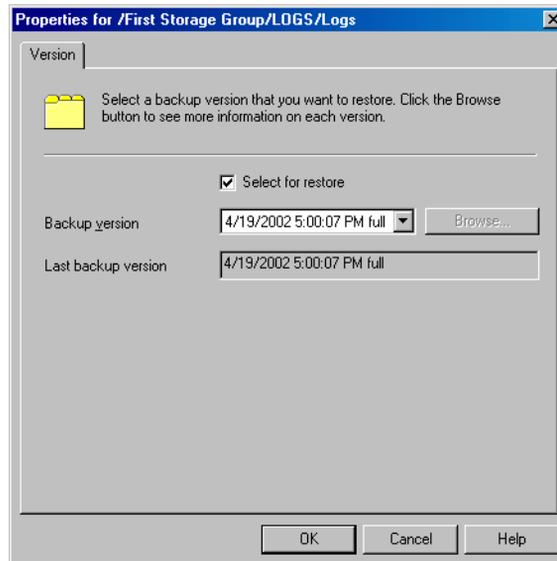


To select a backup version to be restored, right-click the object and select Properties.

IMPORTANT

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

Selecting a Backup Version



Note that restoring a database to a certain state often requires a multiphase restore. This means that multiple versions need to be restored to retrieve data. Because only transaction log files of storage groups are backed up during an incremental backup (without information on physical location of the storage groups), you have to restore the last full backup first and then all transaction log backups made after the last full backup,

IMPORTANT

When you restore from a full database (store) backup, make sure you selected the database files and the transaction log files 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, restore all the versions in *<italic>*: the first full and transaction log backup, the second transaction log backup, and the last transaction log backup. The last transaction log backup has to be restored with the Last restore set (start recovery) option selected.

4. In the Options property page, select the restore options. See Table 2-2 on page 79 for details about these options.
5. Click Devices and then Media to select the devices, verify device information, and set priorities of media to be used for the restore.

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

IMPORTANT

If the devices used for the restore are not those used for the backup, select the same number of devices in the Devices property page as you used when you backed up the object.

6. Click Restore MS Exchange 2000/2003 Server. Review your selection, and then click Finish to start the restore session.

If the Mount databases after recovery option was not specified for the restore, mount the dismounted Information Stores using the Exchange System Manager after the restore has finished.

Table 2-2

Microsoft Exchange Restore Options

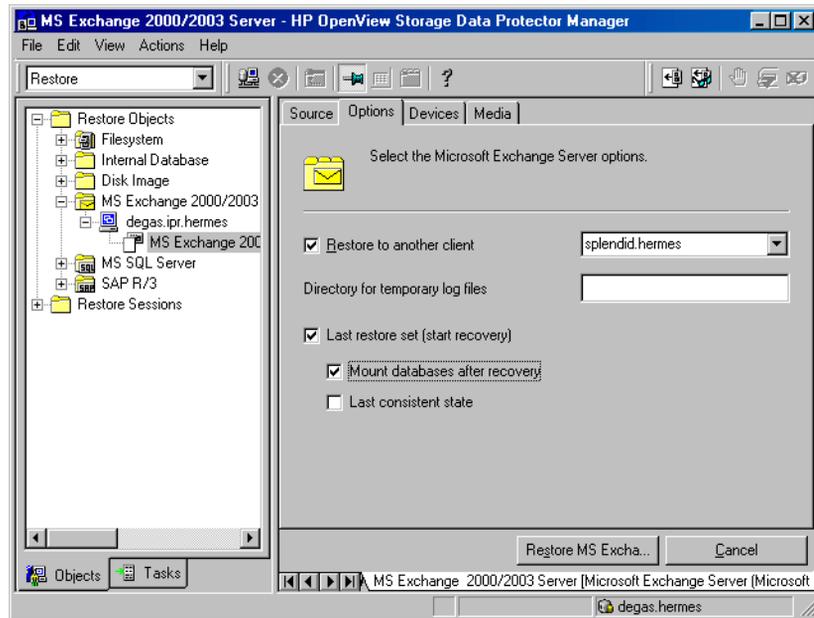
Restore to another client	By default, the target Data Protector Microsoft Exchange client is the Microsoft Exchange Server from which the application data was backed up. Nevertheless, the databases can be restored to a Microsoft Exchange Server other than the one the backup was made from. The new target Microsoft Exchange Server <i>must</i> be a part of the Data Protector cell and have the MS Exchange Integration software component installed.
----------------------------------	---

Table 2-2

Microsoft Exchange Restore Options

Directory for temporary log files	Specifying this option, you set the temporary directory for log files. Data Protector restores the log files to this directory. Using this directory, Microsoft Exchange Server then recovers the database - this operation is referred to as hard recovery .
Last restore set (start recovery)	If this option is set, a hard recovery is performed after the restore. Use this option if you are restoring the last set of files. If you do not set this option, you need to start the recovery manually by running the <code>eseutil /cc /t</code> utility from the appropriate subdirectory of the directory for temporary log files.
Mount databases after recovery	If you specify this option, the restored databases will be automatically mounted after the hard recovery.
Last consistent state	If this option is set, the database will be restored to its last consistent state. The latest log files, created after the backup, will be applied to the restored database during the recovery process.

Figure 2-11 Restore Options



Restoring the Database to Another Client

Follow the steps below to restore a database to another client:

1. Install the same version of Microsoft Exchange Server on a separate system.

NOTE

The new system name can be different.

2. On the newly installed Microsoft Exchange Server, install the same Microsoft Exchange Server Service Pack version(s).
3. On the newly installed Microsoft Exchange Server, create *all* the storage groups that existed on the Microsoft Exchange Server that was backed up. For *every* storage group, use the *same* name, the *same* location and the *same* parameters as on the Microsoft Exchange Server that was backed up.

Restoring a Microsoft Exchange Server Database

4. For *every* newly created storage group, create *all* the Stores (databases) that existed in this particular storage group on the Microsoft Exchange Server that was backed up. When creating a Store (database), use the *same* name, the *same* location and the *same* parameters as used for this particular Store (database) on the Microsoft Exchange Server that was backed up.
5. Install the Data Protector Microsoft Exchange integration on this system.
6. Restore the last full backup of the Microsoft Exchange Server database. Follow the normal procedure for restore of Microsoft Exchange Server database using the Data Protector GUI and set the following options in the Options property page:
 - Select `Restore to another client` and specify the target client name.
 - Specify the directory for temporary log files on the target client, for example `c:\EsseRestore`.
 - Select `Last restore set (start recovery)` if you are restoring the last set of files (if you do not have any incremental backups of the last full backup).

Refer to Table 2-2 on page 79 for details on these restore options.

7. 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 the `Last restore set (start recovery)` option to initiate automatic hard recovery of the Microsoft Exchange Server database after the restore of the Microsoft Exchange Server databases. If you do not set this option, start the recovery manually by running the `eseutil /cc /t` utility from the directory for temporary log files.

See also Table 2-2 on page 79 for detailed information on Microsoft Exchange Server restore options.

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

Restore Using the CLI

Use the `omnir` command to restore a Microsoft Exchange Server database. Refer also to Table 2-2 on page 79 for additional description of the parameters, limitations and prerequisites. This is the syntax of the `omnir` command when used to restore a Microsoft Exchange Server database:

Syntax

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

Options

where:

- msese specifies a Microsoft Exchange Server database for restore
- barhost <ClientName> specifies the system where the Data Protector Microsoft Exchange client that was backed up is installed
- destination <ClientName> specifies the target client for restore
- appname <full_application_name> Specifies a Microsoft Exchange Information Store, Site Replication Service or Key Management Service for the restore. The name of the Store/Service (<full_application_name>) must be provided in double quotes as follows:
 - for the Information Store: Microsoft Exchange Server (Microsoft Information Store)
 - for the Site Replication Service: Microsoft Exchange Server (Microsoft Site Replication Service)
 - for the Key Management Service: Microsoft Exchange Server (Microsoft Key Management Service)
- base <DBName> specifies the Microsoft Exchange store or logs for restore

Integrating Microsoft Exchange Server and Data Protector

Restoring a Microsoft Exchange Server Database

`-session <SessionID>` Specifies the session to be used for restore. Provide the *SessionID* of the backup session. In case of object copies, do not use the copy session ID, but the object's backup ID, which equals the object's backup session ID.

This option must be set for every `-base` option specified.

`-logpath <Path>` By specifying this option, you set the temporary directory for the Microsoft Exchange log files. Data Protector restores the log files to this directory. Using this directory, the Microsoft Exchange then recovers the database - this operation is referred to as hard recovery.

`-mount` the restored Microsoft Exchange databases will be automatically mounted after the soft or hard recovery

`-last` Hard recovery is performed after the restore of the Microsoft Exchange Server databases. Use this option if you are restoring the last set of files. If you do not set this option, start the recovery manually by running the `eseutil /cc /t` utility from the directory for temporary log files. If this option is not specified, soft recovery is performed after the restore.

Example

The Microsoft Information Store with the `/First Storage Group/STORE/Public Folder Store` store and `/First Storage Group/LOGS/Logs` logs is to be restored to the system called `computer.company.com` (where it was backed up), using the Data Protector session with the session ID `2003/07/07-13`. The Microsoft Exchange log files are to be restored to `c:\temp` directory, the hard recovery is to be performed after the restore has finished. The database is to be mounted after the hard recovery. Run the following command:

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

Refer to the `omnir` man page for more information on usage of the command.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector Microsoft Exchange integration. Start at “Backup Problems” on page 86 and if you cannot find a solution there, perform general checks and verifications.

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

Before You Begin

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

Checks and Verifications

- Try to run a backup and restore without using Data Protector. Use Windows Backup to back up and restore the Microsoft Exchange Server Information Store.
- 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
```
- If you perform an incremental backup, ensure that the Enable circular logging option on the Microsoft Exchange Server is disabled. You can check this option by starting the Exchange System manager and selecting Properties from the storage group you are backing up.
- Check if the environment variable Path also includes the `<Exchange_home>\bin` directory. For instructions, see “Prerequisites and Limitations” on page 60.

Backup Problems

Backup Fails

1. Check if the Microsoft Exchange Server services are running. To perform any kind of backup (MailBox Store, Public Folder Store or both), the following services must be running:
 - ✓ Microsoft Exchange System Attendant
 - ✓ Microsoft Exchange Information Store
2. Check on the Microsoft Exchange System manager if all the stores that need to be backed up are mounted.
3. Check if the Cell Manager is correctly set on the Microsoft Exchange Server client by checking the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\Omni  
Back II\Site
```

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

4. Create a Microsoft Exchange Server 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 the backup devices. See the *HP OpenView Storage Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.
5. Check if the filesystem backup of the problematic client works. It is much easier to troubleshoot a filesystem backup.
6. Examine the errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft Exchange Server.
7. Examine the errors logged in the Windows Event log.
8. Try to restart the Microsoft Exchange Server and start the backup again.

Restore Problems

- Check if the following Microsoft Exchange Server services are running:

1. Microsoft Exchange System Attendant
 2. Microsoft Exchange Information Store
- Using the Exchange System Manager, check whether all the stores that need to be restored are dismounted.
 - If you cannot mount the storage after a successful restore, check if the LOGS storage on the same storage group has also been restored.
 - To restore, for example a Second Storage Group, the `<MS_Exchange_Server_home>\Second Storage Group` directory must exist on your drive.
 - Define a directory for temporary log files in the restore context. Check whether the directory you have specified exists. If it does not, create it or specify another existing directory.
 - Check if the Cell Manager is correctly set on the Microsoft Exchange Server client by checking the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\Omni  
Back II\Site
```

Its name and value should be `CellServer` and "`<Cell Manager hostname>`", respectively.

- To restore to another system, make sure that the Microsoft Exchange Server is installed on that system and has the same organization and site names as the restored server.
- Ensure that the filesystem restore of the problematic client works. It is much easier to troubleshoot a filesystem restore. See the *HP OpenView Storage Data Protector Troubleshooting Guide* for information on troubleshooting filesystem restores.
- Examine the errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft Exchange Server.
- Examine the errors logged in the Windows Event log.
- **The restore session fails.**

Message

```
[Critical]
```

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

Description

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

Action

At the command prompt, run the `eseutil` tool 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 the `<Storage_group_name>` subdirectory:

```
eseutil /cc /t
```

- If the storage group name consists of Unicode characters, proceed as follows:

1. One of the subdirectories in the temporary log file directory contains an empty file whose filename equals the name of the storage group you are restoring. Identify the subdirectory where the file is located. The subdirectory name conforms to the following template:

```
Storage Group <Number>
```

2. Run the following commands:

```
<Drive_letter>:
```

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

```
eseutil /cc /t
```

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 3-1

Exchange Single Mailbox Backup Types

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

You can restore Exchange items:

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

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

Integration Concepts

The main component of the Data Protector Exchange Single Mailbox integration is `mbx_bar.exe`, installed on the Exchange Server system, which channels communication between the Data Protector Session Manager, and, via the MAPI interface, the Exchange Server. Figure 3-1 shows the architecture of the Data Protector Exchange Single Mailbox integration.

Figure 3-1 Exchange Single Mailbox Integration Architecture

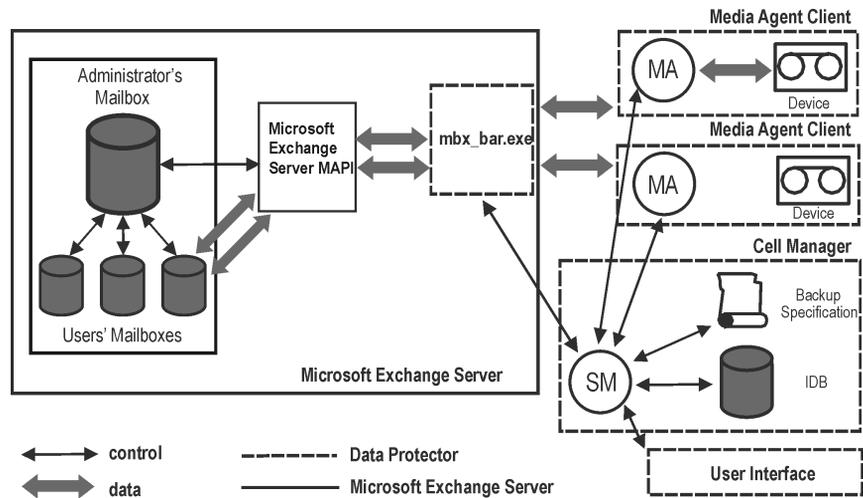


Table 3-2 Legend:

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

Table 3-2

Legend:

IDB	The Data Protector internal database.
-----	---------------------------------------

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

Configuring the Integration

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

Prerequisites

- Ensure that you have correctly installed and configured Exchange Server.
 - For supported versions, platforms, devices, and other information, see the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or <http://www.hp.com/support/manuals>.
 - For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HP OpenView Storage 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 2000/2003 Integration component installed.

Limitations

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

Before You Begin

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

Cluster-Aware Clients

Configure the integration on all cluster nodes.

Configuring Exchange Server Users

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

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

Configuring Exchange Servers

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

IMPORTANT

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

Prerequisites

- Ensure that the Exchange Server is online.

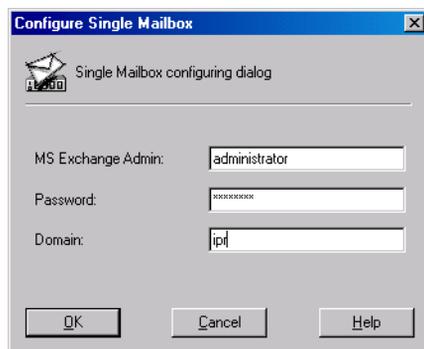
Configure the Exchange Server using the Data Protector Manager.

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

Click Next.

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

Figure 3-2 **Configuring the Exchange Server**



Click OK.

6. The Exchange Server is configured. Exit the GUI or proceed with creating the backup specification at step 6 on page 97.

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 3-3

Exchange Single Mailbox Backup Types

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

Limitations

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

IMPORTANT

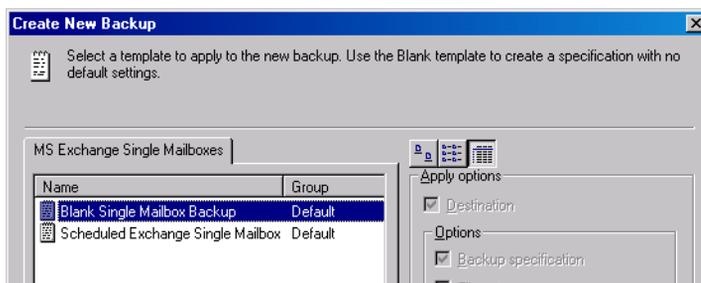
Do not use Data Protector Exchange Single Mailbox backups as a replacement for Data Protector Exchange Server backups. The latter are still needed to successfully recover a crashed system. For information, see “Backing Up Microsoft Exchange Server” on page 72.

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 3-3 **Selecting a Template**

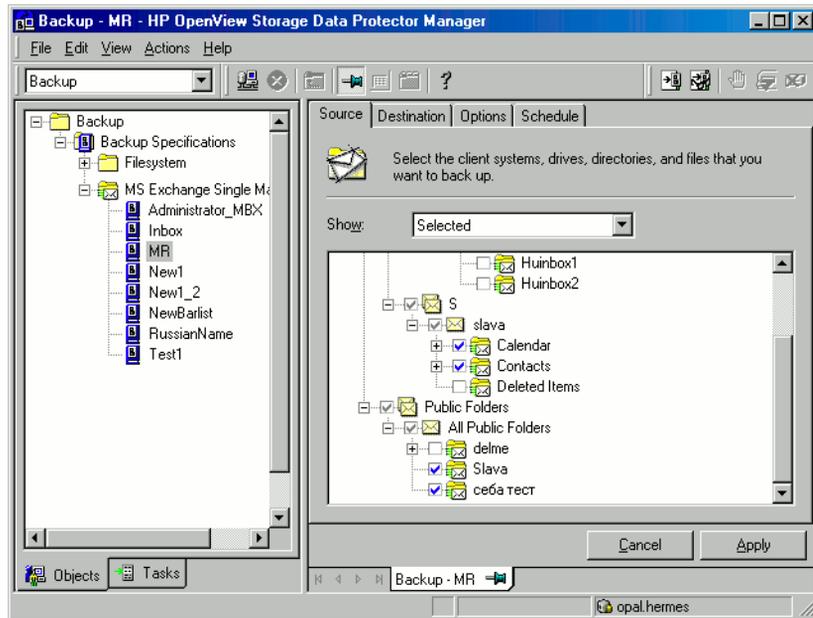


4. In Client, select the Exchange Server system. In a cluster environment, select the virtual server.
Click Next.
5. If the Exchange Server is not configured for use with Data Protector, the Configure Single Mailbox dialog box is displayed. Configure it as described in “Configuring Exchange Servers” on page 94.
6. Select the Exchange items you want to back up.
Mailboxes are organized alphabetically. For example, mailboxes starting with the letter S are collected under the S folder.
To back up all mailboxes and Public Folders, select the Exchange Server system at the top. Or you can browse for and select individual mailboxes and Public Folders or individual Exchange items from different mailboxes and Public Folders.

NOTE

Empty folders will not be backed up.

Figure 3-4 **Selecting Exchange Items for Backup**



Click Next.

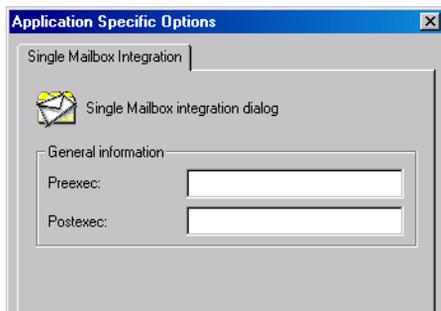
7. Select devices to use for the backup.

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

Click Next.

8. Set backup options. For information on application specific backup options (Figure 3-5), see Table 3-4 on page 100.

Figure 3-5 Exchange Single Mailbox Specific Backup Options



Click Next.

9. Optionally, schedule the backup. See “Scheduling Backup Specifications” on page 100.

Click Next.

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

TIP

Preview your backup specification before using it for real. See “Previewing Backup Sessions” on page 101.

Table 3-4 Exchange Single Mailbox Specific Backup Options

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

Modifying Backup Specifications

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

Scheduling Backup Specifications

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

Scheduling Example

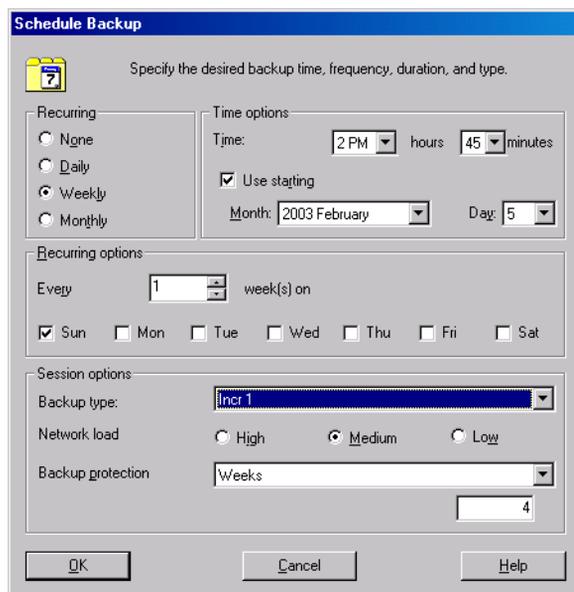
To perform `Incr1` backups of selected Exchange items at 2.45 p.m., 6.00 p.m., and 8.00 p.m. on Sundays:

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 `2 PM 45`. Under `Recurring Options`, select `Sun`. Under `Session Options`, select the `Incr1` backup type. See Figure 3-6.

Click `OK`.

3. Repeat steps 1 and 2 to schedule backups at 6 p.m. and 8 p.m.
4. Click `Apply` to save the changes.

Figure 3-6 Scheduling a Backup Specification



Previewing Backup Sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

Run:

```
omnib -mbx_list <backup_specification_name> -test_bar
```

What Happens During the Preview?

The following are tested:

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

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

Starting Backup Sessions

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

Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

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

Using the Data Protector CLI

On the Exchange Server system, run:

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

where <mailbox_mode> is one of the following:

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

For <list_options>, see the omnib man page.

Example

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

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

Restore

Restore Exchange items using the Data Protector GUI or CLI.

Before You Begin

- ✓ If you intend to restore Exchange items to another mailbox, ensure that the destination mailbox exists on the destination Exchange Server.
- ✓ If you intend to restore Exchange items to another Exchange Server system, ensure that the destination Exchange Server system has the MS Exchange 2000/2003 Integration component installed and that the Exchange Server is configured for use with Data Protector.

Restoring Using the Data Protector GUI

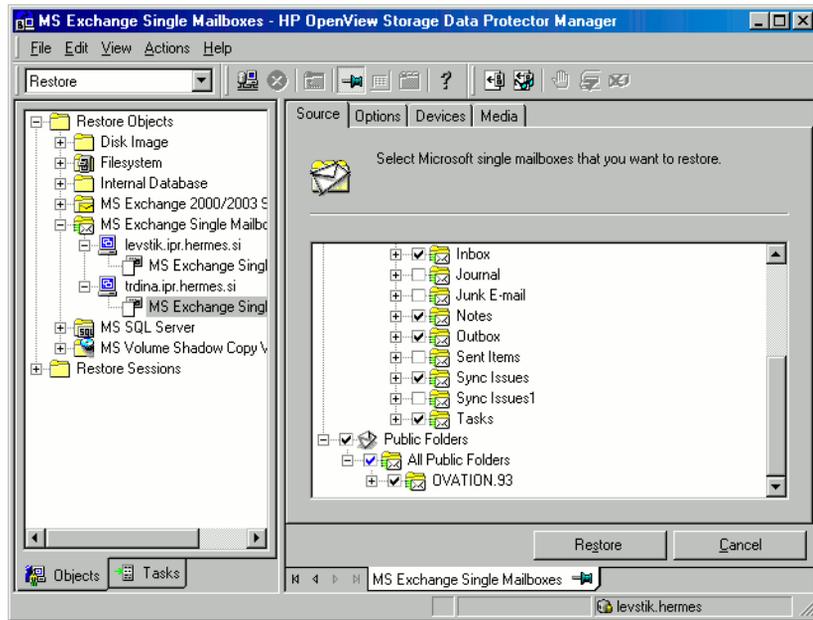
1. In the Context List, click `Restore`.
2. In the Scoping Pane, expand `MS Exchange Single Mailboxes`, the client from which the data to be restored was backed up, and then click `MS Exchange Single Mailboxes`.
3. In the Source page, browse for and select Exchange items to restore.

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

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

See Figure 3-7.

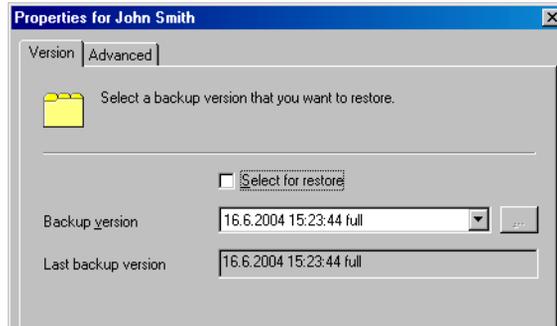
Figure 3-7 **Selecting Exchange Items for Restore**



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

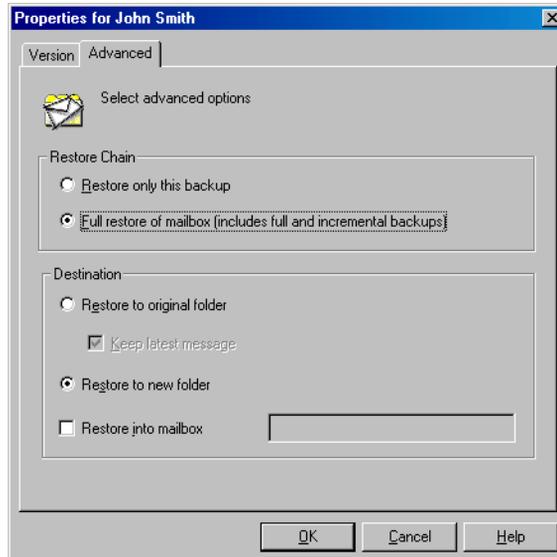
By default, the last backup session is used for restore. To restore from another session, right-click the relevant mailbox or Public Folders, and click Properties. See Figure 3-8.

Figure 3-8 **Version Properties**



To specify the restore destination and the chain of backup sessions to be used, click the Advanced tab. See Figure 3-9.

Figure 3-9 **Advanced Properties**



For details on these options, see Table 3-5 on page 108.

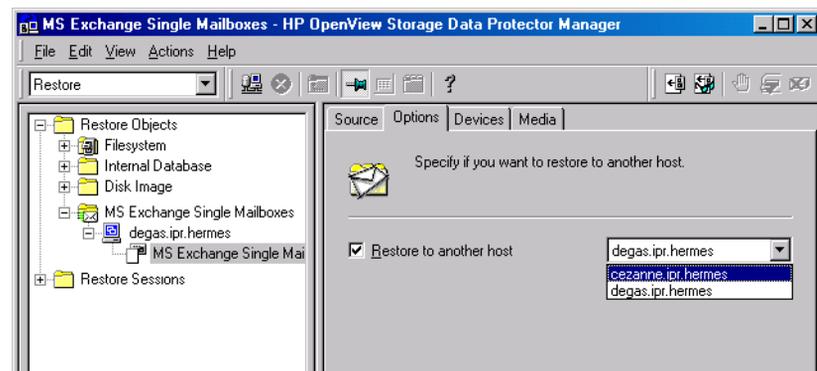
NOTE

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.

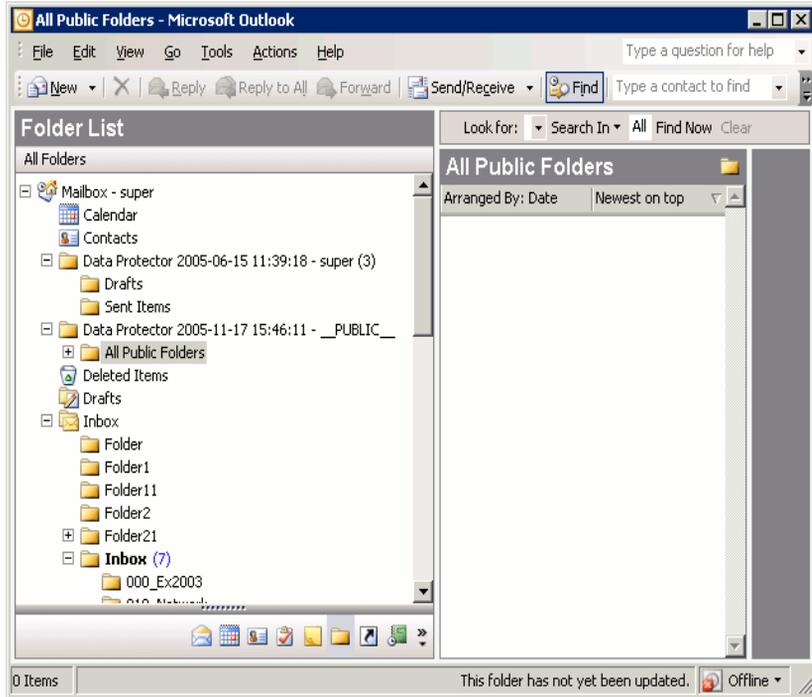
Figure 3-10**Selecting the Destination Exchange Server System**

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

Click Finish to start the restore.

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

Figure 3-11 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 3-5 Exchange Single Mailbox Restore Options

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

Table 3-5 Exchange Single Mailbox Restore Options

Option	Description
Full restore of mailbox	<p>Selected by default. Data is restored, not only from the selected backup session, but also from the latest full, the latest incremental1 (if it exists), and any incremental backups from the last incremental1 up to the selected backup version.</p> <p>Note that any Exchange item that was backed up in any of these sessions is displayed and can be selected for restore.</p>
Restore to original folder	<p>Data Protector restores Exchange items to the same location from which they were backed up.</p> <p>If <code>Keep latest message</code> is selected, existing messages in the destination mailbox or Public Folders are not restored even if they differ from their backed up versions.</p> <p>If <code>Keep latest message</code> is not selected, all messages are restored, replacing their current versions (if they exist). If different versions of the same message exist in the mailbox or Public Folders (for example, if you have a copy of the message), only one is replaced with the backed up version and all other versions remain intact.</p> <p>The messages in the mailbox that were not backed up in the specified backup session (or the restore chain of backup sessions) always remain intact.</p> <p>By default, this option is not selected.</p>

Table 3-5 Exchange Single Mailbox Restore Options

Option	Description
Restore to new folder	<p>Selected by default. Data Protector creates a new folder in the root of the mailbox (or in the root of All Public Folders) and restores Exchange items into it. See Figure 3-11 on page 108.</p> <p>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.</p> <p>If you restore from the same backup several times, a number is appended to the folder name. For example, in the second restore session of a mailbox, the folder Data Protector <backup_date> <backup_time> (1) is created.</p>
Restore into mailbox	<p>By default, Exchange items from a mailbox are restored to the original mailbox. Select this option to specify a different destination mailbox. Note that you can restore Exchange items from different mailboxes to the same mailbox.</p> <p>For privacy protection, you cannot restore Exchange items from mailboxes to Public Folders.</p>
Restore to another host	<p>By default, Exchange items are restored to the original Exchange Server system. Select this option, to specify a different destination Exchange Server system.</p>

Restoring Using the Data Protector CLI

On the Exchange Server system, run:

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

Select among the following <mailbox_options>:

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

Select among the following <Public_Folders_options>:

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

NOTE

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

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

Parameter Description

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

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

<session_ID> Backup version ID. For object copies, use the object's backup ID (which equals the object's backup session ID). Do not use the object's copy session ID.

<mailbox_name> 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".

<dest_mailbox_name> 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.

Restore

If the pathname contains a space, put the pathname in quotes. For example, "Inbox\My folder".

`<ex_folder>` 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 `2005/01/10-1` from the Exchange Server system `infinity.ipr.hermes`, to a new folder in the mailbox `TEMP` on the same Exchange Server system, run:

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

Example 2

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

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

Example 3

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

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

Example 4

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

```
omnir -mbx -barhost exchange.hp.com -public -session  
2005/06/10-19 -folder "All Public Folders\Administration"  
-folder "All Public Folders\Addresses" -mailbox "My Mailbox"  
-originalfolder -keep_msg
```

Monitoring Sessions

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

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

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

Performance Tuning

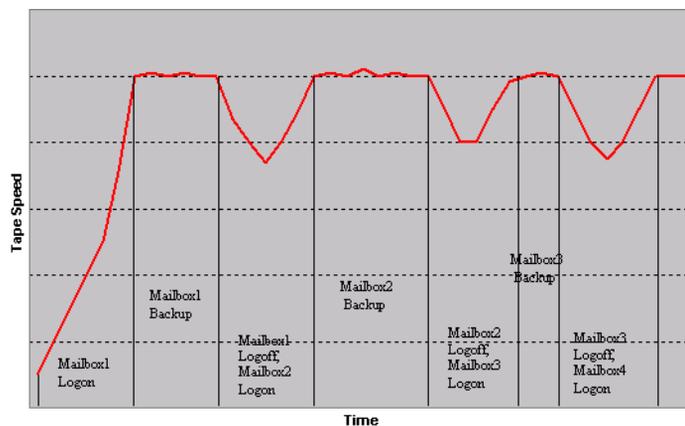
Performance tuning means customizing Exchange Server and Data Protector to achieve better backup and restore results.

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

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

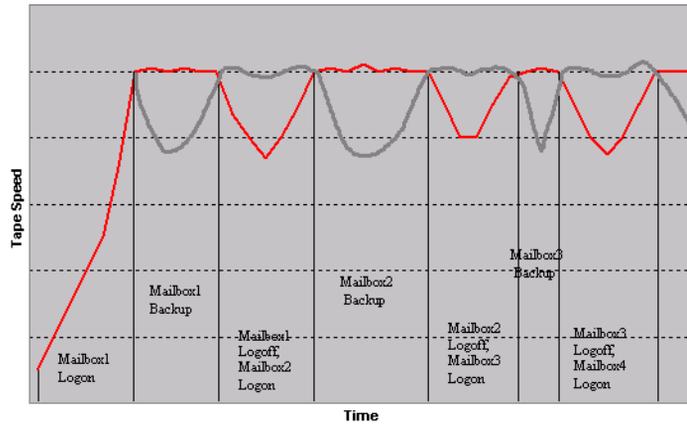
Figure 3-12

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 3-13 Example of Backup with Concurrency 2



NOTE

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

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector Exchange Single Mailbox integration. Start at “Problems” on page 118. If you cannot find a solution there, perform general checks and verifications.

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

Before You Begin

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

Checks and Verifications

If your configuration, backup, or restore failed:

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

Additionally, if your backup or restore failed:

- ✓ Ensure that the Cell Manager is correctly specified on the Exchange Server system: ensure that a value entry with the name CellServer and the value "<Cell Manager>" exists under the key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\Omni  
Back II\Site
```

- ✓ Examine errors logged in the Windows Event log.

Additionally, if your backup failed:

- ✓ Preview the Data Protector Exchange Single Mailbox backup.

If the Exchange Server part of the preview fails, ensure that the Exchange Server is online.

If the Data Protector part of the preview fails:

- Ensure that the Exchange Server is configured for use with Data Protector. See “Configuring Exchange Servers” on page 94.
- Create an Exchange Single Mailbox backup specification to back up to a null or file device.

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

Problems

Problem

You do not have permissions to log in to the system

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

```
Error = 596
```

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

or:

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

Action

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

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

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

If domain-level policy settings are defined:

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

If domain-level policy settings are not defined:

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

5. Restart the Data Protector Inet service.

Problem

Configuration of the Exchange Server fails

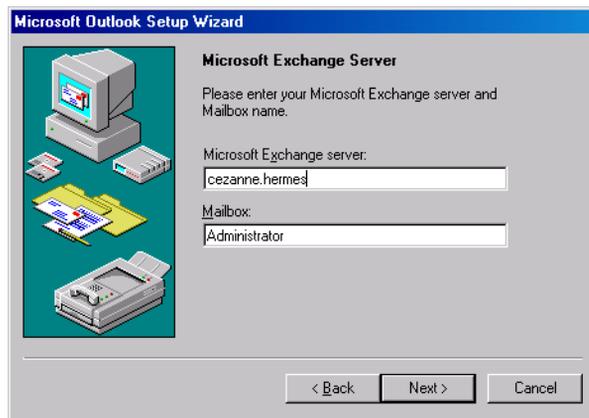
<Data_Protector_home>\log\debug.log on the Exchange Server system contains the following message:

An error has occurred while creating a profile administration object.

Action

1. Log in to the Exchange Server system.
2. Delete the incorrect administrator's profile:
`mbx_bar.exe delete`
3. Manually create a new profile:
`mbx_bar.exe create`
4. In the Choose Profile page, click New.
5. Follow the setup wizard. Type \$\$\$Data Protector for the profile name. Specify the Exchange Server system and the name of the Exchange Server administrator's mailbox. See Figure 3-14 on page 121.

Figure 3-14 Specifying the Exchange Server Administrator's Mailbox



Problem **Restore to another client fails**

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

Problem **Restore to another mailbox fails**

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

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft Volume Shadow Copy integration.

The chapter is organized into the following chapters:

“Introduction” on page 125

“Prerequisites and Limitations” on page 129

“Integration Concepts” on page 131

“Configuring the Integration” on page 136

“Writers Specifics” on page 139

“Backing Up Writers Data” on page 144

“Restoring Writers Data” on page 154

“Monitoring a VSS Backup and Restore” on page 165

“Troubleshooting” on page 168

“User Scenario For Microsoft Exchange Server 2003 Backup and Restore” on page 274

Introduction

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

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

Volume Shadow Copy Service

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

HP OpenView Storage Data Protector supports the integration with the Microsoft Volume Shadow Copy service (VSS).

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

Figure 4-1 and Figure 4-2 show the differences between the traditional backup model and the model with the Data Protector MS Volume Shadow Copy integration.

Figure 4-1

Actors of the Traditional Backup Model

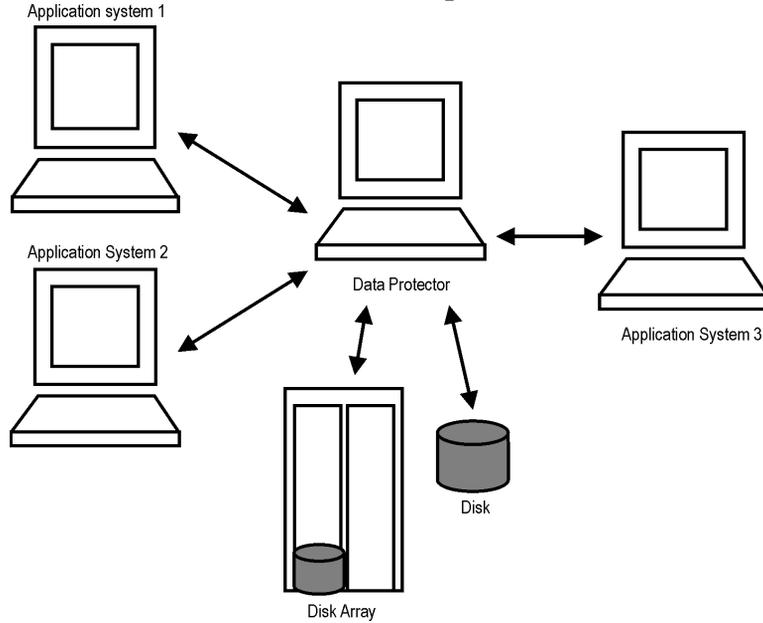
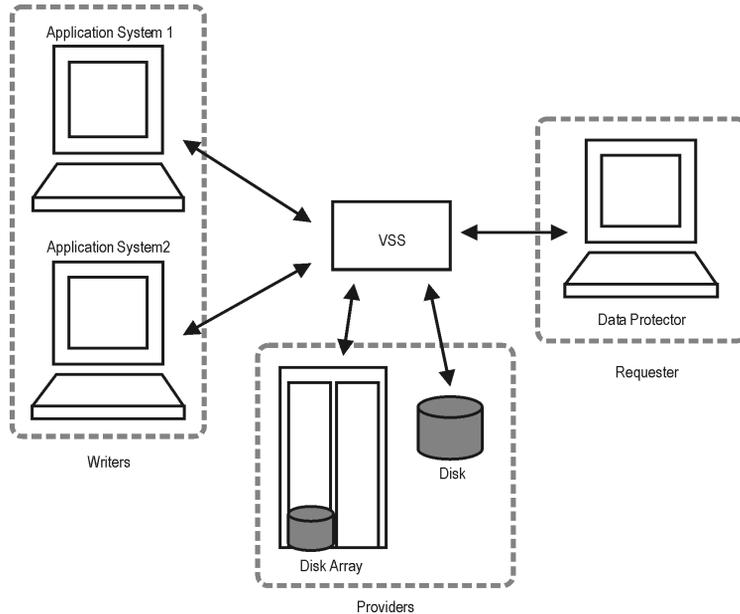


Figure 4-2

Actors of the Data Protector VSS Integration Backup Model



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

Advantages

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

- Central Management for all backup operations

The administrator can manage backup operations from a central point.

- Media Management

Data Protector has an advanced media management system that allows users to monitor media usage and set protection for stored data, as well as to organize and manage devices in media pools.

- Backup Management

Backed up data can be duplicated during or after the backup to increase fault tolerance of backups, to improve data security and availability, or for vaulting purposes.

- Scheduling

Data Protector has a built-in scheduler that allows the administrator to automate backups to run periodically. Using the Data Protector Scheduler, the backups you configure run unattended at specified times, as long as the devices and media are properly set.

- Device Support

Data Protector supports a wide range of devices: files, standalone drives, very large multiple drive libraries, etc.

- Reporting

Data Protector has reporting capabilities that allow you to get information on your backup environment. You can schedule reports to be issued at a specific time or to be attached to a predefined set of events, such as the end of a backup session or a mount request.

Introduction

- **Monitoring**

Data Protector has a feature that allows the administrator to monitor currently running sessions and view finished sessions from any system that has the Data Protector GUI installed.

All backup sessions are logged in the IDB, which provides the administrator with a history of activities that can be queried later.

Prerequisites and Limitations

This is a list of prerequisites and limitations for the Data Protector MS Volume Shadow Copy integration:

Prerequisites

- Before you begin, ensure that you have correctly installed and configured Data Protector, writers and shadow copy providers. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* on how to install Data Protector on various architectures and how to install the Data Protector MS Volume Shadow Copy integration.
 - Writers and shadow copy providers documentation for instructions on how to install and configure writers and providers on your system.

Limitations

See the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for a list of general Data Protector limitations. The integration-specific limitations are the following:

- Maximum 64 shadow copies per volume are allowed. The number of shadow copies per volume is limited by system resources.
- To run a VSS integration backup, the writer's data must be on an NTFS filesystem.
- The VSS integration backup of writers which store their data on network shared volumes is not supported.
- The Data Protector MS VSS integration does not provide any restore method for writers requesting a custom restore method. These writers are by default not presented by Data Protector.

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

Prerequisites and Limitations

- Preview is only possible for VSS filesystem backup and restore sessions.

Integration Concepts

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

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

Benefits of Using the Integration

The advantages of using the Data Protector VSS integration are the following:

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

VSSBAR Agent

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

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

NOTE

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

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

- restoring writer's data from media

Backup

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

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

NOTE

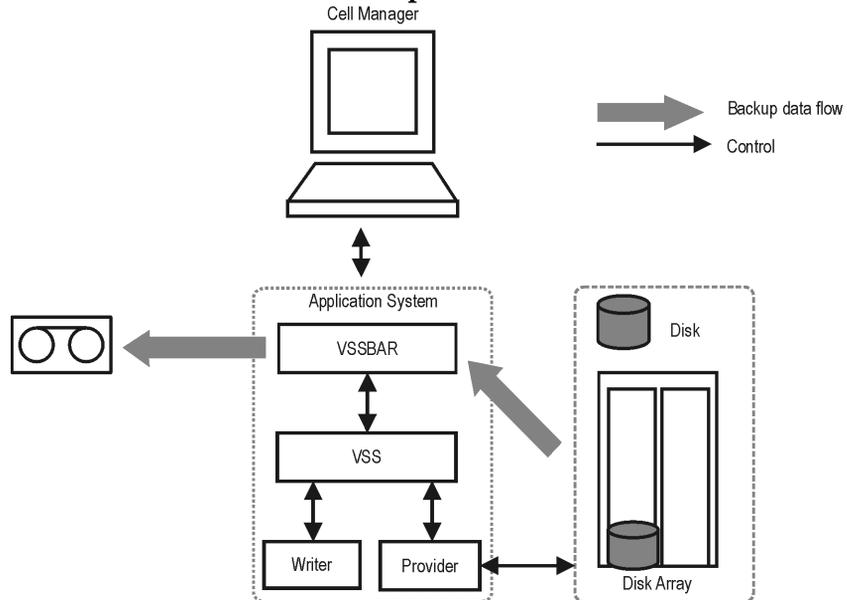
When the VSSBAR agent creates a shadow copy of the volume, this volume is marked in order to avoid attempts to simultaneously create another shadow copy of the same volume. In order to prevent any deadlocks arising from volume locking, only a single VSSBAR agent at a time is allowed to define a shadow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

6. When the writers are fully prepared for the consistent shadow copy backup, the VSSBAR agent passes shadow copy creation requests to VSS.

7. After a shadow copy is created, the VSS service returns the related information to the Data Protector..
8. Data Protector backs up the data from the shadow copy to media and then notifies the VSS service that the shadow copy can be released. VSS issues a command to the shadow copy provider to destroy the shadow copy that has been already backed up. Figure 4-3 shows the relations between the actors of a local or network VSS backup.

Figure 4-3

Local or Network VSS Backup



Restore

Data Protector offers two restore modes:

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

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

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

Restoring Components

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

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

Restoring Files

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

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

IMPORTANT

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

Configuring the Integration

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

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

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

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

Configuring the Data Protector VSS Cluster-Aware Integration

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

1. Configuring an VSS cluster-aware client. Refer to “Configuring a VSS Cluster-Aware Client” on page 136.
2. Configuring a cluster-aware VSS integration backup. Refer to “Configuring Backups for a Cluster-Aware VSS Client” on page 137

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

Configuring a VSS Cluster-Aware Client

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

Configuring Backups for a Cluster-Aware VSS Client

To configure backups for a cluster-aware VSS client, create a Data Protector VSS backup specification, as explained in “Creating Backup Specifications Using GUI” on page 144 taking into account the VSS infrastructure specifics described below.

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

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

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

Example 4-1

VSS Cluster Specifics

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

You have `node_A` and `node_B`, and MS Exchange Server 2003 running on a virtual host `exchsvr`. When creating a backup specification, you can select, among others, MS Exchange Writer and Event Log Writer. Suppose, at the time of a backup, Exchange is running on `node_A`. If you create just one backup specification for both writers, the following problems will occur:

- If you select `node_A` as your source host, you have Event Log Writer and Exchange Server 2003 associated with `node_A`. While it is true, that Event Log Writer is a property of the physical node, it is wrong to associate Exchange Server with it, as it is a property of the virtual server.

Suppose that after a failover, MS Exchange Server 2003 is running on `node_B`. When you try to restore the data to `node_A`, the restore will fail because Exchange disks are now owned by `node_B` and you cannot write to them. However, the restore of Event Log Writer will succeed.

Configuring the Integration

- If you select `exchsvr` as your source host, you have Event Log Writer and Exchange Server associated with a virtual server `exchsvr`. While it is true, that Exchange Server 2003 is a property of the virtual server, it is wrong to associate Event Log Writer with it, as it is a property of the physical node.

Suppose that after a failover, MS Exchange Server 2003 is running on `node_B`. When you try to restore Event Log Writer data to `exchsvr`, it will overwrite (or try to overwrite) the data in the Event Log of `node_B` with the data from the Event Log of `node_A`. The restore of Event Log Writer will fail, while Exchange Server will be restored successfully.

Writers Specifics

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

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

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

NOTE

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

Table 4-1 provides a description of VSS writers.

Table 4-1 **Writer description**

Writer Name	Description	Restore Method
Certificate Authority Writer	This is a system writer, used to back up and restore Certificate Authority (CA) Service database. This service issues, revokes, and manages certificates employed in public key-based cryptography technologies.	Files are restored after a reboot.

Table 4-1 **Writer description**

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

Table 4-1 Writer description

Writer Name	Description	Restore Method
FRS Writer	This VSS writer using a custom API, is used to back up and restore File Replication Service data. File Replication Service is a multithreaded replication engine that replicates system policies and logon scripts stored in System Volume (SYSVOL). FRS can also replicate data for Distributed File System (Dfs), copy and maintain shared files and folders on multiple servers simultaneously.	Custom restore method
IIS Metabase Writer	This is a system writer, used to back up and restore Microsoft Internet Information Server (IIS). IIS is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).	Files are restored after a reboot.
MSDE Writer	This is a writer used to back up and restore Microsoft SQL Server. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language.	Refer to “MSDE Writer Restore Specifics” on page 158.
Microsoft Data Protection Manager 2006 Writer	This is a writer used to back up and restore Microsoft Data Protection Manager 2006. Microsoft Data Protection Manager is a server that creates and stores replicas of clients and uses them for recovering the data on clients.	Refer to “Microsoft Data Protection Manager 2006 Writer Restore Specifics” on page 162.

Table 4-1 Writer description

Writer Name	Description	Restore Method
Microsoft Exchange Server 2003 Writer	This is a writer used to back up and restore Microsoft Exchange Server 2003. Microsoft Exchange Server 2003 is a mail and groupware server.	Refer to “Microsoft Exchange Server 2003 Writer Restore Specifics” on page 159.
NTDS Writer	This is a system writer used to back up and restore Microsoft Active Directory on Windows servers. Active Directory Service is a Windows server directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides methods for storing directory data and making this data available to network users and administrators.	To restore Active Directory, boot into Directory restore mode. Files will be restored if they can be overwritten.
Registry Writer	This VSS writer using a custom API, is used to back up and restore Windows Registry. Windows Registry is a database repository of information containing the Windows system configuration.	Custom restore method
Remote Storage Writer	This is a system writer used to back up and restore Remote Storage Service (RSS). RSS is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened.	Files are restored after a reboot.

Table 4-1 **Writer description**

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

Backing Up Writers Data

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

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

Configuration Steps

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

Creating Backup Specifications Using GUI

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

- For Microsoft Exchange Server 2003 specifics, see “Microsoft Exchange Server 2003 Writer Specifics” on page 147.
- For Microsoft Data Protection Manager 2006 specifics, see “Microsoft Data Protection Manager 2006 Writer Specifics” on page 149.

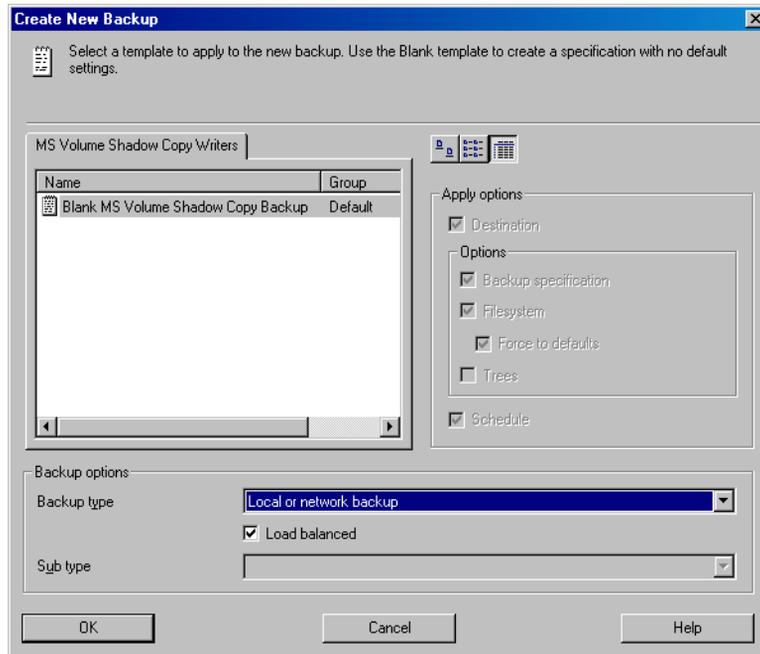
You need to configure different backup specifications for cluster-aware and non-cluster-aware writers. Refer to “Configuring the Data Protector VSS Cluster-Aware Integration” on page 136.

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

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications.
3. Right-click MS Volume Shadow Copy Writers and then click Add Backup. The Create New Backup dialog box is displayed.

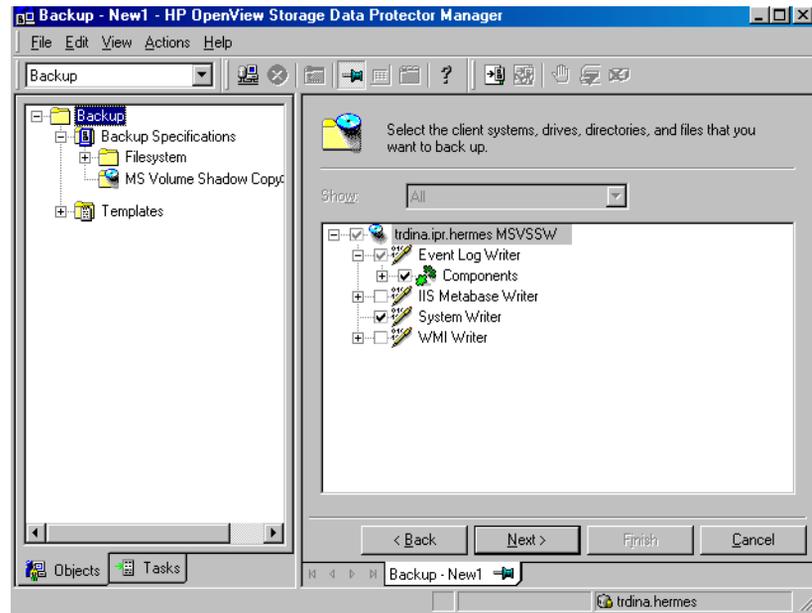
4. In the Create New Backup dialog box, click Blank MS Volume Shadow Copy Backup to select a template.

Figure 4-4 Selecting a Blank Template and Local or Network backup



- Select Local or network backup for the backup type.
5. Specify the name of the client that has the VSSBAR agent installed.
When backing up cluster-aware writers (such as SQL Server via the MSDE Writer), specify the name of the VSS client system as the virtual server name given in the particular writer resource group.
Click Next.
 6. Select the backup objects you want to back up.

Figure 4-5 Selecting Backup Objects



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

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

In case a writer requires all of its components to be backed up, lower-level items are disabled and you cannot select them. If you select such a writer for backup, all its components will be backed up.

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

7. Following the wizard, select the devices, backup options, and schedule your backup.

Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**.

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

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

TIP

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

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

8. Once you have defined all backup options and the schedule, you need to name and save the newly-created backup specification.

You have now completed the creation of a MS Volume Shadow Copy Writers backup specification.

9. You can review the newly-created and saved backup specification in the Backup context, under the specified group of backup specifications.
10. You can run backup using one of the following methods:
 - Schedule the backup of an existing MS Volume Shadow Copy Writers backup specification using the Data Protector Scheduler.
 - Start an interactive backup of an existing MS Volume Shadow Copy Writers backup specification.

Microsoft Exchange Server 2003 Writer Specifics

The Microsoft Exchange Server 2003 Writer supports the following Microsoft Exchange backup types:

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

Backing Up Writers Data

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

Limitations

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

Rollforward Recovery

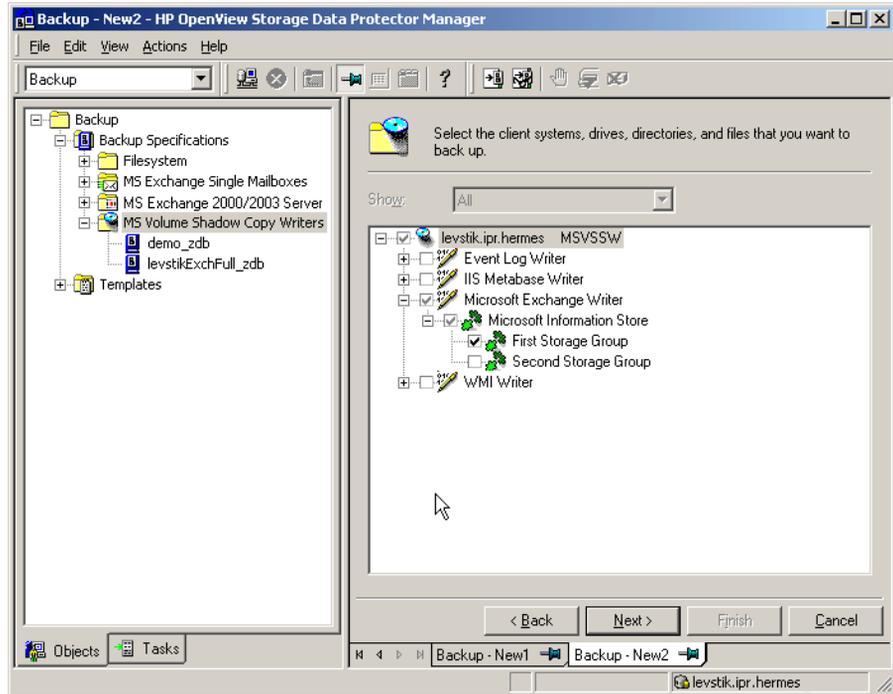
Transaction logs must be backed up to be able to perform the rollforward operation..

Consistency Check

The database can be successfully backed up only if the consistency check of the replicated datafiles succeeded.

To disable consistency checking, set the `OB2VSS_EXCHANGE_DISABLE_CONSISTENCY_CHECK` omnirc variable to 1.

Figure 4-6 **Selecting Microsoft Exchange Server 2003 Storage Groups**



Microsoft Data Protection Manager 2006 Writer Specifics

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

The Data Protection Manager writer is used to back up:

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

IMPORTANT

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

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

Two backup types are supported:

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

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

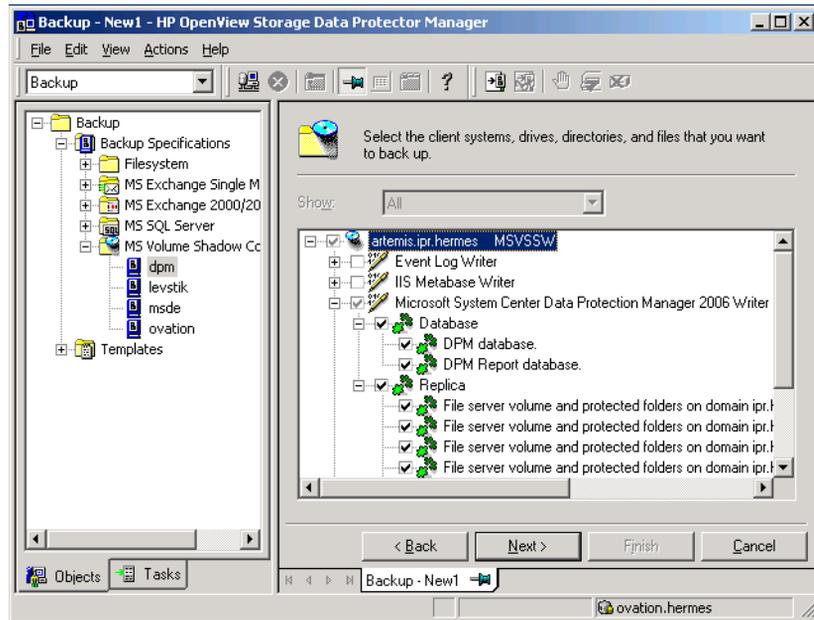
Prerequisite

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

Limitations

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

Figure 4-7 **Selecting Microsoft Data Protection Manager Database and Replicas**



Scheduling the Backup

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

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

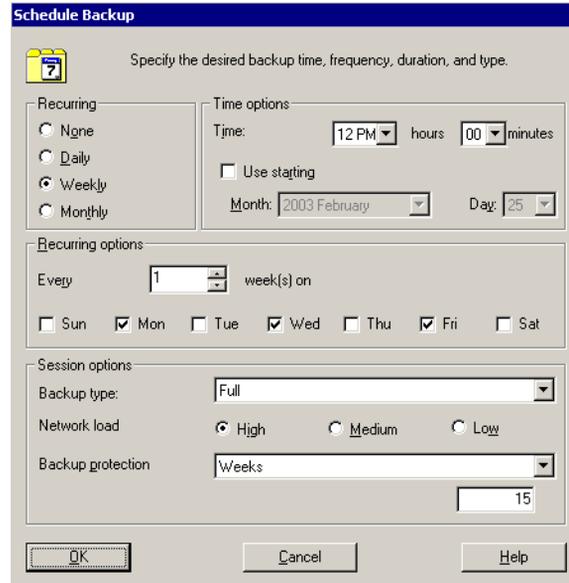
1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click MS Volume Shadow Copy Writers.

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

3. Double-click the backup specification you want to schedule and click the Schedule tab to open the Schedule property page.

4. In the Schedule property page, select a date in the calendar and click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options.

Figure 4-8 Scheduling a Backup



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

Running an Interactive Backup

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

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup; then expand the Backup Specifications and the MS Volume Shadow Copy Writers items.
3. Right-click the backup specification you want to run, and then select Start Backup from the pop-up menu.

The Start Backup dialog box appears.

Select the backup type and the network load {High|Medium|Low}.

Refer to online Help for a description of network load.

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

Restoring Writers Data

You can restore the Data Protector MS Volume Shadow Copy integration objects using the Data Protector GUI.

NOTE

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

Limitations for Custom Restore

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

NOTE

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

Restore Procedure

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

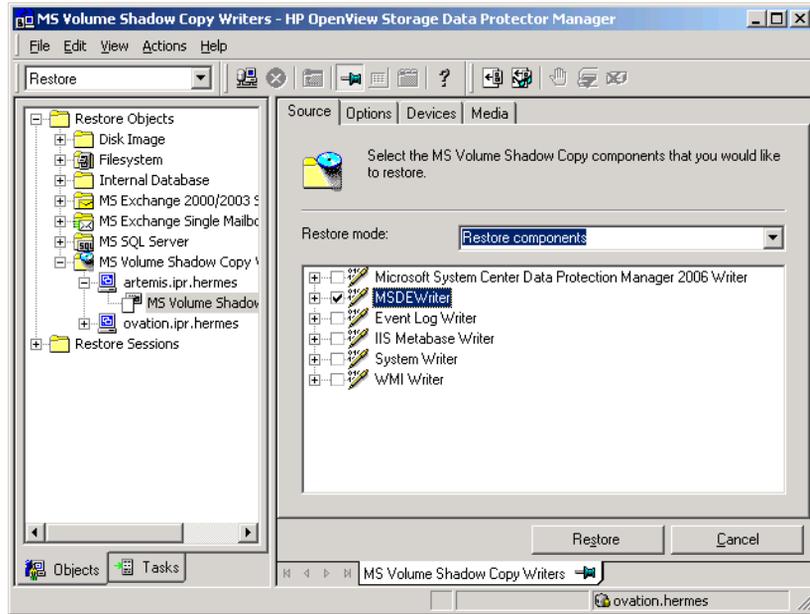
- For Microsoft Exchange Server 2003 Writer specifics see “Microsoft Exchange Server 2003 Writer Restore Specifics” on page 159.
- For MSDE Writer specifics see “MSDE Writer Restore Specifics” on page 158.

- For Microsoft Data Protection Manager 2006 Writer specifics see “Microsoft Data Protection Manager 2006 Writer Restore Specifics” on page 162.

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

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data. In the Results Area, a list of writers, which were backed up on this client, is displayed.
3. Select the restore mode:
 - To restore components using the Volume Shadow Copy Service, select Restore Components.
 - To restore individual files or a group of files without using the Volume Shadow Copy service, select Restore files.
4. In the Results Area, select the writers or writers’ components (for component restore) or files or a group of files (for file restore mode).

Figure 4-9 Restore Objects



You can select the top-level item (full writer restore) or only specific components. If you select a full writer restore, but some components of this writer were not backed up in the same session, the unavailable components are shaded and you cannot select them.

To select the version (the date of a backup), right-click the object name and click **Properties**. The last backup version is selected by default, however, you can select a different version from the drop-down list.

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. Refer to “Restore Options” on page 157.
6. In the **Devices** and **Media** property pages, the devices and media for restore are automatically selected.

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

7. Click the `Restore MS Volume...` button. Review your selection, and then click `Finish` to start a restore session.

The restore session messages are displayed in the Results Area.

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

Restore Options

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

Restore to another client

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

Restore into the following directory

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

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

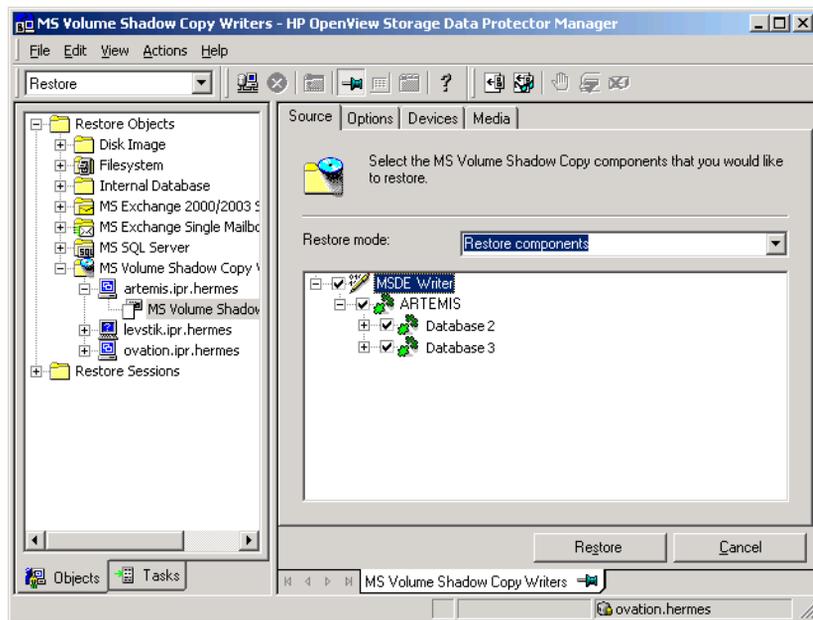
MSDE Writer Restore Specifics

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

IMPORTANT

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

Figure 4-10 MSDE writer



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

IMPORTANT

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

NOTE

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

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

Microsoft Exchange Server 2003 Writer Restore Specifics

Microsoft Exchange Server 2003 Writer is used to restore Microsoft Exchange Server 2003 database files.

When restoring from a Microsoft Exchange 2003 backup, the following two scenarios are possible:

- One or more databases are corrupted, but the log files are not damaged. In this case the database is restored and transaction logs are applied. See “Rollforward Recovery from the Loss of One or More Databases” on page 160.
- The log files are corrupted or missing. In this case all databases and log files need to be restored. A rollforward recovery of the database is not possible. See “Point-in-Time Restore After Loss of a Log File” on page 161.

Limitations

The following limitations apply when restoring Microsoft Exchange Server 2003:

- Shadow copies cannot be restored to alternate locations on the backup system.
- You cannot restore the shadow copy to the Recovery Storage Group.
- Rollforward recovery cannot be performed after a point-in-time restore.

Rollforward Recovery from the Loss of One or More Databases

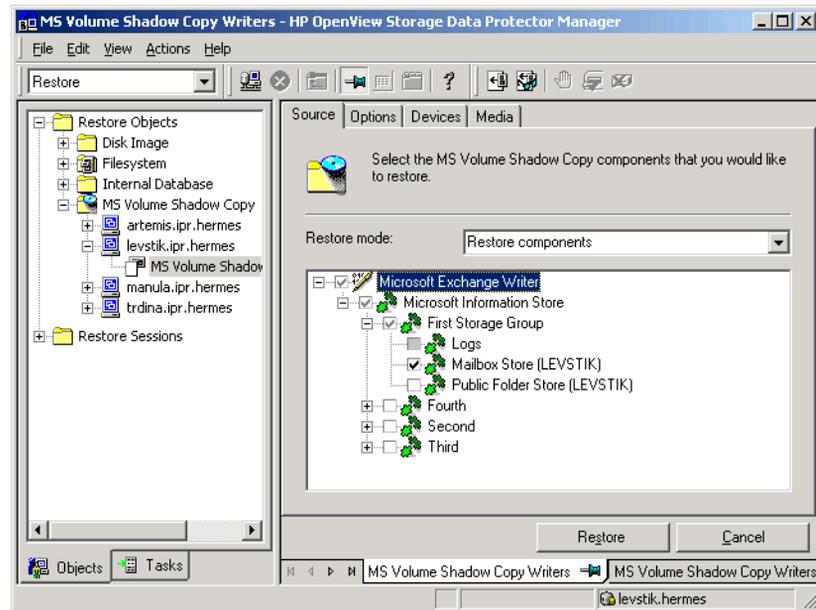
For a rollforward recovery:

1. Dismount all stores from the storage group in which the target store resides using Microsoft Exchange System Manager.
2. In the Data Protector GUI switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data

In the Results Area, expand the Microsoft Exchange Server 2003 writer and select the stores you want to recover. The Logs component is shaded and cannot be selected. You cannot select versions of individual stores, because a rollforward recovery is performed only to the current storage group state.

Figure 4-11

Selecting Microsoft Exchange Server 2003 Stores for Rollforward Recovery



3. Proceed as with general VSS writer restore. See “Restore Procedure” on page 154 for the general VSS writer restore procedure.

4. Mount all stores from the storage group in which they reside using Exchange System Manager. Selected stores are recovered.

Point-in-Time Restore After Loss of a Log File

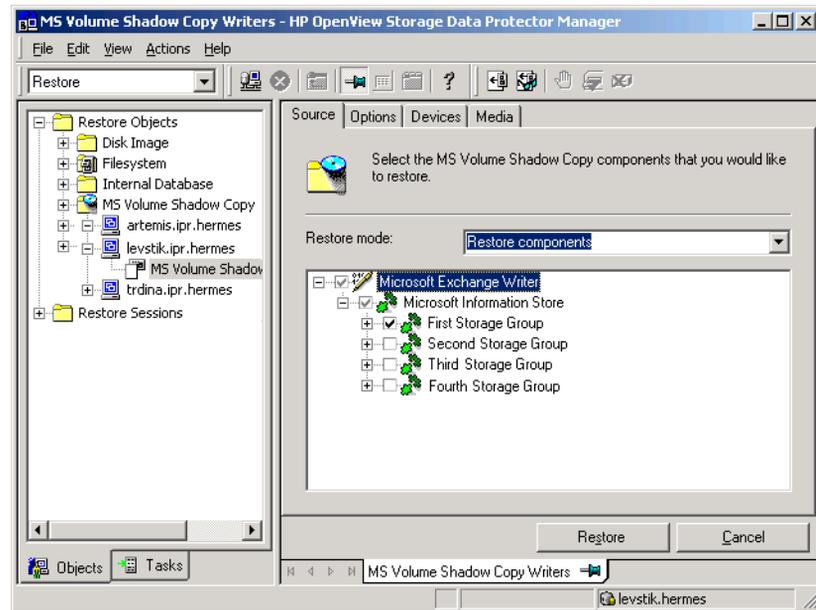
To perform a point-in-time restore:

1. Start Exchange System Manager and check if the storage group is already unmounted. If not, unmount the whole group.
2. Switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.

In the Results Area, expand the Microsoft Exchange 2003 writer and select the whole storage group. Do not select individual stores.

Figure 4-12

Selecting Microsoft Exchange Server 2003 Stores for Point-in-Time Restore



3. Proceed as with general VSS writer restore. See “Restore Procedure” on page 154 for the general VSS writer restore procedure.

Restoring Writers Data

4. Mount the stores from the storage group in which the target stores reside using Exchange System Manager. All stores are mounted and put in the state as they were at the last selected full, incremental, or differential backup.

Microsoft Data Protection Manager 2006 Writer Restore Specifics

When restoring the DPM writer, you can:

- Restore the DPM *server* first and then use the DPM to restore clients.
In case of a disaster, when the entire DPM server is lost, perform a standard disaster recovery procedure first and continue with restoring the DPM server. See “Restore the DPM Server First” on page 162
- Restore individual DPM *clients* directly, without using the DPM server (for example, if you cannot restore the DPM server or if you want to avoid the additional step of recreating the DPM snapshot). When restoring the DPM clients directly you can select between component restore and file restore modes. See “Restore the DPM Clients Directly” on page 164.

NOTE

Although the Data Protection Manager databases could also be restored using the MSDE writer, this method is not recommended, because DPM is *not* shut down automatically as with the DPM writer. If you really need to use this writer, shut down the DPM server manually.

Limitations

- Restore to another server is not supported by the Data Protection Manager writer.
- Parallel restore to different clients is not supported.

Restore the DPM Server First

1. Start the DPM administrator console and add disks to the storage pool so that you have enough free space to restore the replicas.
Ensure that the DPM writer (service) is started.

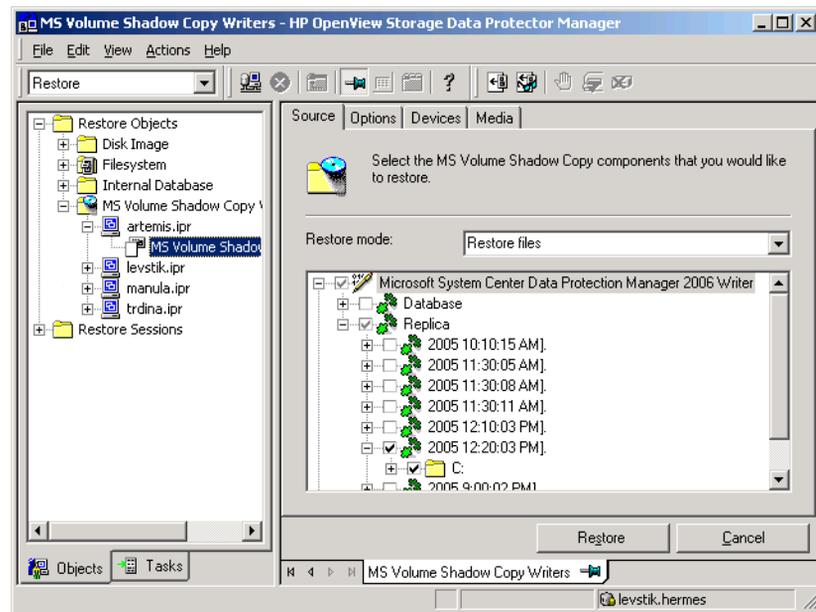
2. Switch to the Data Protector Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.
3. In the Results Area, expand the DPM writer and select *only* the Data Protection Manager databases.

Proceed as with general VSS writer restore. See “Restore Procedure” on page 154 for the general VSS writer restore procedure.

4. Run the DPM command `DpmSync -Sync` to reallocate replicas.
5. Switch back to the Data Protector Restore context, and select and restore the necessary *replicas*.

Figure 4-13

Restoring the Microsoft Data Protection Manager 2006 Client



6. Use the DPM to restore individual clients.

IMPORTANT

The DPM console does not automatically check for new or restored snapshots. Before you can start the restore of clients, you must use the Data Protection Manager to recreate a DPM snapshot.

Restoring Writers Data

- a. In the DPM console, open the Recovery context. Under the Browse tab, select the server, right click on the restored replica, and select Create shadow copy now.
- b. Select and restore the new snapshot to the client.

Restore the DPM Clients Directly

1. Switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.
2. Select the restore modes:
 - Restore Components
Use this mode *only* if the client to which you want to restore supports VSS, for example if you restore to Windows 2003 clients. You can restore only entire replicas.
 - Restore Files
The client does not need to support VSS and you can restore individual folders or files.
3. When selecting the DPM writer for restore, select *only* the Replica components. Do not select the DPM database.
4. Click the Options tab, and under Restore to another client enter the name of the target client. Click Next.
5. Proceed as with general VSS writer restore. See “Restore Procedure” on page 154 for the general VSS writer restore procedure.

Monitoring a VSS Backup and Restore

The Data Protector GUI enables you to monitor current or view previous backup and restore sessions.

Monitoring is automatically activated when you start a restore or a backup interactively.

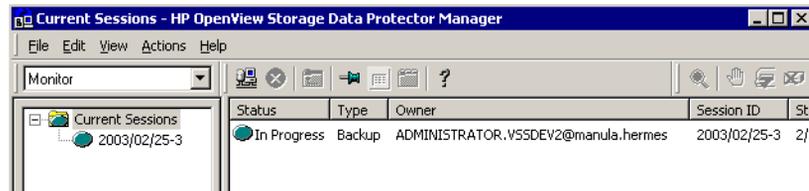
Monitoring Current Sessions

To monitor a currently running session using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Monitor**.
In the Results Area, all currently running sessions are listed. See Figure 4-14.
2. Double-click the session you want to monitor.

Figure 4-14

Monitoring a Current Session



Clearing Sessions To remove all completed or aborted sessions from the Results Area of the Monitor context, proceed as follows:

1. In the Scoping Pane, click **Current Sessions**.
2. In the Actions menu, select **Clear Sessions**. Or click the **Clear Sessions** icon on the toolbar.

To remove a particular completed or aborted session from the current sessions list, right-click the session and select **Remove From List**.

NOTE

All completed or aborted sessions are automatically removed from the Results Area of the Monitor context if you restart the Data Protector GUI.

For detailed information on a completed or aborted session, see “Viewing Previous Sessions”.

Viewing Previous Sessions

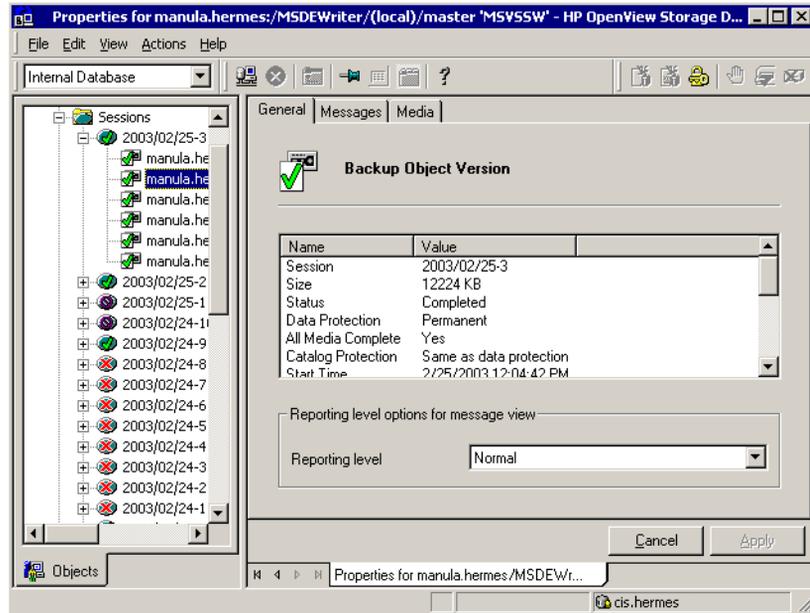
To view a previous session using the Data Protector GUI, proceed as follows:

1. In the Context List, click `Internal Database`.
2. In the Scoping Pane, expand `Sessions` to display all the sessions stored in the IDB.

The sessions are sorted by date. Each session is identified by a session ID consisting of a date in the YY/MM/DD format and a unique number.

3. Right-click the session and select `Properties` to view details on the session.
4. Click the `General`, `Messages` or `Media` tab to display general information on the session, session messages, or information on the media used for this session, respectively. See Figure 4-15.

Figure 4-15 Viewing a Previous Session



Troubleshooting

This section lists problems you might encounter when using the Data Protector Microsoft Volume Shadow Copy integration.

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

Before You Begin

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

Backup Problems

Problem

Microsoft Exchange Writer backup failed

When you start a backup of a Microsoft Exchange Server 2003 Writer, the following error is displayed:

```
[MAJOR]: Writer 'Microsoft Exchange Writer' failed to  
prepare files for backup.
```

The reason may be the failure of the previous backup and, due to a Microsoft Exchange Server 2003 issue, the Exchange Writer could not perform a proper cleanup.

Action

Restart the Information Store.

Restore Problems

Problem

After the restore of system writers was aborted, the Windows operating system is corrupted when you restart it.

If the restore of some system writers (for example, System Writer) is aborted for any reason (hardware or software failure, manually aborted, etc.), the Windows operating system may be corrupted after the restart (for example, the GUI or some system services cannot be started, etc.).

Action Depending on the nature of the corruption, repair or re-install the operating system from the Windows installation CD-ROM.

Problem **Some components are not restored during the restore session.**

If a component cannot be restored to the location specified in Writer Metadata Document (for example, if this location is locked or it is not possible to perform regular restore), this component will be skipped during the restore procedure.

Action Specify a location, where skipped files will be redirected in case of failure, by setting the OB2VSS_DUMPTO environmental variable in the `<Data_Protector_home>\omnirc` file. Restart the Data Protector services to apply the changes in the `omnirc` file.

Example If you want the files that are skipped during the restore to be copied to the `F:\Restore` directory, set `OB2VSS_DUMPTO=F:\Restore` in the `omnirc` file. In case the SQL component `Company` was skipped during the restore, it will be copied to the specified directory as follows:

```
F:\Restore\2002-12-09-23\G\SQL\Log\Company.ldf  
F:\Restore\2002-12-09-23\G\SQL\Log\Company.mdf
```

The pathname includes the backup session ID and the pathname to the original location.

Integrating Microsoft Volume Shadow Copy Service with Data Protector
Troubleshooting

Glossary

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.

AML (*EMASS/GRAU specific term*)

Automated Mixed-Media library.

application agent

A component needed on a client to back up or restore online database integrations.

See also Disk Agent.

application system (*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.

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.

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.

ASR Set

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

Glossary

(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 `<Data_Protector_home>\Config\Server\dr\asr` on a Windows Cell Manager or in `/etc/opt/omni/server/dr/asr/` on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger

See **library**

autoloader

See **library**

Automatic Storage Management

(Oracle specific term)

Automatic Storage Management is an Oracle 10g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.

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.

Glossary

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 entity or a disk image (rawdisk).

A backup object is defined by:

- **Client name:** hostname of the Data Protector client where the backup object resides.
- **Mount point:** the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- **Description:** uniquely defines backup objects with identical client name and mount point.
- **Type:** backup object type (for example filesystem or Oracle).

backup owner

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

backup session

A process that creates a copy of data on storage media. The activities are

specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.

*See also **incremental backup** and **full 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, 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

Glossary

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 target volumes of one or multiple application systems. The backup system is typically connected to a ZDB disk arraybackup device to perform the backup of the data in a replica.

See also **application system, target volume, and replica.**

backup types

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

backup 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 (*HP StorageWorks Disk Array XP specific term*)

The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system.

See also **HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.**

BC EVA (*HP StorageWorks EVA specific term*)

Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware.

See also **replica, source volume, snapshot, and CA+BC EVA.**

Glossary

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.

BC VA (*HP StorageWorks Virtual Array specific term*)

Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

See also HP StorageWorks Virtual Array LUN, application system, and backup system.

BCV (*EMC Symmetrix specific term*)

Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror.

The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also BC and BC Process.

Boolean operators

The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query consistency checkmanual disaster recovery is equivalent to consistencymanual AND checkdisaster 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 SAPDBA, BRBACKUP and BRRESTORE.

Glossary

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 **SAPDBA**, **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 **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM

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

CA (*HP StorageWorks Disk Array XP specific term*)

Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CA+BC EVA (*HP StorageWorks EVA specific term*)

The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.

See also **BC EVA**, **replica**, and **source volume**.

CAP (*StorageTek specific term*)

Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

Glossary

catalog protection

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

See also data protection.

CDB

The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.

See also MMDB.

CDF file (*UNIX specific term*)

A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell

A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. 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.

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'

Glossary

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 writers and filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or client system

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

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses ...).

CMD Script for Informix Server

(*Informix Server specific term*)

A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other

Glossary

Data Protector cells is highly recommended
See also MoM.

COM+ Registration Database

(Windows specific term)

The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface

A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, and restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*

The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView Storage Management Appliance, and is accessed by a Web browser.
See also HP StorageWorks EVA SMI-S Agent.

concurrency

See Disk Agent concurrency.

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 StorageWorks EVA specific term)*

A pair that consists of the source volumes on a local EVA and their replica on a remote EVA.

See also source volume, replica, and CA+BC EVA.

CRS

The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM

The Data Protector Copy and Consolidation Session Manager process

Glossary

controls the object copy and object consolidation sessions and runs on the Cell Manager system.

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 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 in the `<Data_Protector_home>\log\server\Ob2EventLog.txt` file on the Cell Manager. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

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.

data stream

Sequence of data transferred over the communication channel.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

database parallelism

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

Data Replication (DR) group (*HP StorageWorks EVA specific term*)

A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common CA EVA log.

See also **copy set**.

Glossary

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.

Dboject (*Informix Server specific term*)

An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory

The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the <Data_Protector_home>\db40 directory on a Windows Cell Manager and in the /var/opt/omni/server/db40 directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF

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

delta backup

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

device

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

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (*EMC Symmetrix specific term*)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to

Glossary

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 **incremental backup**.

differential backup (*MS SQL 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.

direct backup

A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.

See also **XCOPY engine**.

directory junction (*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.

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

Glossary

it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency

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

disk discovery

The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group (*Veritas Volume Manager specific term*)

The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap

images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging

The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format

A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. *See also* **virtual full backup**.

Glossary

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller

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

DR image

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

DR OS

A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See client backup with disk discovery.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)

See Symmetrix Agent (SYMA)

Glossary

emergency boot file (*Informix Server specific term*)

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

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 Logs

Files in which Windows logs all events, such as the starting or stopping of

services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger

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

exporting media

A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. *See also importing media.*

Extensible Storage Engine (ESE)

(*Microsoft Exchange Server specific term*)

A database technology used as a storage system for information exchange in Microsoft Exchange Server.

failover

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

Glossary

failover (*HP StorageWorks EVA specific term*)

An operation that reverses the roles of source and destination in CA+BC EVA configurations.

See also CA+BC EVA.

FC bridge

See **Fibre Channel bridge**

Fibre Channel

An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional 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 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.

Glossary

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 StorageWorks Disk Array XP specific term*)

HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.

See also **Primary Volume**, and **MU numbers**.

flash recovery area (*Oracle specific term*)

Flash recovery area is an Oracle 10g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and recovery (recovery files).

See also **recovery files**.

fnames.dat

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

formatting

A process that erases any data contained on a medium and prepares it for use with

Data Protector. Information about media (media 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 backed upstreamed to tape, even if there

Glossary

are no changes from the previous backup.

See also **incremental ZDB**.

global options file

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

group (*Microsoft Cluster Server specific term*)

A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI

A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

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: `/etc/opt/omni/server/Holidays` on the UNIX Cell Manager and `<Data_Protector_home>\Config\Server\holidays` on the Windows Cell Manager.

host backup

See **client backup with disk discovery**.

hosting system

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

Glossary

HP ITO

See OVO.

HP OpC

See OVO.

HP OpenView SMART Plug-In (SPI)

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

HP OVO

See OVO.

HP StorageWorks Disk Array XP LDEV

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.

See also BC (HP StorageWorks Disk Array XP specific term), CA (HP StorageWorks Disk Array XP specific term), and replica.

HP StorageWorks EVA SMI-S Agent

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

See also Command View (CV) EVA, and HP StorageWorks SMI-S EVA provider.

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.

See also HP StorageWorks EVA SMI-S Agent and Command View (CV) EVA.

Glossary

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.

See also **BC VA** and **replica**.

HP VPO

See **OVO**.

ICDA (*EMC Symmetrix specific term*)

EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

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

IDB recovery file

An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify

IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

importing media

A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.

See also **exporting media**.

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.

Glossary

incremental mailbox backup

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

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 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.

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

Glossary

two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. *See also* **Key Management Service** and **Site Replication Service**.

Informix Server (*Informix Server specific term*)

Refers to Informix Dynamic Server.

initializing

See **formatting**.

Installation Server

A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)

A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, 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/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**.

integrated security (*MS SQL specific term*)

Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

integration object

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

Internet Information Server (IIS)

(*Windows specific term*)

Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext

Glossary

Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address

Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL (*Sybase specific term*)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO

See **OVO**.

jukebox

See **library**.

jukebox device

A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

Key Management Service (*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.

LBO (*EMC Symmetrix specific term*)

A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library

Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation**

A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

Glossary

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.

lock name

You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script (*Informix Server UNIX specific term*)

A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to `<INFORMIXDIR>/etc/no_log.sh`.

logging level

The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless

Glossary

of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID (*MS 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

Glossary

and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet

See **Wake ONLAN**.

mailbox (*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 about 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 StorageWorks Disk Array XP specific term*)

An HP StorageWorks XP disk array that

contains the primary volumes for the Continuous Access configuration and acts as a master device.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)

See **Enterprise Cell Manager**.

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI (*Microsoft Exchange specific term*)

The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

Glossary

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 ID

A unique identifier assigned to a medium by Data Protector.

media label

A user-defined identifier used to describe a medium.

media location

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

media management session

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

media pool

A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set

The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type

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

media usage policy

The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

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

Glossary

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, writer**.

mirror (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

See **target volume**.

mirror rotation (*HP StorageWorks Disk Array XP specific term*)

See **replica set rotation**.

MMD

The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup

Glossary

environment, this part of the database can be common to all cells.
See also **CMMDB, CDB**.

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount request

A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point

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

MSM

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

MU number (*HP StorageWorks Disk Array XP specific term*)

A Mirror Unit number is an *integer*

number (0, 1 or 2), used to indicate a first level mirror.

See also **first level mirror**.

multi-drive server

A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat

See **IDB recovery file**.

OBDR capable device

A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object

See **backup object**

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session

A process that merges a restore chain of a backup object, consisting of a full

Glossary

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.

offline backup

A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.
- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

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

Glossary

standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log

See **archived redo log**

ON-Bar (*Informix Server specific term*)

A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

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

ONCONFIG (*Informix Server specific term*)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file in the

directory <INFORMIXDIR>\etc (on Windows) or <INFORMIXDIR>/etc/ (on UNIX).

online backup

A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.
- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

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**.

Glossary

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**.

OpC

See **OVO**.

OpenSSH

A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard (*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**.

OVO

HP OpenView Operations for Unix provides powerful capabilities for operations management of a large

Glossary

number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations. *See also* **merging**.

ownership

The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: `root.sys@<Cell Manager>`, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

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

`<Data_Protector_home>\Config\Server\dr\p1s` directory on a Windows Cell Manager or in `/etc/opt/omni/server/dr/p1s` directory on a UNIX Cell Manager 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 StorageWorks Disk Array XP specific term*)

A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- **COPY** - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- **PAIR** - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.

Glossary

- **SUSPENDED** - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism

The concept of reading multiple data streams from an online database.

physical device

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

post-exec

A backup option that executes a command or script after the backup of

an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
*See also **pre-exec**.*

pre- and post-exec commands

Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list

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

pre-exec

A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
*See also **post-exec**.*

Primary Volume (P-VOL) (*HP StorageWorks Disk Array XP specific term*)

Glossary

Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.

See also **Secondary Volume (S-VOL)**.

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

RAID

Redundant Array of Inexpensive Disks.

RAID Manager Library (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup

See **disk image backup**.

RCU (*HP StorageWorks specific term*)

The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS

Relational Database Management System.

Glossary

RDF1/RDF2 (*EMC Symmetrix specific term*)

A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS

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

Recovery Catalog (*Oracle specific term*)

A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

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

Recovery Catalog Database (*Oracle specific term*)

An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files (*Oracle specific term*)

Recovery files are Oracle 10g specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces.

See also **flash recovery area**.

RecoveryInfo

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) (*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.

Glossary

recycle

A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log (*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 (*HP*

StorageWorks Disk Array XP specific term)

The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database (*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/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 a virtual copy (for example, a snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object is replicated. *See also snapshot, snapshot creation, split mirror, and split mirror creation.*

Glossary

replica set (*ZDB specific term*)

A group of replicas, all created using the same backup specification.

See also **replica** and **replica set rotation**.

replica set rotation (*ZDB specific term*)

The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.

See also **replica** and **replica set**.

restore chain

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

restore session

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

RMAN (*Oracle specific term*)

See **Recovery Manager**.

RSM

The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM (*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.

SAPDBA (*SAP R/3 specific term*)

An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan

A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Glossary

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 StorageWorks Disk Array XP specific term*)

Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. *See also* **Primary Volume (P-VOL)**.

session

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

session ID

An identifier of a backup, restore, object copy, object consolidation, or media management session, consisting of the date when the session ran and a unique number.

session key

This environment variable for the Pre- 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 CLI commands.

shadow copy (*MS 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**.

shadow copy provider (*MS 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 (*MS VSS specific term*)

A collection of shadow copies created at the same point in time.

See also **shadow copy**.

shared disks

A Windows disk on another system that has been made available to other users

Glossary

on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF

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

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

The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.

See also **Information Store** and **Key Management Service**.

slot

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB

See **split mirror backup**.

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object

consolidation, and media management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.

See also **replica** and **snapshot creation**.

snapshot backup (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

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

snapshot creation (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for use. However background

Glossary

copying processes normally continue after creation.

See also **snapshot**.

source (R1) device (*EMC Symmetrix specific term*)

An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.

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

source volume (*ZDB specific term*)

A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are:
-A matrix in which some or much of the data contains zeros
-files from image applications
-high-speed databases
If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.
See also **replica** and **split mirror creation**.

split mirror backup (*EMC Symmetrix specific term*)

See **ZDB to tape**.

split mirror backup (*HP StorageWorks Disk Array XP specific term*)

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

split mirror creation (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.

See also **split mirror**.

split mirror restore (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.

See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

Glossary

sqlhosts file (*Informix Server specific term*)

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

SRD file

The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF (*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 (*HP StorageWorks Disk Array XP specific term*)

A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP

StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file

The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file

The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device

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

Glossary

standard security (*MS SQL specific term*)

Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.

See also **integrated security**.

Storage Group

(*Microsoft Exchange Server specific term*)

A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

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.

storage volume (*ZDB specific term*)

A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management

systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover

See **failover**

Sybase Backup Server API (*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.

Symmetrix Agent (SYMA) (*EMC Symmetrix specific term*)

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

synthetic backup

A backup solution that produces a synthetic full backup, an equivalent to a

Glossary

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 (sybssystemprocs)
- model database (model).

system disk

A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition

A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

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 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/

Glossary

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.

tablespace

A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup (*ZDB specific term*)

See ZDB to disk.

target database (*Oracle specific term*)

In RMAN, the target database is the database that you are backing up or restoring.

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 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 crashed 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 (*MS 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.

Glossary

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 logs (*Data Protector specific term*)

Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table (*Sybase specific term*)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (*MS VSS specific term*)

A shadow copy that is created on the

Glossary

application system and can be presented to the backup system which performs the backup.

See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file (*Novell NetWare specific term*)

A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation

See **lights-out operation**.

user account

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas

NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data

Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group

Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile (*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.

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

Glossary

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 StorageWorks EVA specific term)

The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.

*See also **Command View (CV) EVA**.*

Virtual Device Interface (MS SQL Server specific term)

This is a SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk (HP StorageWorks EVA specific term)

A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array

snapshot functionality.

*See also **source volume** and **target volume**.*

virtual full backup

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

virtual 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.

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

Glossary

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**.

VPO

See **OVO**.

VSS

See **Microsoft Volume Shadow Copy service**.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN

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

Web reporting

The Data Protector functionality that allows you to view reports on backup 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.

Glossary

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

(MS 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.

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).

XCopy engine *(direct backup specific term)*

A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through

XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also **direct backup**.

ZDB

See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*

A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, and instant recovery, and split mirror restore. *See also* **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.

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

ZDB to disk+tape *(ZDB specific term)*

A form of zero downtime backup where the replica produced is kept on the disk

Glossary

array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.

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

ZDB to tape (*ZDB specific term*)

A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.

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

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data

to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.

See also **ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.**

Glossary

A

- advantages
 - MS Exchange Server integration, 58
 - MS SQL Server integration, 3
 - VSS integration, 127, 131
- architecture
 - MS Exchange Server integration, 62
 - MS Exchange Single Mailbox integration, 91
 - MS SQL Server integration, 9
 - VSS integration, 126, 133

B

- backing up MS Exchange Server, 72–74
 - backup flow, 61
 - backup options, 70
 - backup specifications, creating, 66
 - full backups, 61
 - incremental backups, 61
 - scheduling backups, 72
 - starting backups, 73
- backing up MS Exchange Single Mailbox, 96–103
 - backup options, 100
 - backup specifications, creating, 96
 - backup specifications, modifying, 100
 - backup types, 90
 - full backups, 90
 - incremental backups, 90
 - performance tuning, 115
 - previewing backups, 101
 - scheduling backups, 100
 - scheduling backups, example, 100
 - starting backups, 102
 - starting backups, example, 103
- backing up MS SQL Server, 25–26
 - backup flow, 8
 - backup methods, 25
 - backup options, 21
 - backup specifications, creating, 18
 - backup types, 3, 6
 - concepts, parallelism, 9
 - differential backups, 6
 - full backups, 6
 - scheduling backups, 25
 - starting backups, 26
 - transaction logs backups, 7
- backing up VSS, 144–153
 - architecture, 126, 133

- backup flow, 132
- backup specifications, creating, 144
- scheduling backups, 151
- starting backups, 152
- backup
 - problems, 168
 - troubleshooting, 168
- backup flow
 - MS Exchange Server integration, 61
 - MS SQL Server integration, 8
 - VSS integration, 132
- backup methods
 - MS SQL Server integration, 25
- backup options
 - MS Exchange Server integration, 70
 - MS Exchange Single Mailbox integration, 100
 - MS SQL Server integration, 21
- backup specifications, creating
 - MS Exchange Server integration, 66
 - MS Exchange Single Mailbox integration, 96
 - MS SQL Server integration, 18
 - VSS integration, 144
- backup specifications, modifying
 - MS Exchange Single Mailbox integration, 100
- backup specifications, scheduling
 - MS Exchange Server integration, 72
 - MS Exchange Single Mailbox integration, 100
 - MS SQL Server integration, 25
 - VSS integration, 151
- backup types
 - MS Exchange Single Mailbox integration, 90
 - MS SQL Server integration, 3, 6

C

- checking configuration
 - MS Exchange Single Mailbox integration, 95
 - MS SQL Server integration, 17
- concepts
 - MS Exchange Server integration, 61–62
 - MS Exchange Single Mailbox integration, 91
 - MS SQL Server integration, 6–11
 - VSS integration, 131–135

Index

configuration files

- MS SQL Server integration, 12

- configuring MS Exchange Server, 64–71

- configuring MS Exchange Single Mailbox, 93–95

- checking configuration, 95

- configuring MS SQL Server, 14–24

- checking configuration, 17

- configuration files, 12

- configuring VSS, 136–138

- conventions, ix

- creating backup specifications

- MS Exchange Server integration, 66

- MS Exchange Single Mailbox integration, 96

- MS SQL Server integration, 18

- VSS integration, 144

D

- differential backups

- MS SQL Server integration, 6

- disaster recovery

- MS SQL Server integration, 35

E

- examples, MS Exchange Server integration

- restoring using CLI, 84

- examples, MS Exchange Single Mailbox integration

- restoring, 112

- scheduling backups, 100

- starting backups, 103

- examples, MS SQL Server integration

- restoring using CLI, 34

F

- full backups

- MS Exchange Server integration, 61

- MS Exchange Single Mailbox integration, 90

- MS SQL Server integration, 6

I

- incremental backups

- MS Exchange Server integration, 61

- MS Exchange Single Mailbox integration, 90

- interactive backups

- MS Exchange Server integration, 73

- MS Exchange Single Mailbox integration, 102

- MS SQL Server integration, 26

- VSS integration, 152

- introduction

- MS Exchange Server integration, 57

- MS Exchange Single Mailbox integration, 90

- MS SQL Server integration, 3

- VSS integration, 125

L

- limitations

- MS Exchange Server integration, 60

- MS SQL Server integration, 5

- VSS integration, 129

M

- modifying backup specifications

- MS Exchange Single Mailbox integration, 100

- monitoring sessions

- MS Exchange Single Mailbox integration, 114

- MS SQL Server integration, 44

- VSS integration, 165

- MS Exchange Server backup, 72–74

- backup flow, 61

- backup options, 70

- backup specifications, creating, 66

- full backups, 61

- incremental backups, 61

- scheduling backups, 72

- starting backups, 73

- MS Exchange Server configuration, 64–71

- MS Exchange Server integration

- advantages, 58

- architecture, 62

- backup, 72–74

- concepts, 61–62

- configuration, 64–71

- introduction, 57

- limitations, 60

- prerequisites, 60

- restore, 75–84

- troubleshooting, 85–88

- MS Exchange Server restore, 75–84

- restore flow, 61

- restore options, 79, 83

- using CLI, 83
 - using GUI, 76
 - MS Exchange Server troubleshooting, 85–88
 - MS Exchange Single Mailbox backup, 96–103
 - backup options, 100
 - backup specifications, creating, 96
 - backup specifications, modifying, 100
 - backup types, 90
 - full backups, 90
 - incremental backups, 90
 - performance tuning, 115
 - previewing backups, 101
 - scheduling backups, 100
 - scheduling backups, example, 100
 - starting backups, 102
 - starting backups, example, 103
 - MS Exchange Single Mailbox configuration, 93–95
 - checking configuration, 95
 - MS Exchange Single Mailbox integration
 - architecture, 91
 - backup, 96–103
 - concepts, 91
 - configuration, 93–95
 - introduction, 90
 - monitoring sessions, 114
 - restore, 104–113
 - troubleshooting, 117–121
 - MS Exchange Single Mailbox restore, 104–113
 - examples, 112
 - restore options, 108
 - using CLI, 110
 - using GUI, 104
 - MS Exchange Single Mailbox troubleshooting, 117–121
 - MS SQL Server backup, 25–26
 - backup flow, 8
 - backup methods, 25
 - backup options, 21
 - backup specifications, creating, 18
 - backup types, 3, 6
 - concepts, parallelism, 9
 - differential backups, 6
 - full backups, 6
 - scheduling backups, 25
 - starting backups, 26
 - transaction logs backups, 7
 - MS SQL Server configuration, 14–24
 - checking configuration, 17
 - configuration files, 12
 - MS SQL Server integration
 - advantages, 3
 - architecture, 9
 - backup, 25–26
 - backup types, 6
 - concepts, 6–11
 - configuration, 14–24
 - disaster recovery, 35
 - introduction, 3
 - limitations, 5
 - monitoring sessions, 44
 - performance tuning, 39
 - prerequisites, 5
 - restore, 27–38
 - troubleshooting, 47–53
 - MS SQL Server restore, 27–38
 - disaster recovery, 35
 - restore flow, 8
 - restore options, 3, 31
 - using CLI, 34
 - using GUI, 27
 - MS SQL Server troubleshooting, 47–53
- ## P
- performance tuning
 - MS Exchange Single Mailbox integration, 115
 - MS SQL Server integration, 39
 - prerequisites
 - MS Exchange Server integration, 60
 - MS SQL Server integration, 5
 - VSS integration, 129
 - previewing backups
 - MS Exchange Single Mailbox integration, 101
 - problems
 - backup, 168
- ## R
- restore flow
 - MS Exchange Server integration, 61
 - MS SQL Server integration, 8
 - VSS integration, 133
 - restore options
 - MS Exchange Server integration, 79, 83
 - MS Exchange Single Mailbox integration, 108
 - MS SQL Server integration, 3, 31

- VSS integration, 157
- restoring MS Exchange Server, 75–84
 - restore flow, 61
 - restore options, 79, 83
 - using CLI, 83
 - using GUI, 76
- restoring MS Exchange Single Mailbox, 104–113
 - examples, 112
 - restore options, 108
 - using CLI, 110
 - using GUI, 104
- restoring MS SQL Server, 27–38
 - disaster recovery, 35
 - restore flow, 8
 - restore options, 3, 31
 - using CLI, 34
 - using GUI, 27
- restoring VSS, 154–164
 - restore flow, 133
 - restore options, 157
 - using GUI, 154
- running backups *See* starting backups

S

- scheduling backups
 - MS Exchange Server integration, 72
 - MS Exchange Single Mailbox integration, 100
 - MS SQL Server integration, 25
 - VSS integration, 151
- starting backups
 - MS Exchange Server integration, 73
 - MS Exchange Single Mailbox integration, 102
 - MS SQL Server integration, 26
 - VSS integration, 152

T

- transaction logs backups
 - MS SQL Server integration, 7
- troubleshooting
 - backup, 168
- troubleshooting MS Exchange Server, 85–88
- troubleshooting MS Exchange Single Mailbox, 117–121
- troubleshooting MS SQL Server, 47–53
- troubleshooting VSS, 168–169
- typographical conventions, ix

V

- Volume Shadow Copy service
 - VSS integration, 125
- VSS backup, 144–153
 - architecture, 126, 133
 - backup flow, 132
 - backup specifications, creating, 144
 - scheduling backups, 151
 - starting backups, 152
- VSS configuration, 136–138
- VSS integration
 - advantages, 127, 131
 - architecture, 126, 133
 - backup, 144–153
 - concepts, 131–135
 - configuration, 136–138
 - introduction, 125
 - limitations, 129
 - monitoring sessions, 165
 - prerequisites, 129
 - restore, 154–164
 - troubleshooting, 168–169
- VSS restore, 154–164
 - restore flow, 133
 - restore options, 157
 - using GUI, 154
- VSS troubleshooting, 168–169

W

- writers specifics, VSS integration, 139
 - MS Data Protection Manager 2006, 149, 162
 - MS Exchange Server 2003, 147, 159
 - MSDE, 158