HP Data Protector A.06.10

Integration guide for Sybase, Network Node Manager, Network Data Management Protocol, and VMware

[Build 500]

T B D

Part number: TBD
First edition: TBD Month 2008

invent

Legal and notice information

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Printed in the US

Contents

Figures

Tables

Publication history

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

Table 1 Edition history

Part number	Guide edition	Product
B6960-90111	October 2004	Data Protector Release A.05.50
B6960-96010	July 2006	Data Protector Release A.06.00
TBD	TBD	Data Protector Release A.06.10

About this guide

This guide describes how to configure and use Data Protector with Sybase, Network Node Manager, Network Data Management Protocol, and VMware.

Intended audience

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

- Basic Data Protector functionality
- Administration of the respective application

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

Documentation set

Other documents and online Help provide related information.

Guides

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

You can find these documents from the Manuals page of the HP Business Support Center website:

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

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

- *HP Data Protector concepts guide*

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

- *HP Data Protector installation and licensing guide*

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

- *HP Data Protector troubleshooting guide*

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

- *HP Data Protector disaster recovery guide*

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

- *HP Data Protector integration guides*

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

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

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

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

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

  - *HP Data Protector integration guide for IBM applications: Informix, DB2, and Lotus Notes/Domino*

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

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

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

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

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

- *HP Data Protector integration guide for HP Reporter*

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

- *HP Data Protector integration guide for HP Operations Manager for UNIX*

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

- *HP Data Protector integration guide for HP Operations Manager for Windows*

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

  There are two versions of the guide:

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

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

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

- *HP Data Protector zero downtime backup concepts guide*

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

- *HP Data Protector zero downtime backup administrator's guide*

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

- *HP Data Protector zero downtime backup integration guide*

  This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases. The guide also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

- *HP Data Protector MPE/iX system user guide*

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

- *HP Data Protector Media Operations user guide*

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

- *HP Data Protector product announcements, software notes, and references*

  This guide gives a description of new features of HP Data Protector A.06.10. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.hp.com/support/manuals

  There are also four other *Product announcements, software notes and references*, which serve a similar purpose for the following:

  - HP Operations Manager software software UNIX integration
  - HP Operations Manager software software Windows integration
  - HP Service Information Portal and HP Reporter software
  - HP Performance Manager software and HP Performance Agent integration
  - HP Media Operations

Online help

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

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

- ***Windows***: Unzip `DP_help.zip` and open `DP_help.chm`.
- ***UNIX***: Unpack the zipped tar file `DP_help.tar.gz`, and access the online help system through `DP_help.htm`.

Documentation map

Abbreviations

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

Abbreviation	Guide
CLI	Command line interface reference
Concepts	Concepts guide
DR	Disaster recovery guide
GS	Getting started guide
Help	Online Help
IG-IBM	Integration guide—IBM applications
IG-MS	Integration guide—Microsoft applications
IG-O/S	Integration guide—Oracle, SAP R/3, and SAP DB/MaxDB
IG-OMU	Integration guide—HP Operations Manager software software, UNIX
IG-OMW	Integration guide—HP Operations Manager software software, Windows
IG-PM/PA	Integration guide—Performance Manager and Performance Agent software
IG-Report	Integration guide—HP Reporter software
IG-SIP	Integration guide—HP Service Information Portal
IG-Var	Integration guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and licensing guide
MO GS	Media Operations getting started guide

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

Map

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

Integrations

Look in these guides for details of the following integrations:

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

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

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: Table 2 on page 22	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
text	Emphasized monospace text

△ CAUTION:

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

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Data Protector graphical user interface

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. You can use the original Data Protector GUI or the Data Protector Java GUI. Refer to the online Help for information about the Data Protector graphical user interface.

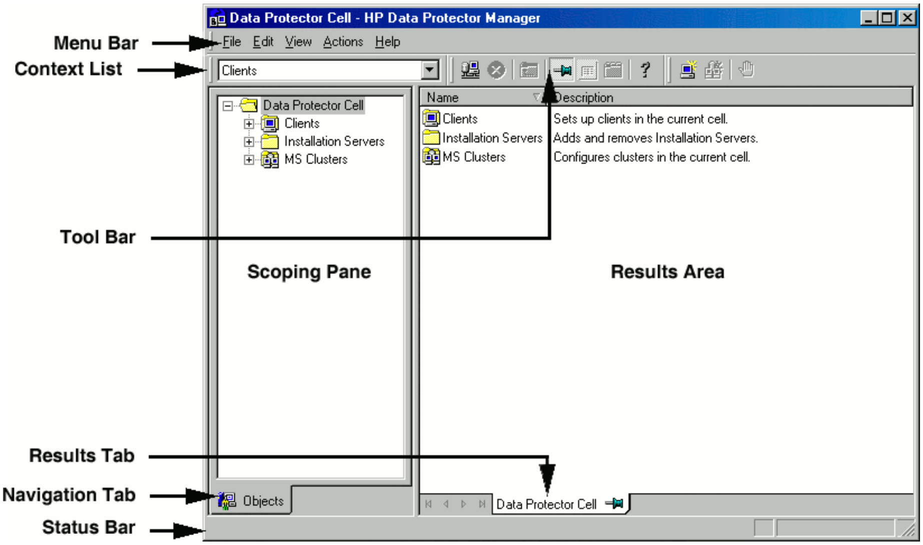

Name	Description
Clients	Sets up clients in the current cell.
Installation Servers	Adds and removes Installation Servers.
MS Clusters	Configures clusters in the current cell.

Figure 1 Data Protector graphical user interface

General Information

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

HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

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

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to AppRM.DocFeedback@hp.com. All submissions become the property of HP.

1 Integrating Sybase Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Sybase Adaptive Server (**Sybase Server**) integration. It describes concepts and methods you need to understand to back up and restore Sybase databases.

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

Table 3 Backup types

Full	Backs up all selected Sybase databases and transaction logs.
Trans	Backs up changes made to the transaction logs since the last backup of any type.

During backup, the database is online and actively used.

Sybase databases are restored using the `isql` utility. You can restore a database:

- To a specific point in time
- To a new database
- To another Sybase instance

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

Integration concepts

Data Protector integrates with Sybase Backup Server through the Data Protector Database Library based on a common library called Data Protector **BAR** (Backup And Restore). The Data Protector Database Library channels communication between the Data Protector Session Manager, and, via the **Sybase Backup Server API**, the

Sybase Server **isql** utility. Figure 2 on page 28 shows the architecture of the Data Protector Sybase integration.

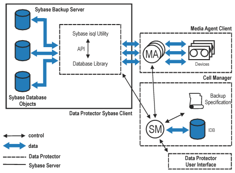

Figure 2 Sybase integration architecture

Table 4 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
API	Sybase Backup Server Application Programming Interface.
Database Library	A set of Data Protector executables that enable data transfer between the Sybase Backup Server and Data Protector.
MA	Data Protector General Media Agent.
Backup Specification	A list of objects to be backed up, backup devices, and options to be used.
IDB	The Data Protector Internal Database.

The `isql` utility sends backup and restore commands (issued through the Data Protector GUI or CLI, or the Sybase `isql` command line interface) to Sybase Backup Server, initiating data transfer between Sybase databases and Data Protector media.

While Sybase Backup Server is responsible for read/write operations to disk, Data Protector manages devices and media used for backup and restore.

Data Protector CLI commands

Run the Data Protector CLI commands from the following directories:

Windows: `Data_Protector_home\bin`

UNIX:

Command	Directory
`omnib`	`opt/omni/bin`
`omnidb`	
`syb_tool`	
`testbar`	
`omnigetmsg`	`opt/omni/lbin`
`util_cmd`	
`util_sybase.pl`	

To run the commands, you must have appropriate Data Protector user rights. For information, see the online Help index: “user groups” and “adding users”.

If the names of the database or database instances are in a non-ASCII encoding, set the `OB2_CLI_UTF8` environment variable to `1` to enable unicode output of the Data Protector Sybase CLI utilities. The terminal application must also use a UTF-8 locale.

Configuring the integration

You need to configure Sybase users and every Sybase Adaptive Server instance (**Sybase instance**) you intend to back up from or restore to.

Prerequisites

- Ensure that you have correctly installed and configured Sybase Server.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on the Sybase Server, see the *Adaptive Server Enterprise System Administration Guide* and *Adaptive Server Enterprise Installation and Configuration Guide*.

  Every Sybase instance and its default Sybase Backup Server must be configured on the same system.

- Ensure that you have correctly installed Data Protector. On how to install the Data Protector Sybase integration in various architectures, see the *HP Data Protector installation and licensing guide*.

  Every Sybase Server system you intend to back up from or restore to must have the Data Protector `Sybase Integration` component installed.

Before you begin

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

Cluster-aware clients

Configure Sybase instances only on one cluster node, since the configuration files reside on the Cell Manager.

If you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name.

Configuring Sybase users

On UNIX, add user `root` and the Sybase Server administrator (the owner of the `isql` utility) to the Data Protector `admin` or `operator` user group. For information, see the online Help index: “adding users”.

This chapter assumes that the Sybase Server administrator is user `sybase` in the group `sybase`.

Configuring Sybase instances

Provide Data Protector with Sybase instance configuration parameters:

- Pathname of the Sybase Server home directory
- Pathname of the Sybase `isql` utility
- Sybase instance name
- Sybase instance user
- Password of the Sybase instance user
- Name of the Sybase *`SYBASE_ASE`* directory

- Name of the Sybase *SYBASE_OCS* directory

Data Protector then creates the Sybase instance configuration file on the Cell Manager and verifies the connection to the Sybase Backup Server.

To configure a Sybase instance, use the Data Protector GUI. On UNIX, you also use the Data Protector CLI.

Before you begin

- Ensure that the default Sybase Backup Server of the Sybase instance is online.

Using the Data Protector GUI

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

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

   In **Application database**, type the Sybase instance name.

   UNIX only: Type `sybase` in both **Username** and **Group name**. This user will be the backup owner.

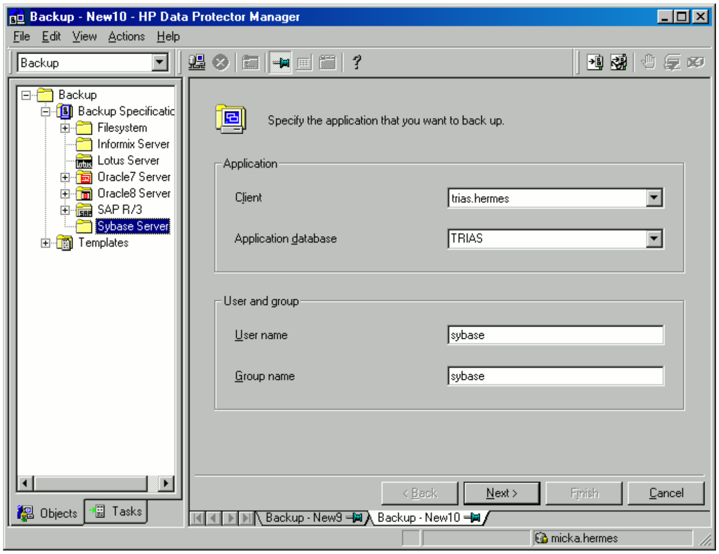

Specify the application that you want to back up.

Application

Client: trias.hermes

Application database: TRIAS

User and group

User name: sybase

Group name: sybase

Figure 3 Specifying the Sybase instance

Click **Next**.

5. In the **Configure Sybase** dialog box, review and, if necessary, correct the configuration parameters that are filled in automatically. On Windows, all configuration parameters are determined automatically. On UNIX, you need to set the Sybase Server home directory, and username and password of the Sybase instance user with the Sybase right to back up and restore databases.

   command, username and password of the Sybase instance user with the Sybase right to back up and restore databases, and the names of the and directories. See Figure 4 on page 33 and Figure 5 on page 33.

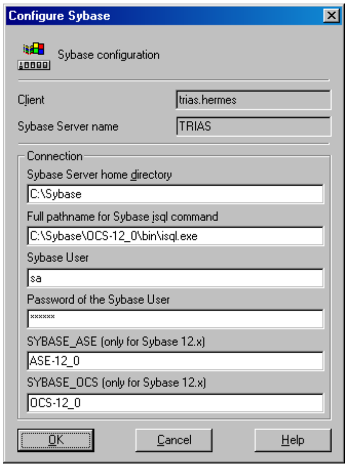

Client: trias.hermes

Sybase Server name: TRIAS

Sybase Server home directory: C:\Sybase

Full pathname for Sybase isql command: C:\Sybase\OCS-12_0\bin\isql.exe

Sybase User: sa

Password of the Sybase User: ******

SYBASE_ASE (only for Sybase 12.x): ASE-12_0

SYBASE_OCS (only for Sybase 12.x): OCS-12_0

Figure 4 Configuring a Sybase instance (Windows)

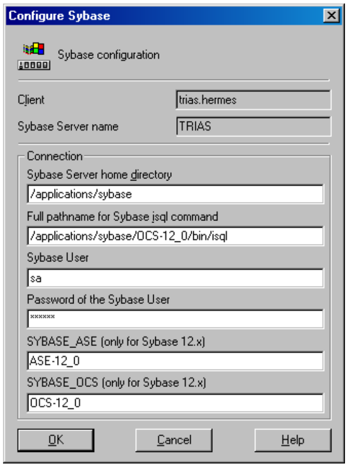

Client: trias.hermes

Sybase Server name: TRIAS

Sybase Server home directory: /applications/sybase

Full pathname for Sybase isql command: /applications/sybase/OCS-12_0/bin/isql

Sybase User: sa

Password of the Sybase User: ******

SYBASE_ASE (only for Sybase 12.x): ASE-12_0

SYBASE_OCS (only for Sybase 12.x): OCS-12_0

Figure 5 Configuring a Sybase instance (UNIX)

Click **OK**.

6. The Sybase instance is configured. Exit the GUI or proceed with creating the backup specification at Step 6 on page 37.

Using the Data Protector CLI

Run:

Windows: `perl -I..\lib\perl util_sybase.pl -CONFIG` *Sybase_instance Sybase_home isql_path Sybase_user Sybase_password Sybase_ASE Sybase_OCS*

UNIX: `util_sybase.pl -CONFIG` *Sybase_instance Sybase_home isql_path Sybase_user Sybase_password Sybase_ASE Sybase_OCS*

Parameter description

Sybase_instance	Name of the Sybase instance.
Sybase_home	Pathname of the Sybase Server home directory.
isql_path	Pathname of the Sybase `isql` command.
Sybase_user	Sybase instance user with the Sybase right to back up and restore databases.
Sybase_password	Password of the Sybase instance user.
Sybase_ASE	Name of the Sybase *Sybase_ASE* directory.
Sybase_OCS	Name of the Sybase *Sybase_OCS* directory.

Example 1

To configure the Sybase instance `mysybase`, run:

```
util_sybase.pl -CONFIG mysybase /applications/sybase.12/
/applications/sybase.12/OCS-12_0/bin/isql sa " " ASE-12_0
OCS-12_0
```

Successful configuration returns `*RETVAL*0`. Otherwise, `*RETVAL*`*`error_number`* is returned.

To get the error description, run:

`omnigetmsg 12` *`error_number`*.

Checking the configuration

You can check the configuration of a Sybase instance after you have created at least one backup specification for the Sybase instance. Use the Data Protector GUI. On UNIX, you can also use the Data Protector CLI.

Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Sybase Server**. Click the backup specification to display the Sybase instance to be checked.
3. Right-click the instance and click **Check configuration**.

Using the Data Protector CLI

Run:

Windows: `perl -I..\lib\perl util_sybase.pl -CHKCONF` *`Sybase_instance_name`*

UNIX: `util_sybase.pl -CHKCONF` *`Sybase_instance_name`*

Backup

The Data Protector Sybase integration provides online backup of the following types:

Table 5 Backup types

Full	Backs up all selected Sybase databases and transaction logs.
Trans	Backs up changes made to the transaction logs since the last backup of any type.

To be prepared for hardware or software failures on your system:

- Regularly back up Sybase system databases.

Back up the `master` database every time you create, alter, or delete a device or database. Back up the `model` database and `system procedure` database every time you change them.

- Keep a copy of the following system tables:
  - `sysusages`
  - `sysdatabases`
  - `sysdevices`
  - `sysloginroles`
  - `syslogins`

Creating backup specifications

Create a backup specification using the Data Protector GUI.

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

   In **Application database**, type the Sybase instance name.

   UNIX only: Type **sybase** in both **Username** and `Group name`. This user is the backup owner.

   Click **Next**.
5. If the Sybase instance is not configured for use with Data Protector, the **Configure Sybase** dialog box is displayed. Configure it as described in "Configuring Sybase instances" on page 30.

6. Select the databases you want to back up.

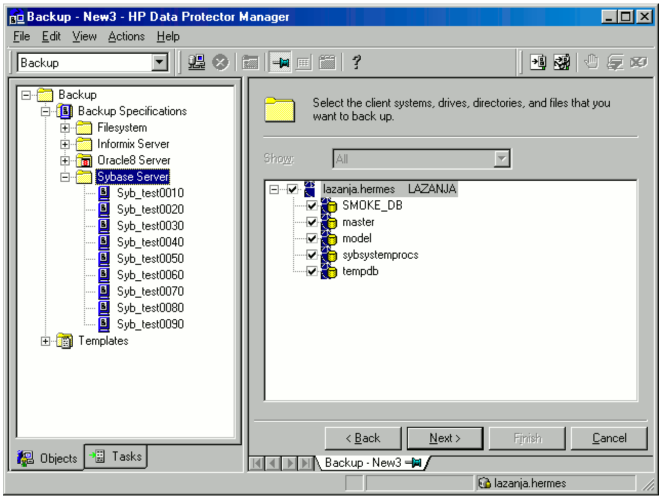

Select the client systems, drives, directories, and files that you want to back up.

Figure 6 Selecting backup objects

Click **Next**.

7. Select devices to use for the backup.

   To specify device options, right-click the device and click **Properties**.

   Click **Next**.

8. Set backup options. For information on application specific options, see Table 6 on page 40.

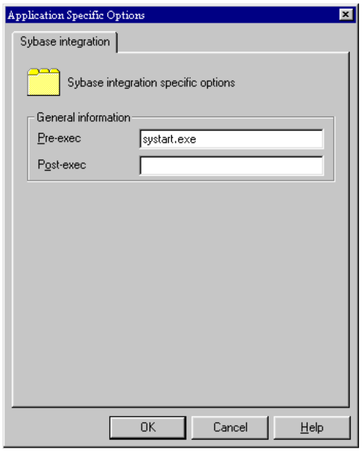

Application Specific Options

Sybase integration

Sybase integration specific options

General information

Pre-exec: systart.exe

Post-exec:

Figure 7 Pre- and post-exec commands (Windows)

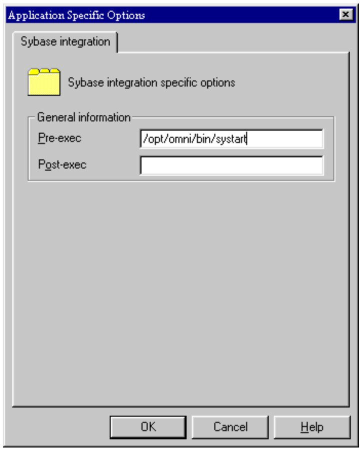

Application Specific Options

Sybase integration

Sybase integration specific options

General information

Pre-exec: /opt/omni/bin/systart

Post-exec:

Figure 8 Pre- and post-exec commands (UNIX)

Click **Next**.

9. Optionally, schedule the backup. For more information, see “Scheduling backup specifications” on page 40.

   Click **Next**.

10. View the properties of objects selected for backup. If you have selected only specific databases, not the whole instance, you can specify the number of concurrent data streams for backing up a particular database: right-click the database and click **Properties**.

    This option is equivalent to Sybase *dump striping*.

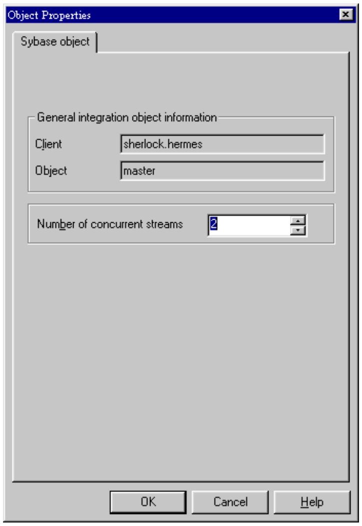

General integration object information

Client: sherlock.hermes

Object: master

Number of concurrent streams: 2

Figure 9 Specifying the number of concurrent streams

    The Sybase Backup Server then splits the database into approximately equal parts and sends the parts concurrently to devices according to device concurrency values.

    If the total sum of device concurrencies is big enough, two or more databases can be backed up simultaneously.

    Click **Next**.

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

TIP:

Preview your backup specification before using it for real. See "Previewing backup sessions" on page 41.

Table 6 Sybase backup options

`Pre-exec`, `Post-exec`	Specify a command that will be started by `ob2sybase.exe` (Windows) or `ob2sybase.pl` (UNIX) on the Sybase Server system before the backup of every selected database (`pre-exec`) or after it (`post-exec`). Do not use double quotes. ***Windows***: Provide only the name of the command. The command must reside in the *`Data_Protector_home`*`\bin` directory. See Figure 7 on page 38. ***UNIX***: Provide the pathname of the command. See Figure 8 on page 38.

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".

Example

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

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. Under **Session options**, select the **Full** backup type. See Figure 10 on page 41. Click **OK**.
3. Repeat Step 1 on page 40 and Step 2 on page 40 to schedule another backup at 13:00, and another one at 18:00.
4. Click **Apply** to save the changes.

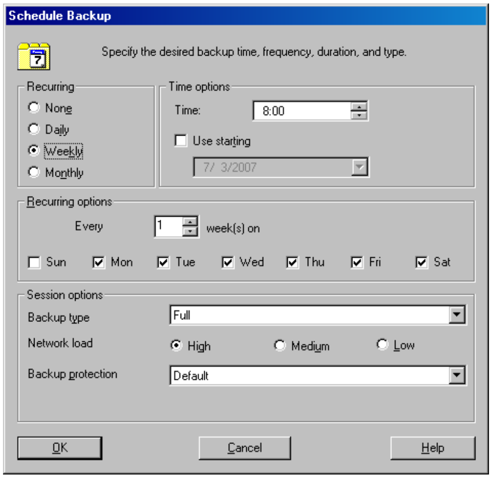

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring
○ None
○ Daily
● Weekly
○ Monthly

Time options
Time: 8:00
□ Use starting
7/ 3/2007

Recurring options
Every 1 week(s) on
□ Sun ☑ Mon ☑ Tue ☑ Wed ☑ Thu ☑ Fri ☑ Sat

Session options
Backup type: Full
Network load: ● High ○ Medium ○ Low
Backup protection: Default

Figure 10 Scheduling a backup specification

Previewing backup sessions

Preview the backup session to test it. 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 **Sybase Server**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify the **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

Run:

```
omnib -sybase_list backup_specification_name -test_bar
```

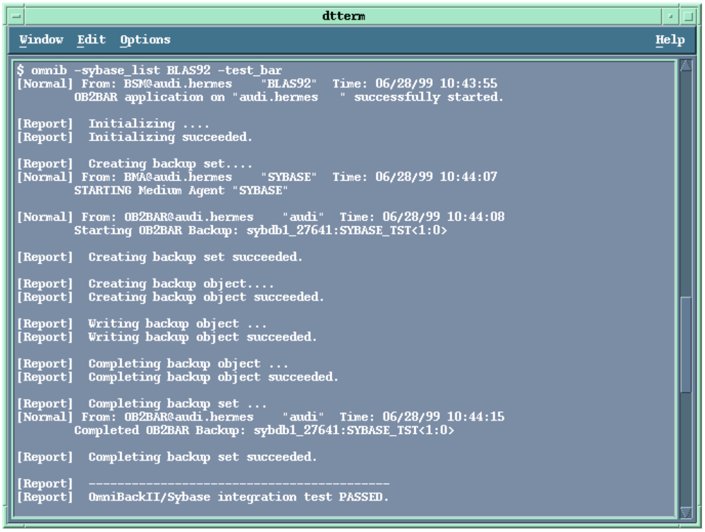

```
$ omnib -sybase_list BLAS92 -test_bar
[Normal] From: BSM@audi.hermes    "BLAS92"  Time: 06/28/99 10:43:55
        OB2BAR application on "audi.hermes   " successfully started.

[Report]  Initializing ....
[Report]  Initializing succeeded.

[Report]  Creating backup set....
[Normal] From: BMA@audi.hermes    "SYBASE"  Time: 06/28/99 10:44:07
        STARTING Medium Agent "SYBASE"

[Normal] From: OB2BAR@audi.hermes    "audi"  Time: 06/28/99 10:44:08
        Starting OB2BAR Backup: sybdb1_27641:SYBASE_TST<1:0>

[Report]  Creating backup set succeeded.

[Report]  Creating backup object....
[Report]  Creating backup object succeeded.

[Report]  Writing backup object ...
[Report]  Writing backup object succeeded.

[Report]  Completing backup object ...
[Report]  Completing backup object succeeded.

[Report]  Completing backup set ...
[Normal] From: OB2BAR@audi.hermes    "audi"  Time: 06/28/99 10:44:15
        Completed OB2BAR Backup: sybdb1_27641:SYBASE_TST<1:0>

[Report]  Completing backup set succeeded.

[Report]  ------------------------------------------
[Report]  OmniBackII/Sybase integration test PASSED.
```

Figure 11 Example of previewing a backup

What happens during the preview?

The following are tested:

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

Starting backup sessions

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

Start a backup in any of the following ways:

- Use the Data Protector GUI.
- Use the Data Protector CLI.
- Use the Sybase `isql` utility.

Using the Data Protector GUI

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

Successful backup displays the message `Session completed successfully`.

Using the Data Protector CLI

Run:

```
omnib -sybase_list backup_specification [-barmode sybase_mode]
[options]
```

Parameter description

backup_specification	Name of the Data Protector Sybase backup specification.
sybase_mode	Backup type. Select among {`full` |`trans`}.
options	For information, see the `omnib` man page.

Example

To perform a full backup using the backup specification `FullSybase`, run:

```
omnib -sybase_list FullSybase -barmode full
```

Using Sybase commands

To start a database backup from the client where the database is located, using the Sybase `isql` utility:

1. Check if the devices to be used contain formatted (initialized) media with enough free space.
2. Verify the backup options in the Data Protector Sybase backup specification.
3. Log in to the Sybase Server system as user `sybase`.

4. Run the Sybase `isql` command:

   ```
   isql -SSybase_instance -USybase_user -PSybase_password
   dump database database to "ob2syb::backup_specification"
   ```

   Parameter description

Sybase_instance	Sybase instance name.
Sybase_user	Sybase instance user.
Sybase_password	Password of the Sybase instance user.
database	Name of the database to be backed up.
backup_specification	Name of the Data Protector Sybase backup specification.

Restore

Restore Sybase databases using the Sybase `isql` utility.

To restore a Sybase database:

1. Restore a full backup of the Sybase database.
2. Restore subsequent transaction backups (if they exist).

Localized database names

If the names of backed up objects contain characters that cannot be displayed using the current language group (on Windows) or code page (on UNIX):

1. Set the encoding used on the terminal to UTF-8.
2. ***Windows only***: Set the environment variable `OB2_CLI_UTF8` to `1`.

3. When gathering information for restore, redirect the output of the `syb_tool` or `omnidb` command to a text file.

   If you need to edit the file containing the load command, use a UTF-8 aware editor that does not set the first byte ("BOM"), since such a file is not supported by `isql`. Note that the Windows Notepad editor cannot be used.

   For details, see “Finding information for restore” on page 45.

4. When restoring the objects, add the `-i` *`file_name`* `-J utf8` options to the `isql` command, where *`file_name`* is the file with the load command.

   For details, see “Restoring using the Sybase isql command” on page 50.

Finding information for restore

To restore a corrupted database, first find the necessary media and the session ID of the last full backup. If you have backed up the database using several streams, also determine the number of streams.

Use the Data Protector GUI or CLI.

Using the Data Protector GUI

In the Internal Database context, expand `Objects` or `Sessions`. To view details on a session, right-click the session and click `Properties`.

Using the Data Protector CLI

Use the Data Protector `syb_tool` command or the standard Data Protector CLI commands.

Using the Data Protector syb_tool command

The Data Protector `syb_tool` command returns the exact Sybase `load` command needed for restore.

The syntax of the `syb_tool` command is:

```
syb_tool database Sybase_instance
-date YYYY/MM/DD.hh:mm:ss
  [ -new_db new_database  ]
 [ -new_server new_Sybase_instance ]
 [ -file file  ]
 [ -media    ]
```

Parameter description

`database`	Database to be restored.
`Sybase_instance`	Sybase instance from which the database to be restored was backed up.
`date`	Point in time. The first backup version created after this point in time is restored. Use the `0-24h` time format.
`new_database`	Target database to which to restore.
`new_Sybase_instance`	Target Sybase instance to which to restore.
`file`	Pathname of a file to which the `load` command or command sequence is recorded.
`-media`	Lists media needed for the restore.

To define the time interval between the closure of transaction logs and the start of a backup session, set the global variable `OB2SybaseTransLogDelay`. The default value is 20 seconds.

Example 1

To get the `load` command that restores `database1` of the Sybase instance `audi` from the first backup performed after 12.00 noon on June 1, 1999, and to get the necessary media, run:

```
syb_tool database1 audi -date 1999/06/01.12:00:00 -media
```

See Figure 12 on page 47.

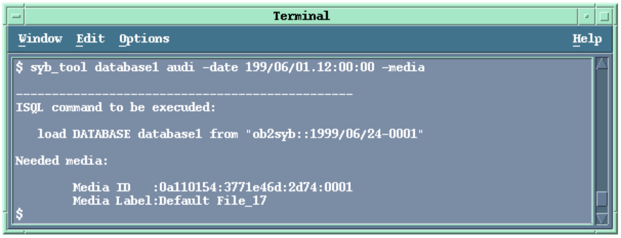

```
$ syb_tool database1 audi -date 199/06/01.12:00:00 -media

-----------------------------------------------
ISQL command to be execuded:

   load DATABASE database1 from "ob2syb::1999/06/24-0001"

Needed media:

        Media ID    :0a110154:3771e46d:2d74:0001
        Media Label:Default File_17
$
```

Figure 12 Running the syb_tool command

Example 2

To get the `load` command that restores `database1` of the Sybase instance `sherlock` from the first backup performed after 12.00 noon on June 1, 1999, to get the necessary media, and to record the `load` command to the file `c:/tmp/isqlfile` (Windows), run:

```
syb_tool database1 sherlock -date 1999/06/01.12:00:00 -file
c:\tmp\isqlfile -media
```

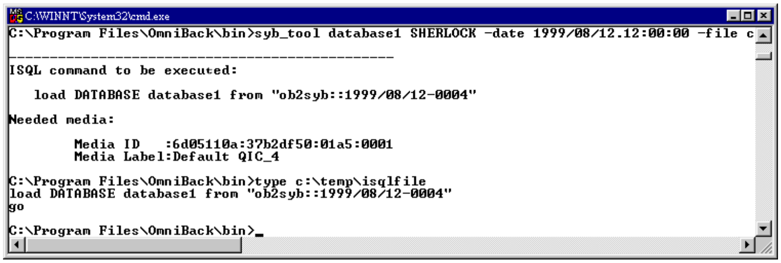

```
C:\Program Files\OmniBack\bin>syb_tool database1 SHERLOCK -date 1999/08/12.12:00:00 -file c

-----------------------------------------------
ISQL command to be executed:

   load DATABASE database1 from "ob2syb::1999/08/12-0004"

Needed media:

        Media ID    :6d05110a:37b2df50:01a5:0001
        Media Label:Default QIC_4

C:\Program Files\OmniBack\bin>type c:\temp\isqlfile
load DATABASE database1 from "ob2syb::1999/08/12-0004"
go

C:\Program Files\OmniBack\bin>_
```

Figure 13 Running the syb_tool command with the -file and -media options

Example 3

To get the `load` command that restores `database1` to `database2` from the first backup performed after 12.00 noon on June 1, 1999, run:

```
syb_tool database1 audi -date 1999/06/01.12:00:00 -new_db
database2 -media
```

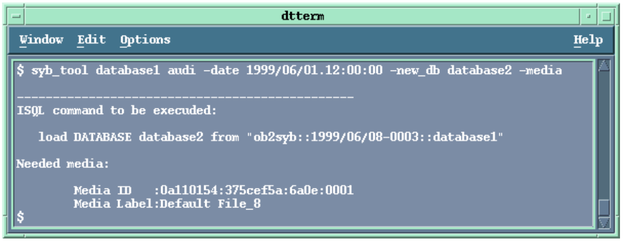

```
$ syb_tool database1 audi -date 1999/06/01.12:00:00 -new_db database2 -media

-----------------------------------------------
ISQL command to be executed:

   load DATABASE database2 from "ob2syb::1999/06/08-0003::database1"

Needed media:

        Media ID   :0a110154:375cef5a:6a0e:0001
        Media Label:Default File_8
$
```

Figure 14 The load command for restore to a different database

Example 4

To get the `load` command that restores `database1` of the Sybase instance `audi` to the Sybase instance `toplarna`, run:

```
syb_tool database1 audi -date 1999/06/01.12:00:00 -new_server
toplarna -file /tmp/isql -media
```

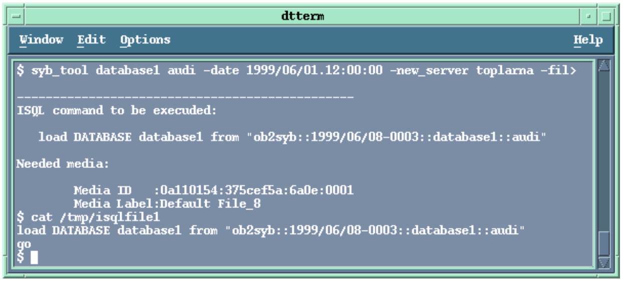

```
$ syb_tool database1 audi -date 1999/06/01.12:00:00 -new_server toplarna -fil>

-----------------------------------------------
ISQL command to be executed:

   load DATABASE database1 from "ob2syb::1999/06/08-0003::database1::audi"

Needed media:

        Media ID   :0a110154:375cef5a:6a0e:0001
        Media Label:Default File_8
$ cat /tmp/isqlfile1
load DATABASE database1 from "ob2syb::1999/06/08-0003::database1::audi"
go
$ █
```

Figure 15 The load command for restore to a different server

Example 5

To get the `load` command that restores `database1` of the Sybase instance `audi` from the first backup performed after 14:28 on July 7, 1999, and to record the `load` command to the file `/tmp/dudule`, run:

```
syb_tool database1 audi -date 1999/07/07.14:28:00 -file
/tmp/dudule
```

You see in Figure 16 on page 49 that you need to restore one full backup and four transaction log backups, the last one backed up with concurrency 3.

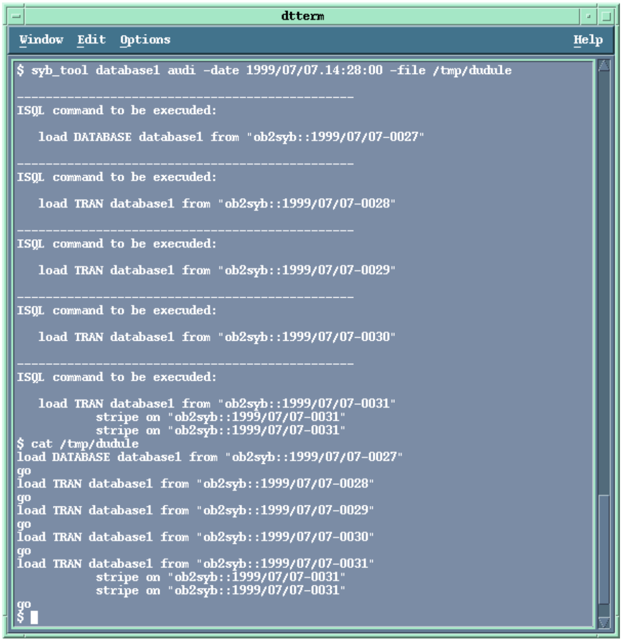

```
$ syb_tool database1 audi -date 1999/07/07.14:28:00 -file /tmp/dudule

-----------------------------------------------
ISQL command to be execuded:

   load DATABASE database1 from "ob2syb::1999/07/07-0027"

-----------------------------------------------
ISQL command to be execuded:

   load TRAN database1 from "ob2syb::1999/07/07-0028"

-----------------------------------------------
ISQL command to be execuded:

   load TRAN database1 from "ob2syb::1999/07/07-0029"

-----------------------------------------------
ISQL command to be execuded:

   load TRAN database1 from "ob2syb::1999/07/07-0030"

-----------------------------------------------
ISQL command to be execuded:

   load TRAN database1 from "ob2syb::1999/07/07-0031"
           stripe on "ob2syb::1999/07/07-0031"
           stripe on "ob2syb::1999/07/07-0031"
$ cat /tmp/dudule
load DATABASE database1 from "ob2syb::1999/07/07-0027"
go
load TRAN database1 from "ob2syb::1999/07/07-0028"
go
load TRAN database1 from "ob2syb::1999/07/07-0029"
go
load TRAN database1 from "ob2syb::1999/07/07-0030"
go
load TRAN database1 from "ob2syb::1999/07/07-0031"
           stripe on "ob2syb::1999/07/07-0031"
           stripe on "ob2syb::1999/07/07-0031"

go
$ █
```

Figure 16 Loading transaction logs from multiple backups

Using the standard Data Protector CLI commands

1. Get a list of backed up Sybase databases:

   `omnidb -sybase`

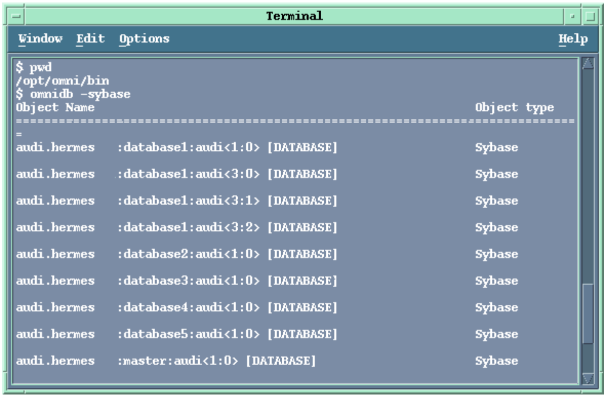

```
$ pwd
/opt/omni/bin
$ omnidb -sybase
Object Name                                                 Object type
=============================================================================
=
audi.hermes   :database1:audi<1:0> [DATABASE]              Sybase

audi.hermes   :database1:audi<3:0> [DATABASE]              Sybase

audi.hermes   :database1:audi<3:1> [DATABASE]              Sybase

audi.hermes   :database1:audi<3:2> [DATABASE]              Sybase

audi.hermes   :database2:audi<1:0> [DATABASE]              Sybase

audi.hermes   :database3:audi<1:0> [DATABASE]              Sybase

audi.hermes   :database4:audi<1:0> [DATABASE]              Sybase

audi.hermes   :database5:audi<1:0> [DATABASE]              Sybase

audi.hermes   :master:audi<1:0> [DATABASE]                 Sybase
```

Figure 17 Example of a list of backed up Sybase databases

2. Get a list of backup sessions for a specific object, including the session ID:

   ```
   omnidb -sybase "object_name"
   ```

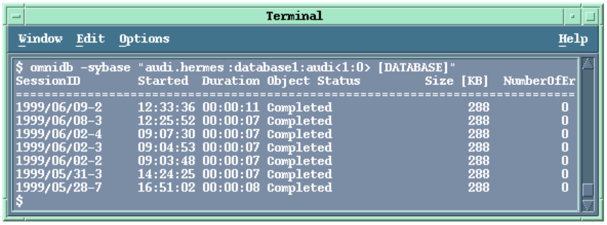

```
$ omnidb -sybase "audi.hermes:database1:audi<1:0> [DATABASE]"
SessionID          Started  Duration Object Status           Size [KB]  NumberOfEr
==================================================================================
1999/06/09-2       12:33:36 00:00:11 Completed                     288           0
1999/06/08-3       12:25:52 00:00:07 Completed                     288           0
1999/06/02-4       09:07:30 00:00:07 Completed                     288           0
1999/06/02-3       09:04:53 00:00:07 Completed                     288           0
1999/06/02-2       09:03:48 00:00:07 Completed                     288           0
1999/05/31-3       14:24:25 00:00:07 Completed                     288           0
1999/05/28-7       16:51:02 00:00:08 Completed                     288           0
$
```

Figure 18 Example of a list of backup sessions for a specific object

IMPORTANT:

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.

3. Get a list of media needed for restore:

   ```
   omnidb -session session_id -media
   ```

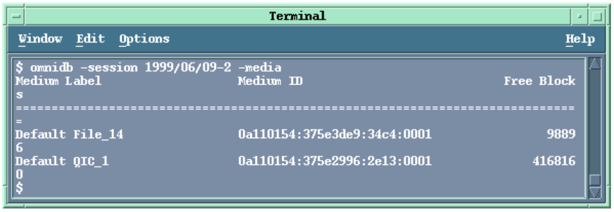

```
$ omnidb -session 1999/06/09-2 -media
Medium Label                      Medium ID                              Free Block
s
=================================================================================
=
Default File_14                   0a110154:375e3de9:34c4:0001                  9889
6
Default QIC_1                     0a110154:375e2996:2e13:0001                416816
0
$
```

Figure 19 Example of finding media needed for restore

For details on the omnidb command, see the omnidb man page.

Restoring using the Sybase isql command

1. On UNIX, log in to the Sybase Server system as user sybase.

2. Run the Sybase `isql` utility:

```
isql -SSybase_instance -USybase_user -PSybase_password [-i
input_file -J utf8]
```

Parameter description

Sybase_instance	Sybase instance name.
Sybase_user	Sybase instance user.
Sybase_password	Password of the Sybase instance user.
input_file	The file to which the `load` parameter was saved. See also "Localized database names" on page 44.

3. If you did not provide the load command in a file, type the desired `load` command in the first line. To run the command(s), type `go` in the last line and press **Enter**.

   The syntax of the Sybase `load` command is:

```
load {database|transaction} new_database from
"ob2syb::version[::database[::Sybase_instance]]"
stripe on
"ob2syb::version[::database[::Sybase_instance]]"
```

Parameter description

`{database|transaction}` Defines whether databases or transaction logs are to be restored.

`version` Session ID of the backup version to restore from. You can also type `latest version` to restore from the latest backup.

`new_database` Target database to which to restore.

`database` Database to be restored.

`Sybase_instance` Sybase instance from which the database to be restored was backed up.

The `stripe` part is needed only when restoring a database backed up with several streams. The number of streams used for backup is displayed in the `Data Protector Monitor` during the backup session.

IMPORTANT:

To restore a database to a new database, first create a new database. The new database should have the same structure as the database to be restored.

For details on the Sybase `load` command, see the *Adaptive Server Enterprise System Administration Guide*.

TIP:

To list all Sybase databases of a particular Sybase instance, run:

Windows:

`perl -I..\lib\perl util_sybase.pl -OBJS0` *`Sybase_instance_name`*

UNIX: `util_sybase.pl -OBJS0` *`Sybase_instance_name`*

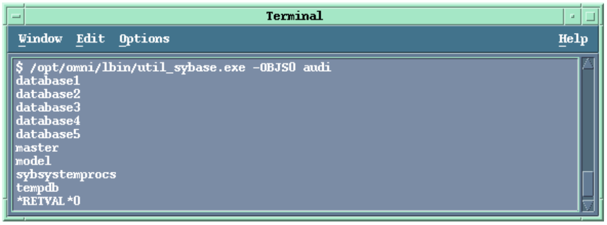

```
$ /opt/omni/lbin/util_sybase.exe -OBJS0 audi
database1
database2
database3
database4
database5
master
model
sybsystemprocs
tempdb
*RETVAL*0
```

Figure 20 Example of a list of Sybase databases

Restore examples

Example 1

To restore the database `database2` from the backup session `1999/06/09-2`, run:

```
1>load database database2 from "ob2syb::1999/06/09-2"
2>go
```

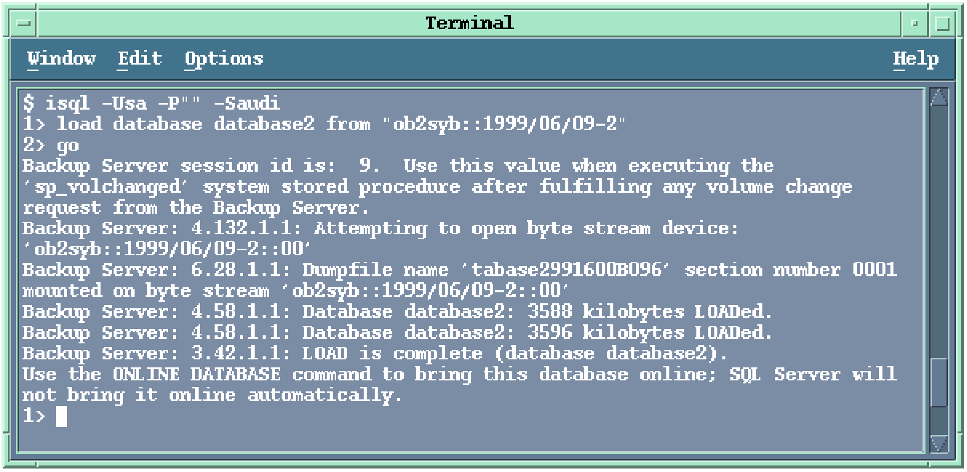

```
$ isql -Usa -P"" -Saudi
1> load database database2 from "ob2syb::1999/06/09-2"
2> go
Backup Server session id is:  9.  Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
request from the Backup Server.
Backup Server: 4.132.1.1: Attempting to open byte stream device:
'ob2syb::1999/06/09-2::00'
Backup Server: 6.28.1.1: Dumpfile name 'tabase2991600B096' section number 0001
mounted on byte stream 'ob2syb::1999/06/09-2::00'
Backup Server: 4.58.1.1: Database database2: 3588 kilobytes LOADed.
Backup Server: 4.58.1.1: Database database2: 3596 kilobytes LOADed.
Backup Server: 3.42.1.1: LOAD is complete (database database2).
Use the ONLINE DATABASE command to bring this database online; SQL Server will
not bring it online automatically.
1>
```

Figure 21 Restoring a database from a specific session

Example 2

To restore the latest version of the database `Sybdata` to a new database, named `Sybdata1`:

1. Create a database device. See Figure 22 on page 54.

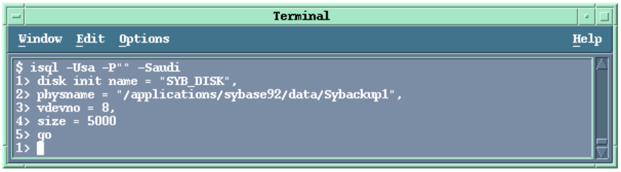

```
$ isql -Usa -P"" -Saudi
1> disk init name = "SYB_DISK",
2> physname = "/applications/sybase92/data/Sybackup1",
3> vdevno = 8,
4> size = 5000
5> go
1>
```

Figure 22 Creating a database device

2. Create an empty database, named `Sybdata1`. See Figure 23 on page 54.

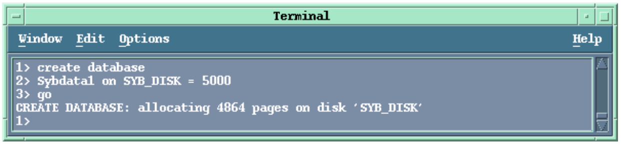

```
1> create database
2> Sybdata1 on SYB_DISK = 5000
3> go
CREATE DATABASE: allocating 4864 pages on disk 'SYB_DISK'
1>
```

Figure 23 Creating an empty database

3. Restore `Sybdata` to `Sybdata1` by running:

```
1>load database Sybdata1 from "ob2syb::latest version::Sybdata"
2>go
```

Example 3

To restore the latest version of the database `database3` backed up with three streams, run:

```
1>load database database3 from "ob2syb::latest version"
2>stripe on "ob2syb::latest version"
3>stripe on "ob2syb::latest version"
4>go
```

Example 4

To start a restore a database from the instance "instance1", which name contains Cyrilic and Latin charaters, and for which the load command was saved in the file `restore_20050609-2.txt`, run :

```
isql -S instance1 -U admin -PSybase_password -J utf8 -i
restore_20050609-2.txt
```

Restoring using another device

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

Specify the new device in the file:

Windows: `Data_Protector_home\Config\server\Cell\restoredev`

UNIX: `/etc/opt/omni/server/cell/restoredev`

Use the format:

`"DEV 1" "DEV 2"`

where `DEV 1` is the original device and `DEV 2` the new device.

IMPORTANT:

Delete this file after use.

On Windows, use the Unicode format for the file.

Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or 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".

Troubleshooting

This section lists general checks and verifications.

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

Before you begin

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

Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors written to `debug.log`, located on the Sybase Server system in:

  Windows: `Data_Protector_home\log`

  UNIX: `/var/opt/omni/log`

- Make a test backup and restore of any filesystem on the problematic client. For information, see online Help.

- In a cluster environment, before performing procedures from the Data Protector CLI, ensure that the environment variable `OB2BARHOSTNAME` is set to the virtual server name. When the Data Protector GUI is used, this is not required.
- Ensure that the Sybase instance and its default Sybase Backup Server are online.
- ***UNIX only:*** Ensure that user `root` and user `sybase` are added to the Data Protector `admin` or `operator` user group.

Additionally, if your configuration or backup failed:

- If you use non-default Sybase settings, ensure that they are registered in:

  Windows: The `System Properties` dialog box, which you access by double-clicking `System` in the `Control Panel`.

  UNIX: The Data Protector Sybase configuration file.

Additionally, if your backup failed:

- Check the configuration of the Sybase instance described in "Checking the configuration" on page 35.
- Test the backup specification as described in "Previewing backup sessions" on page 41.

  If the Data Protector part of the test fails:

  1. ***UNIX only:*** Ensure that the owner of the backup specification is user `sybase` and that it is added to the Data Protector `operator` or `admin` user groups.
  2. Create a Sybase backup specification to back up to a null or file device. If the backup succeeds, then the problem is probably related to devices. For information on troubleshooting devices, see the HP Data Protector troubleshooting guide.

  If the test succeeds, start a backup directly from the Sybase Server. See "Using Sybase commands" on page 43.

Additionally, if your backup or restore failed:

- Test Data Protector data transfer using the `testbar` utility. Log in to the Sybase Server system as user `sybase` and run:
  - If your backup failed:

    ```
    testbar -type:Sybase -appname:Sybase_instance_name
    -bar:backup_specification_name -perform:backup
    ```

  - If your restore failed:

    ```
    testbar -type:Sybase -appname:Sybase_instance_name
    -bar:backup_specification_name -perform:restore
    -object:object_name -version:object_version
    ```

where *object_name* is the name of the object to be restored.

If the test fails:

- Troubleshoot errors. See the text file Trouble.txt located on the Cell Manager in:

  Windows: *Data_Protector_home*\help\enu

  UNIX: /opt/omni/gui/help/C

- On the Sybase Server system, examine system errors, reported in:

  Windows: *Data_Protector_home*\log\debug.log

  UNIX: /var/opt/omni/log/debug.log

Additionally, if your restore failed:

- Ensure that the Data Protector operator user group has the See private objects user right selected. On how to change user rights, see the online Help index: “changing user rights”.

2 Integrating HP OpenView Network Node Manager and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector HP OpenView Network Node Manager (NNM) integration. It describes concepts and methods you need to understand to back up and restore the NNM database.

You can back up or restore NNM objects: the whole database or only parts of it.

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

Table 7 Backup types

Full	Backs up the selected NNM objects.
Incremental	Backs up changes made to the selected NNM objects since the last full backup.

This chapter provides information specific to the Data Protector HP OpenView Network Node Manager integration. For general Data Protector procedures and options, see online Help.

Integration concept

The basic components of the Data Protector NNM integration are the following Perl scripts:

Table 8 Data Protector NNM integration components

`NNMpre.ovpl`	A script without arguments that:

	1. Initiates a special NNM backup, instructing the NNM database to make a direct copy of itself to a location specified in the `solid.ini` file, from which Data Protector backs it up later. **2.** Pauses the eight NNM processes, so that Data Protector can actually back up the NNM data.
`NNMpost.ovpl`	A script without arguments that restarts the NNM processes after the backup completes.
`NNMScript.exe` (Windows only)	A script with a pre- and post- argument that locates the NNM Perl compiler and `NNM pre.ovpl` or `NNMpost.ovpl`, and starts the script.

NOTE:

Files created by the embedded database remain on the disk and are overwritten by future backups. Remove the files manually to free the disk space.

The NNM Perl compiler is used for `NNMpre.ovpl` and `NNMpost.ovpl`.

While HP OpenView Network Node Manager is responsible for read/write operations to disk, Data Protector reads from and writes to devices and manages media.

Configuring the integration

Prerequisites

- Ensure that you have correctly installed and configured NNM.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on backup and recovery strategies and NNM concepts, see the *Reporting and data analysis with HP OpenView network node manager*.
- Ensure that you have correctly installed Data Protector. On how to install the Data Protector NNM integration in various architectures, see the *HP Data Protector installation and licensing guide*.

  Every NNM system you intend to back up from or restore to must have the Data Protector HP OpenView NNM Backup Integration and Disk Agent components installed.

Before you begin

- Configure devices and media for use with Data Protector. For information, see online Help.
- To test whether the NNM system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the NNM system.

Tasks for the NNM administrator

- Communicate the location of the NNM backup directory, specified in the NNM embedded database file `solid.ini`.
- In `solid.ini`, comment out the line beginning with `At=` that schedules a nightly backup of the NNM embedded database.

Backup

The Data Protector NNM integration provides two backup types and two backup modes.

Table 9 Backup types

Full	Backs up all selected NNM objects.
Incremental	Backs up changes made to the selected NNM objects since the last full backup.

Table 10 Backup modes

Offline	The database is taken offline. Consequently, no changes can be made to the database during the backup process, leaving it in a consistent state.
Online	The database is in a paused state and the changes made to the database during the backup process are recorded to temporary files. When the backup completes, the database resumes its normal state and the changes from the temporary files are applied to the database, bringing it to a consistent state.

To perform an offline backup:

1. On the NNM system, take the NNM database offline by running:

   ```
   ovstop
   ```

2. Back up the complete NNM directory using Data Protector.
3. On the NNM system, bring the NNM database online by running:

   ```
   ovstart
   ```

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 **Filesystem**, and click **Add Backup**.
3. Select a template:

   Windows: `NT_NNM_template`

   UNIX: `Unix_NNM_template`

   You can also select the `Blank Filesystem Backup` template or any other template.

   Click **OK**.
4. Select the appropriate client and directories to be backed up from the client.

   Click **Next**.
5. Select devices to use for the backup.

   To specify device options, right-click the device and click **Properties**.

   Click **Next**.
6. Set backup options.

   IMPORTANT:

   If you have selected the NNM template, do not change the default pre- and post-exec options. If you have selected a different template, specify exactly the same pre- and post-exec scripts as specified in the NNM template.

   Click **Next**.

7. Optionally, schedule the backup. For more information, see "Scheduling backup specifications" on page 40.

   Click **Next**.

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

TIP:

Preview your backup specification before using it for real. For details, see the online Help index: "previewing a backup". Then, start a test backup to verify if the pre- and post-exec scripts are functioning.

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".

Example

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

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

2. Under Recurring, select **Weekly**. UnderT ime options , select **8:00**. Under Recurring Options, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See Figure 24 on page 64.

   Click **OK**.

3. Repeat Step 1 on page 63 and Step 2 on page 63 to schedule another backup at 13:00, and another one at 18:00.

4. Click **Apply** to save the changes.

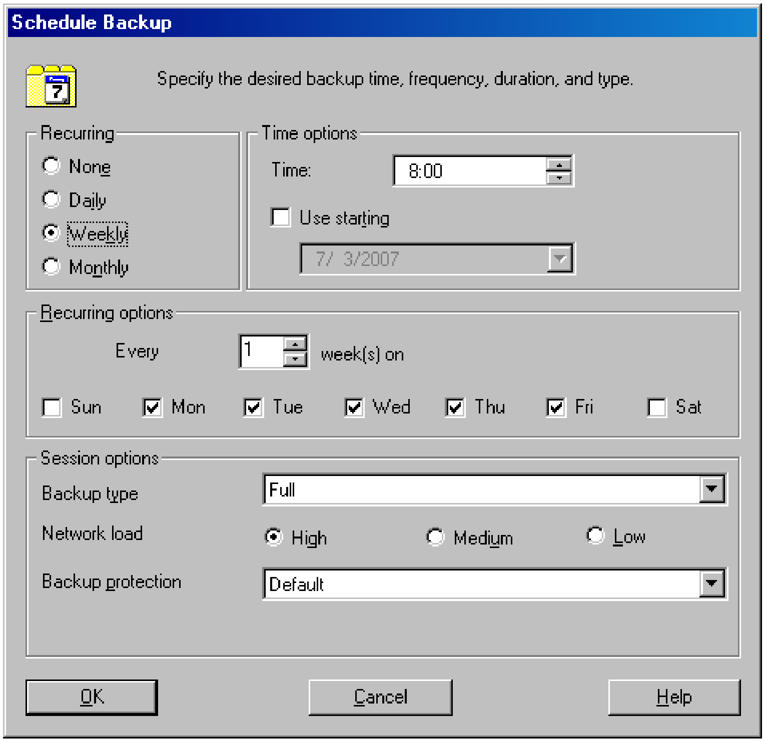

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

○ None
○ Daily
● Weekly
○ Monthly

Time options

Time: 8:00

□ Use starting

7/ 3/2007

Recurring options

Every 1 week(s) on

□ Sun ☑ Mon ☑ Tue ☑ Wed ☑ Thu ☑ Fri □ Sat

Session options

Backup type: Full

Network load: ● High ○ Medium ○ Low

Backup protection: Default

Figure 24 Scheduling a backup specification

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.

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

Restore

To restore NNM objects:

1. Stop all NNM processes.
2. Restore the NNM objects using the Data Protector GUI.
3. Perform the NNM recovery procedures.
4. Restart the NNM processes.

   For details, see the online Help index: “standard restore procedure” and the *NNM reporting and data analysis* manual.

Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or 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”.

Messages generated by scripts, NNM, and Data Protector are logged to the IDB.

Acceptable warnings on Windows

The following warnings, which are likely to occur during an NNM backup, have no impact on the validity of the backup. They are only informational.

Message

```
[Warning] From: session_owner Time: mm/dd/yy hr:mn:sc

[error code] path\HP OpenView\NNM\bin\tcl7.5.dll

Cannot preserve time attributes: ([5] Access is denied.).
```

Description

The file `tcl7.5.dll` is backed up, but the time attributes, which are not significant to Data Protector, are not preserved.

Message

```
[Warning] From: session_owner Time: mm/dd/yy hr:mn:sc
```

```
[error code] path\HP
OpenView\NNM\databases\analysis\default\solid.db Cannot open:
([33] The process cannot access the file ....).
```

Description

The embedded database file referenced in this message has already been backed up as part of the pre-exec script. Its default location is in the `path\HP OpenView\NNM\databases\analysis\default\backup` directory, which is specified in the `solid.ini` file. After the restore, copy the backed up `solid.db` file from that directory to the active `path\HP OpenView\NNM\databases\analysis\default` directory.

Message

```
[Warning] From: session_owner Time: mm/dd/yy hr:mn:sc
```

```
[error code] path\HP
OpenView\NNM\databases\openview\topo\netmon.lock Cannot open:
([33] The process cannot access the file ....).
```

Message

```
[Warning] From: session_owner Time: mm/dd/yy hr:mn:sc
```

```
[error code] path\HP OpenView\NNM\databases\snmpCollect\dblock
Cannot open: ([33] The process cannot access the file ....).
```

Description

These files are not significant to Data Protector.

Troubleshooting

This section lists problems you might encounter when using the Data Protector NNM integration.

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

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.

- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://www.hp.com/support/manuals for an up-to-date list of supported versions, platforms, and other information.

Problems

Problem

The system is already in a paused state

NNM reports:

```
The system is already in a paused state. ’ovpause’ cannot
continue, If a synchronization error has occurred, try
removing the file e:Program Files\HP OpenView\tmp\ovpause.lock
(Windows system) or /var/opt/OV/tmp/ovpause.lock (UNIX system)
and then retrying the ’ovpause’ command.
```

Action

Ensure that the NNM processes are not paused manually before the Data Protector NNM session starts. Otherwise, the pre-exec script `NNNpre.ovpl` fails.

Problem

The system is not in a paused state

NNM reports:

```
The system is not in a paused state. ’ovresume’ cannot
continue. If a synchronization error has occurred, try
creating the empty file e:Program Files\HP
OpenView\tmp\ovpause.lock (Windows systems) or
/var/opt/OV/tmp/ovpause.lock (UNIX systems) and then retrying
the ’ovresume’ command.
```

Action

Ensure that the NNM processes are not resumed manually during the Data Protector NNM session. Otherwise, the post-exec script `NNMpost.ovpl` fails and Data Protector displays the message `Backup completed with errors.`

Problem

ODBC Error: SQLSTATE=HY000

Data Protector reports:

```
ODBC Error:SQLSTATE=HY000 NATIVE ERROR=21306 SOLID
Communication Error 21306: Server ’tcpip 2690’ not found,
connection failed Connect to ODBC data Source "ovdbrun"
failed.
```

Action

Ensure that no NNM processes are paused manually before the Data Protector NNM session starts. Otherwise, the pre-exec script `NNMpre.ovpl` fails because it cannot connect to the NNM embedded database.

Problem

Embedded database is currently in the backup process

NNM reports:

```
Embedded database is currently in the backup process.
```

```
Aborting Data Protector backup.
```

Action

Ensure that the default scheduled backup in the `solid.ini` file is commented out. A Data Protector NNM backup and an active backup of the NNM embedded database cannot be performed simultaneously.

Problem

Wrong number of arguments

On Windows, Data Protector reports:

```
Wrong number of arguments. Please specify pre or post backup.
"NNMScript.exe pre" for pre-exec script "NNMScript.exe post"
for post-backup script.
```

Action

Correct the number of arguments for `NNMScript.exe`, as specified in the pre-exec and post-exec backup options.

Problem

Couldn’t find HP OpenView Network Node Manager key

On Windows, Data Protector reports:

```
Couldn’t find HP OpenView Network Node Manager key in
registry.
```

Action

Ensure that NNM is installed on the target client and that the registry key `HP OpenView Network Node Manager` exists under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView`.

Problem

Couldn’t find the HP OpenView Network Node Manager PathName

On Windows, Data Protector reports:

```
Couldn’t find the HP OpenView Network Node Manager PathName
in registry.
```

Action

Ensure that a registry entry with the name `PathName` exists under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\HP OpenView Network Node Manager` and has a string value.

Problem

Couldn’t find OmniBack II key

On Windows, NNM reports:

```
Couldn’t find OmniBack II key in registry.
```

Action

Ensure that Data Protector with a Disk Agent is installed on the target client and that the registry key `OmniBack II` exists under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView`. Any other name causes problems, potentially requiring reinstallation of the Disk Agent.

Problem

Couldn’t find the Data Protector HomeDir

On Windows, NNM reports:

```
Couldn’t find the Data Protector HomeDir in registry.
```

Action

Ensure that a registry entry with the name `HomeDir` exists under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Common`, having a string value for the Data Protector path. Otherwise, create it or reinstall the Disk Agent.

Problem

Incorrect argument

On Windows, Data Protector reports:

```
Incorrect arguments. Use "pre" or "post".
```

Action

Ensure that NNMScript.exe has correct arguments, as specified in the pre- and post-exec backup options. The arguments are not case-sensitive.

Problem

Failure starting `NNM_perl_compiler_pathData Protector_Home_Dir\bin\*.ovpl.`

On Windows, Data Protector reports:

```
Failure starting NNM_perl_compiler_path
Data_Protector_home\bin\*.ovpl.
```

Action

Ensure that the NNM Perl compiler has not been removed and paths for Data Protector and NNM in the registry are correct.

Problem

Execution of `NNM_perl_compiler_path` *`Data_Protector_home`*`\bin\*.ovpl` **failed**

On Windows, NNM reports:

```
Execution of
NNM_perl_compiler_pathData_Protector_home\bin\*.ovpl failed.
```

Action

Ensure that *path*\HP OpenView\NNM\bin is in the PATH and scripts are in the *Data_Protector_home*\bin directory. Otherwise, the command that starts NNMpre.ovpl or NNMpost.ovpl fails.

3 Integrating NDMP Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Network Data Management Protocol Server integration (**NDMP Server integration**). It describes concepts and methods you need to understand to perform filesystem backups and restores on a Network Attached Storage device.

Network Data Management Protocol (**NDMP**) is a protocol used to manage backup and restore operations on a Network Attached Storage device. NDMP uses a client server model, where the Data Protector NDMP Media Agent client controls the backup, while the NDMP server performs the actual backup operations.

The Data Protector NDMP Server integration offers interactive and scheduled filesystem backups of the following types:

- Full
- Inc1

For information on these backup types, see the *HP Data Protector concepts guide*.

The Data Protector NDMP Server integration offers two restore types:

- Standard filesystem restore
- Direct access restore

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

Integration concept

Data Protector integrates with NDMP Server through the Data Protector NDMP library and the NDMP Media Agent. The Data Protector NDMP library channels communication between the Data Protector Session Manager, and, via the NDMP interfaces, the NDMP Server. Figure 25 on page 74 shows the architecture of the integration.

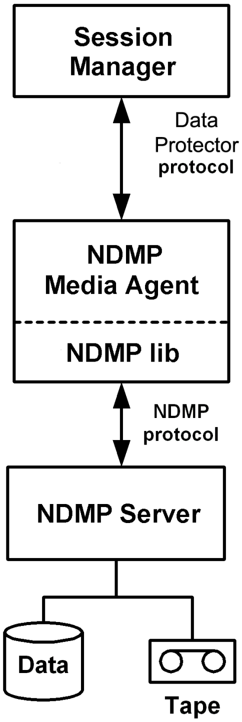

Figure 25 Data Protector NDMP Server integration architecture

Legend	
Session Manager	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore. No Data Protector Disk Agents are involved in the session because the whole functionality is already implemented within the NDMP Media Agent.
NDMP Media Agent	The NDMP client, which contains a layer called the NDMP library. The library enables the NDMP Media Agent to communicate with the NDMP Server through the NDMP interfaces.

For more information on the NDMP protocol and NDMP interfaces, see the NDMP documentation.

Data Protector supports two different NDMP Server types:

- NetApp NAS device (**NetApp**)
- Celerra NAS device (**Celerra**)

In a typical environment (Figure 26 on page 75), the NDMP Server system and the Data Protector client with the NDMP Media Agent installed (**NDMP client**) are connected to the LAN. However, data from the NDMP Server disks does not flow through the LAN, it is backed up to a tape device connected to the NDMP Server system. The NDMP client initiates, monitors, and controls data management and the NDMP Server executes these operations, having a direct control over devices connected to it and over the backup and restore speed.

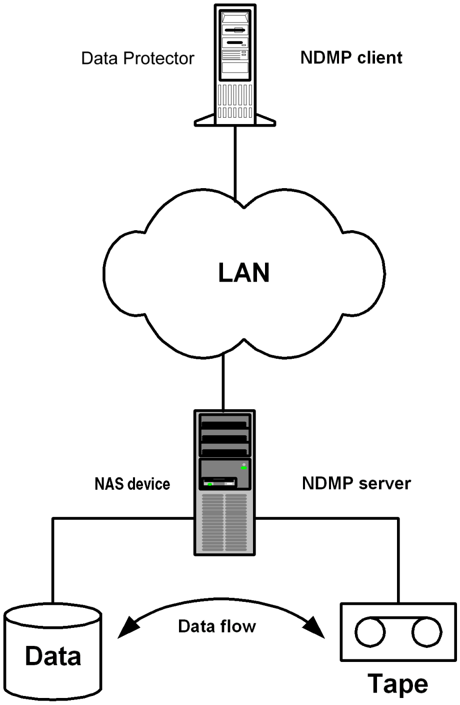

Figure 26 The NDMP environment configuration

Due to the NDMP catalog handling design, Data Protector caches the entire catalog on the NDMP client before storing it to the Data Protector internal database (IDB). Since the catalog can increase in size significantly, the NDMP client caches parts of the catalog into **file history swap files**, located in the following directory:

Windows: `Data_Protector_home\tmp`

UNIX: `/var/opt/omni/tmp`

For more information on file history swap files, see "The NDMP specific omnirc file variables" on page 99.

Configuring the integration

To configure the Data Protector NDMP Server integration:

1. Import the NDMP Server system into the Data Protector cell.
2. Create a media pool for NDMP media.
3. Configure NDMP devices.

Prerequisites

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

  Every NDMP client (Data Protector client that controls the NDMP Server backup) must have the Data Protector `NDMP Media Agent` component installed.

Importing NDMP server systems

Import the NDMP Server system using the Data Protector GUI:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, right-click **Clients** and click **Import Client**.

3. In `Name`, type the name of the NDMP Server system you want to import and select **NDMP Server**.

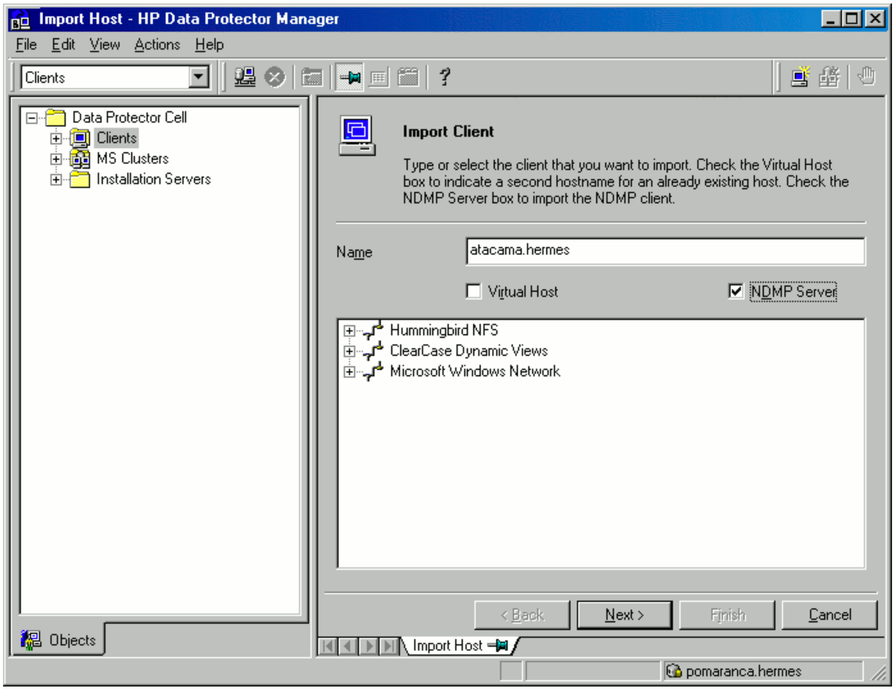

Import Client

Type or select the client that you want to import. Check the Virtual Host box to indicate a second hostname for an already existing host. Check the NDMP Server box to import the NDMP client.

Name: atacama.hermes

□ Virtual Host ☑ NDMP Server

Figure 27 Specifying an NDMP server system

Click **Next**.

4. In **NDMP Type**, select the NAS device type.

In Port, specify the TCP/IP port number of the NDMP Server. The default number is 10000.

Provide the NDMP Server system user account that will be used by Data Protector to connect to the NDMP Server system. This user must have permission to read from and write to the NDMP media.

The Data Protector NDMP integration supports the “none”, “text”, and “MD5” NDMP authentication methods. Data Protector automatically detects and uses the method supported by your NDMP Server.

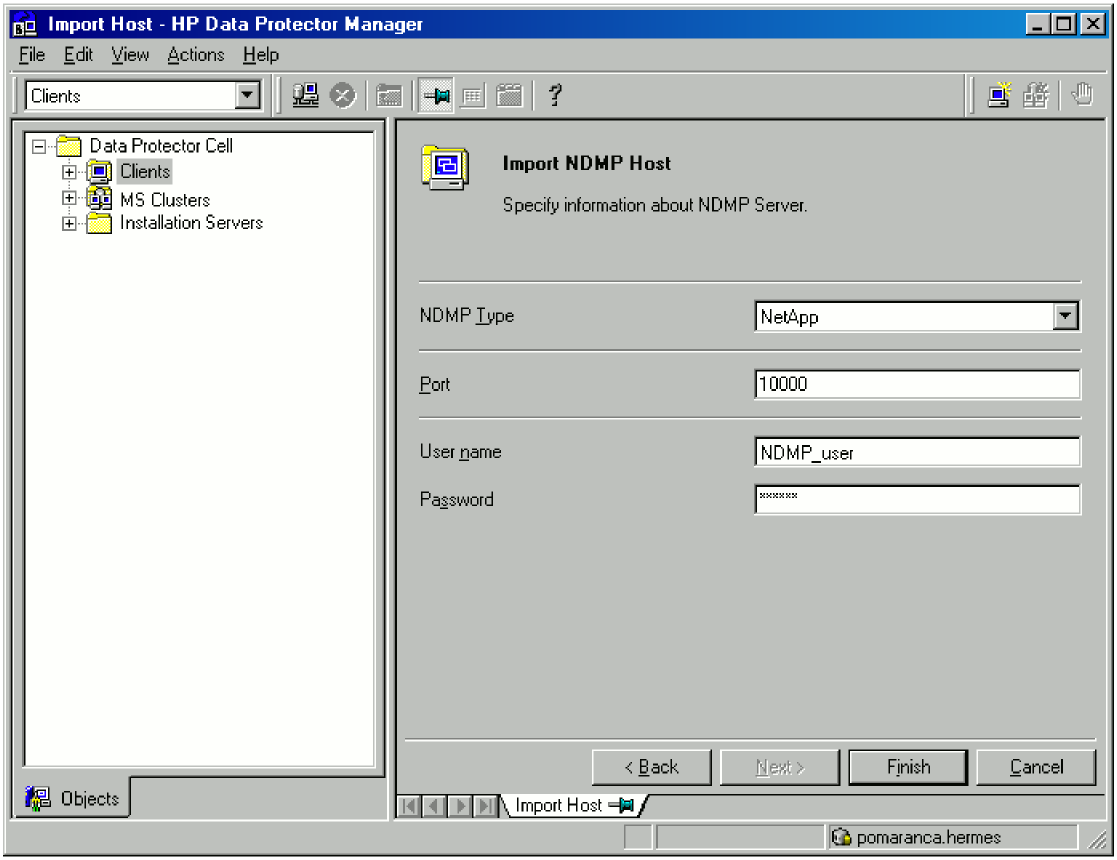

Import NDMP Host

Specify information about NDMP Server.

NDMP Type	NetApp
Port	10000
User name	NDMP_user
Password	******

Figure 28 Specifying an NDMP server system

Click **Finish**.

Creating media pools

Create a special media pool for NDMP media. For information, see the online Help index: “creating media pools”.

The NDMP media pool can only be used by devices using the NDMP data format (**NDMP devices**).

Limitations

- A medium cannot be used by different NMDP Server types. Consequently, data that was backed up from one NDMP Server type cannot be restored to another NDMP Server type.

Configuring NDMP devices

Configure NDMP devices using the Data Protector GUI.

Prerequisites

- The NDMP Server system must have a tape drive connected to it.

  The drive must be supported by both NDMP Server and Data Protector.

Library robotics can be connected to:

- NDMP Server system (Figure 29 on page 80).
- NDMP client (Figure 30 on page 81).
- Data Protector client with the general Media Agent installed (**general Media Agent client**) (Figure 30 on page 81).

If it is connected to the NDMP Server system, the library robotics must be supported by both NDMP Server and Data Protector.

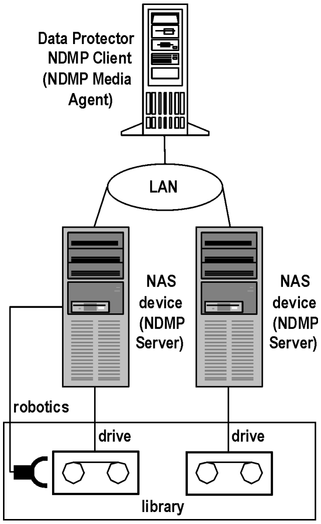

Figure 29 Library configuration—I

Figure 30 Library configuration—II

Several drives can be connected to the NDMP Server system.

If library robotics or drives are connected to the NDMP Server system, they can be controlled only by an NDMP client.

Library drives can be shared between multiple NDMP Server systems and general Media Agent clients, and with other applications. For more information, see the *HP Data Protector concepts guide*.

Limitations

- NDMP devices can only use NDMP media pools.

Configuring tape libraries

To configure a tape library with robotics connected to the NDMP Server system:

1. In the Context list, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices**, and then click **Add Device**.

3. Type a name for the device. Optionally, describe the device. See Figure 31 on page 82.

   In **Device Type**, select **SCSI Library**.

   In **Interface Type**, select the NAS device used.

   In **Client**, select the NDMP client that will control the library through the NDMP Server.

   In **NDMP Server**, select the NDMP Server system with the library robotics connected to it.

   Optionally, in **Management Console URL**, type a valid URL of the library management console. It will enable you to invoke a web browser and load the management console interface directly from the Data Protector GUI.

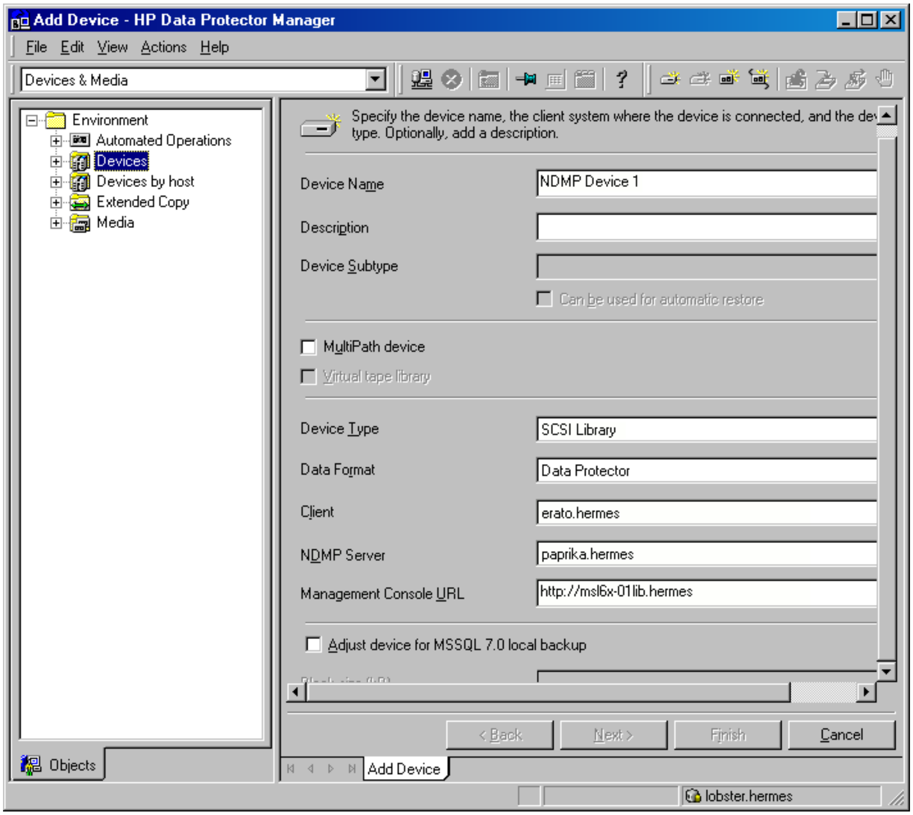

Specify the device name, the client system where the device is connected, and the de type. Optionally, add a description.

Device Name: NDMP Device 1

Description:

Device Subtype:

□ Can be used for automatic restore

□ MultiPath device

□ Virtual tape library

Device Type: SCSI Library

Data Format: Data Protector

Client: erato.hermes

NDMP Server: paprika.hermes

Management Console URL: http://msl6x-01lib.hermes

□ Adjust device for MSSQL 7.0 local backup

Figure 31 Configuring a library

Click **Next**.

4. Specify library robotics' SCSI address and drive handling. For information, see "Network appliance configuration" on page 86 and "EMC Celerra configuration" on page 87.

   Click **Next**.

5. Specify slots to be used by Data Protector.

   Click **Next**.

6. Select the media type used in the library.

   Click **Next**.

7. Click **Finish** and then click **Yes** to configure drives in the library.

8. Type a name for the drive. Optionally, describe the drive.

   In **Client**, select the NDMP client that will control the library through the NDMP Server.

   In **NDMP Server**, select the NDMP Server system with the library robotics connected to it.

   In **Data Format**, select the NAS device used.

   Click **Next**.

9. Specify the drive's NDMP SCSI address. For information, see "Network appliance configuration" on page 86 and "EMC Celerra configuration" on page 87.

   Do not change the drive index number.

   Click **Next**.

10. Specify the media pool.

    To specify advanced device options, click **Advanced**. For information on supported block sizes, see Table 14 on page 88.

NOTE:

Multiplexing data streams is not supported by NDMP Server, limiting device concurrency to 1.

11. Click **Yes** to create another drive or **NO** to finish.

    On how to configure a tape library with robotics connected to a Data Protector NDMP or General Media Agent client and drives connected to the NDMP Server system, see the online Help index: “configuring SCSI libraries”. Then configure the drives as described in Step 8 on page 83 through Step 11 on page 84 inStep 8 on page 83.

Configuring standalone devices

To configure a standalone device:

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices**, and then click **Add Device**.

3. Type a name for the device. Optionally, describe the device.

   In **Device Type**, select **Standalone.**

   In **Data Format**, select the NAS device used.

   In **Client**, select the NDMP client that will control the device through the NDMP Server.

   In **NDMP Server**, select the NDMP Server system to which the standalone device is connected.

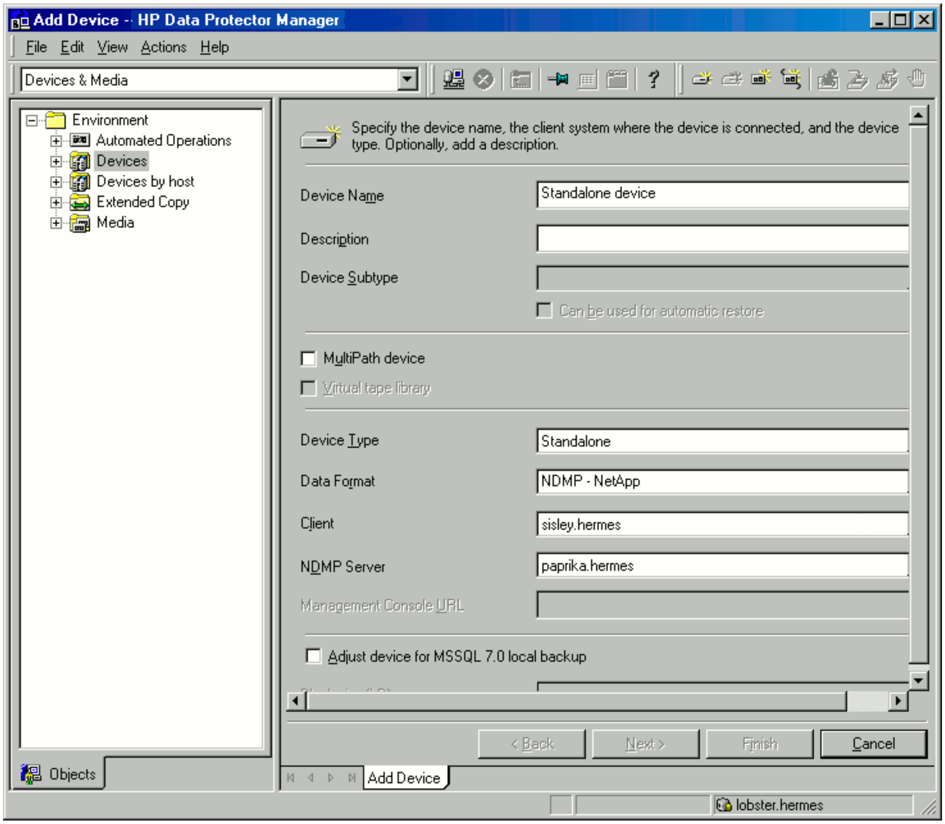

Specify the device name, the client system where the device is connected, and the device type. Optionally, add a description.

Device Name: Standalone device

Description:

Device Subtype:

□ Can be used for automatic restore

□ MultiPath device

□ Virtual tape library

Device Type: Standalone

Data Format: NDMP - NetApp

Client: sisley.hermes

NDMP Server: paprika.hermes

Management Console URL:

□ Adjust device for MSSQL 7.0 local backup

Figure 32 Configuring a standalone device

   Click **Next**.

4. Provide the SCSI address of the device. For information, see “Network appliance configuration” on page 86 and “EMC Celerra configuration” on page 87.

   Click **Next**.

5. Specify the media pool.

   To specify advanced device options, click **Advanced**. For information on supported block sizes, see Table 14 on page 88.

NOTE:

Multiplexing data steams is not supported by NDMP Server, limiting device concurrency to 1.

6. Click **Finish**.

Network appliance configuration

Before you begin

- Ensure that the NDMP Server is online.

Standalone tape devices and drives in a tape library

To get information about standalone tape devices (or drives in a tape library) connected to the NDMP Server system, run:

```
sysconfig -t
```

on the NDMP Server system. The SCSI address is written at the beginning of the output and consists of four parts. See Table 11 on page 86.

Table 11 Analyzing the drive's SCSI address

Parts	Description
`{n|u}`	`no rewind` and `unload/reload` respectively.[1]
`rst`	`Raw SCSI tape` (always present).
`{0	1	2	...}`	Device number.
`{l|m|h|a}`	Data density and compression.

[1] Data Protector supports only the `no rewind` devices.

Example

The output for a DLT 4000 drive is:

```
nrst0m - no rewind device, format is:42500 bpi 6.0GB
```

Library robotics

To get the SCSI address of the library robotics connected to the NDMP Server system, run:

```
sysconfig -m
```

on the NDMP Server system. The SCSI address consists of two parts. See Table 12 on page 87.

Table 12 Analyzing the library Robotics' SCSI address

Parts	Description
`mc`	`Media changer device` (always present).
`{0	1	2	...}`	Device number.

Example

The output for a DLT 4000 library is:

```
mc0
```

EMC Celerra configuration

Before you begin

- Ensure that the NDMP Server is online.

SCSI devices

To get information about SCSI devices (tape drives and library robotics) connected to the EMC Celerra NAS device:

1. Log in to the Celerra control station.
2. Run:

   ```
   server_devconfig server_name -list -scsi -all
   ```

Example

See Table 13 on page 88 for an example list of SCSI devices. `c2t210` and `c2t310` are the SCSI addresses of the drives in the tape library and `c2t010` is the SCSI address of the library robotics.

Table 13 Example of a list of SCSI devices

Name	SCSI Address	Device Type	Information
`jbox1`	`c2t010`	`jbox`	`ATL P1000 62200001.03`
`tape2`	`c2t310`	`tape`	`QUANTUM DLT7000 1624q$`
`ttape2`	`c2t210`	`tape`	`QUANTUM DLT7000 1624q$`

Block size

The integration supports variable tape block sizes. For limitations, see Table 14 on page 88.

Table 14 Supported block sizes

NAS Device	Block size range (KB)
ONTAP < 6.5.3	64
ONTAP ≥ 6.5.3	64 ≤ *Size* ≤ 256
Celerra	64 ≤ *Size* ≤ 256

Prerequisites

- Ensure that the NDMP Server is configured to support variable block size.

The recommended (default) block size is 64 KB. You can set any value between 64 KB and 1024 KB. If the set block size is not supported by the NAS device, and you start a backup, Data Protector displays an error and aborts the session.

NOTE:

Although the Data Protector media formatting completes successfully, that does not guarantee that the NAS device supports the set block size, and backup may still fail.

Limitations

- The device used for restore must have the same or greater block size than the one that was used for backup.
- ***Celerra only:*** Block size value should not be greater than the Celerra `readWriteBlockSizeInKB` parameter.

TIP:

To get the current value of the `readWriteBlockSizeInKB` parameter, run:

```
server_param server_3 -facility PAX -info
readWriteBlockSizeInKB -verbose
```

Backup

Limitations

- Only filesystem backup is supported.
- You cannot store an NDMP backup and a standard Data Protector backup on the same medium.
- Load balancing is not supported.
- Device concurrency is limited to 1.
- You cannot browse devices and filesystems.
- Only `Full` and `Incl` backup types are supported.
- Object copying, object mirroring, and media copying are not supported.
- By default, you cannot select more than 5 million files for backup.

  To enable higher values (up to 20 millions), set the `OB2NDMPMEMONLY omnirc` file variable to `0`. For more information, see "The NDMP specific omnirc file variables" on page 99.

- Once you have selected a directory, you cannot exclude any subdirectories or files from backup. Specifically, the following options are not supported:
  - Data Protector GUI: the `Trees/Filters` set of options: `Trees`, `Excludes`, `Skips`, and `Onlys`.
  - Data Protector `omnib` command: `-trees`, `-exclude`, `-skip`, and `-only`.

Before you begin

- Ensure that media to be used are formatted.
- ***NetApp only:*** Get information about filesystems exported from the NDMP Server system by running `exportfs`.

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 **Filesystem**, and click **Add Backup**.

3. Select a template. In Backup type, select **Data mover backup**. In Sub type, select **NDMP-NetApp** or **NDMP-Celerra**. See Figure 33 on page 91.

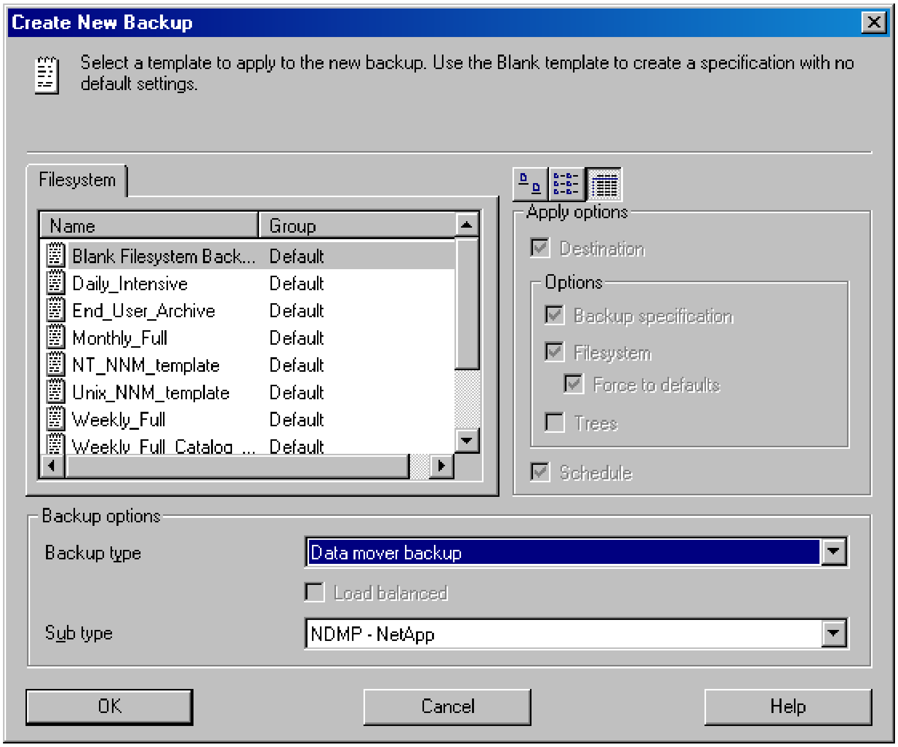

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Filesystem

Name	Group
Blank Filesystem Back...	Default
Daily_Intensive	Default
End_User_Archive	Default
Monthly_Full	Default
NT_NNM_template	Default
Unix_NNM_template	Default
Weekly_Full	Default
Weekly_Full_Catalog ...	Default

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

□ Trees

☑ Schedule

Backup options

Backup type: Data mover backup

□ Load balanced

Sub type: NDMP - NetApp

Figure 33 Selecting a backup template

Click **OK**.

4. Select the NDMP Server system you want to back up and click **Add/Remove**.

   In the Add/Remove Disk Mount Points dialog box, specify the filesystem mountpoints you want to back up: type the pathname of each directory in New mount point and click **Add**. See Figure 34 on page 92.

   Click **OK**.

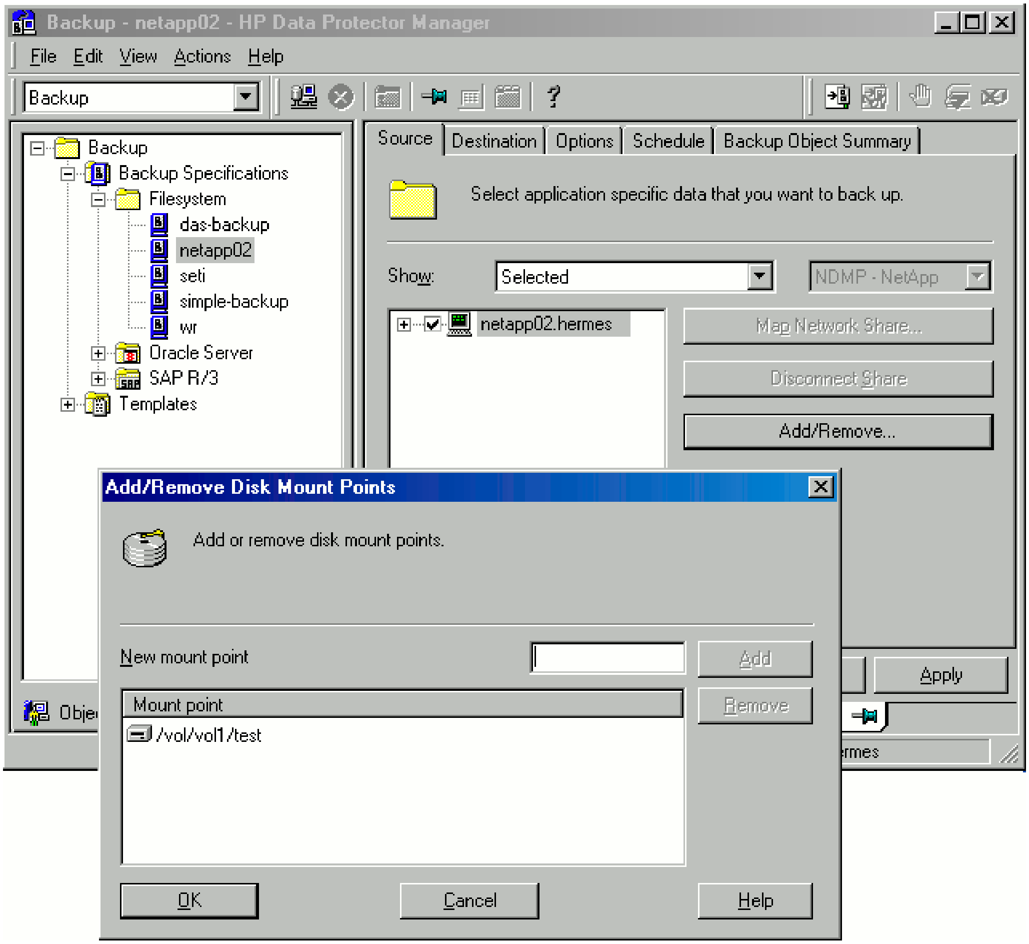

Backup - netapp02 - HP Data Protector Manager

Select application specific data that you want to back up.

Show: Selected

NDMP - NetApp

☑ netapp02.hermes

Add/Remove Disk Mount Points

Add or remove disk mount points.

New mount point

Mount point
/vol/vol1/test

Figure 34 Specifying the NDMP server mountpoints for backup (UNIX)

Click **Next**.

5. Select devices to use for the backup.

   To specify device options, right-click the device and click **Properties**.

   Click **Next**.

6. Set backup options.

   Click **Next**.

7. Optionally, schedule the backup.

   Click **Next**.

8. Review the summary of the backup specification

   To specify the NDMP NetApp options for a specific backup object, right-click the object, click **Properties**, and click the **NDMP** tab.

   For each object, you can specify a new user account that will override the user account specified in the Import NDMP Host dialog box, provided that the access rights are properly set on the NetApp or Celerra NAS device system.

   To set the NDMP environment variables, click **Advanced**. See Figure 35 on page 94. For more information, see "NDMP environment variables" on page 98.

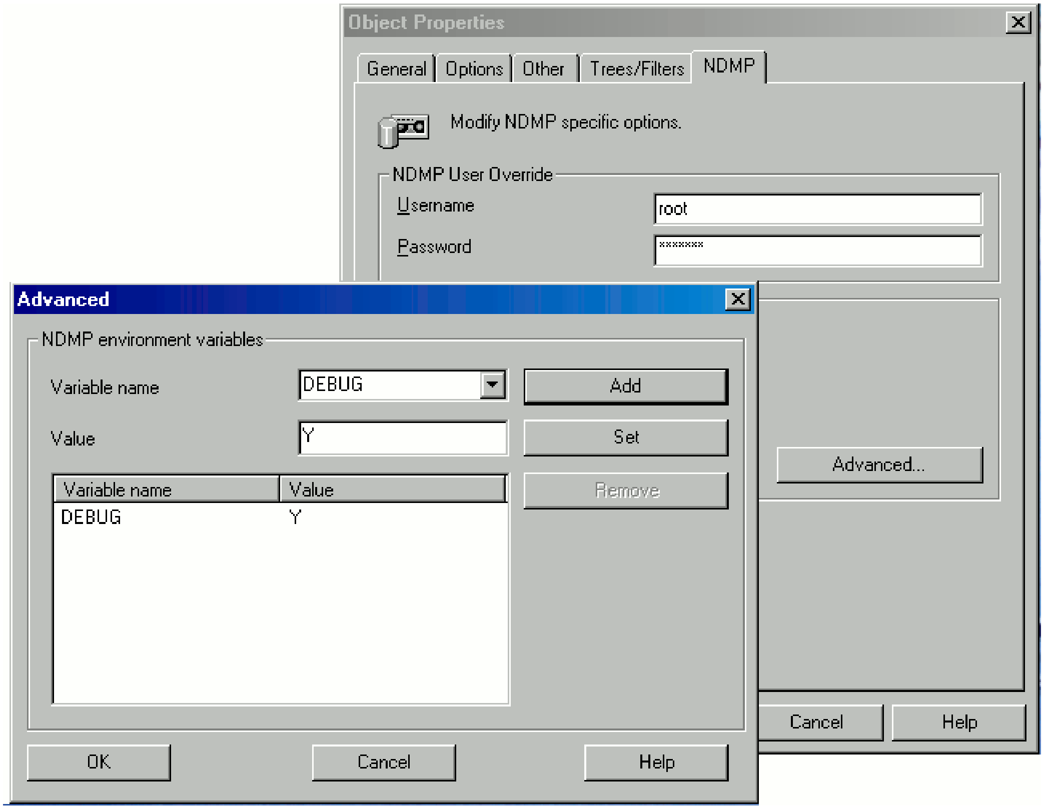

Modify NDMP specific options.

NDMP User Override

Username: root

Password: *******

Advanced...

NDMP environment variables

Variable name: DEBUG

Value: Y

Variable name	Value
DEBUG	Y

Figure 35 Specifying advanced NetApp options

   Click **Next**.

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

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.

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.

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

Restore

Restore filesystems using the Data Protector GUI or CLI.

Limitations

- Once you have selected a directory, you cannot exclude any subdirectories or files from restore. Specifically, the following options are not supported:
  - Data Protector GUI options: `Restore only` and `Skip`.
  - Data Protector CLI `omnir` command: `-only`, `-skip` and `-exclude`.

Restoring using the Data Protector GUI

1. In the Context List, select **Restore**.
2. In the Scoping Pane, expand **Filesystem**, expand the client with the data you want to restore, and then click the object that has the data.
3. In the Source page, browse for and select the objects you want to restore.
4. In the Destination page, specify restore destination for every selected object.

5. In the Options page, specify the NDMP Server system user account that will be used by Data Protector to connect to the NDMP Server system. This user must have permission to read from and write to the NDMP media.

   To specify the NDMP environment variables, click **Advanced** (Figure 36 on page 96). For more information, see "NDMP environment variables" on page 98.

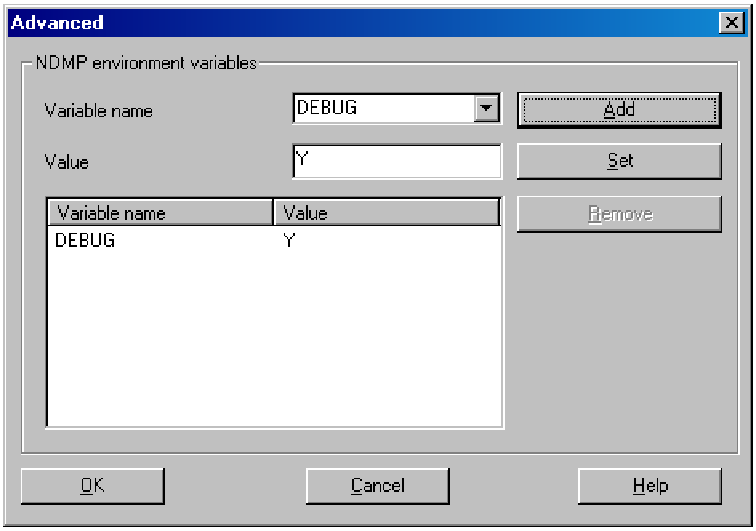

Advanced

NDMP environment variables

Variable name: DEBUG

Value: Y

Variable name	Value
DEBUG	Y

Figure 36 NDMP advanced restore options

6. In the Devices page, select devices you want to use for the restore.
7. Optionally, in the Media page, specify the media allocation priority.
8. Optionally, in the Copies page, specify the media set to restore from.
9. Click **Restore**.
10. In the Start Restore Session dialog box, click **Next**.
11. Specify **Report level** and **Network load**.
12. Click **Finish** to start the restore.

Direct access restore

Direct access restore is an optimized data recovery operation. Backed up data is accessed directly, in the middle of a tape.

This is achieved by partitioning backed up data into segments during backup and recording their start addresses.

During restore, Data Protector first computes which segment contains the requested file or directory, then locates the segment, and finally starts reading through it to locate the beginning of the file or directory.

Prerequisites

File history tracking must be turned on during the backup. On how to enable file history tracking, see "NDMP environment variables" on page 98.

To enable direct access restore, set the NDMP environmental variable `DIRECT` to `Y`. The procedure for the direct access restore is the same as for standard restore. The only difference is that you can browse for and select individual files and directories for restore. See Figure 37 on page 97.

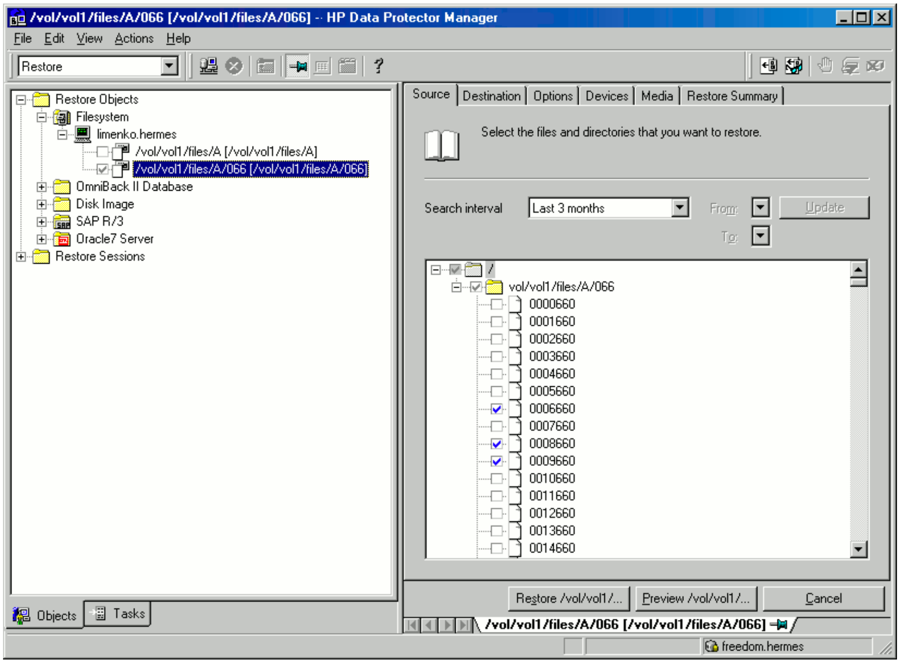

/vol/vol1/files/A/066 [/vol/vol1/files/A/066] -- HP Data Protector Manager

Select the files and directories that you want to restore.

Search interval: Last 3 months

Figure 37 Selecting NDMP Server Data for direct access restore

Limitations

- ***NetApp***:
  - Direct access restore of files is supported on ONTAP v6.1.x and higher.

- Direct access restore of directories is supported on ONTAP v6.4.x and higher. If you select both a directory and individual files from another directory, and start the restore, only the selected files are restored.

- ***Celerra***: Direct access restore of directories is not supported. If you select a directory and start the restore, only the directory without its contents is restored. To restore the whole directory, set `DIRECT=N`.

Restoring using another device

You can restore using a device other than that used for backup. For more information, see online Help.

NDMP environment variables

Set the NDMP environment variables for NetApp and Celerra NAS devices using the Data Protector GUI. See Figure 35 on page 94 and Figure 36 on page 96.

The following tables show the supported NDMP environment variables:

Table 15 NDMP variables for NetApp NAS device

Variable	Value	Function
`HIST`	`y/n`	Turns on/off file history tracking.
`DIRECT`	`y/n`	Enables direct access restore.
`LEVEL`	`0, 1, 2, ... 9`	Backup level (0=full).

Table 16 NDMP variables for Celerra NAS device

Variable	Value	Function
`HIST`	`y/n`	Turns on/off file history tracking.
`DIRECT`	`y/n`	Enables direct access restore.
`LEVEL`	`0, 1, 2, ... 9`	Backup level (0=full)
`BASE_DATE`	*32bit level32bit date*	Incremental backup based on a specific date.

Variable	Value	Function
OPTIONS	LK	Follow symbolic links.
	AT	Preserve access time.
	NT	Save NT attributes.
	MI/MD/MM	Restore collision policy for localization.

NOTE:

You can also set some NDMP environment variables using the omnirc file. For more information, see "The NDMP specific omnirc file variables" on page 99.

The NDMP specific omnirc file variables

On how to set the omnirc variables, see the online Help index: "omnirc options".

NOTE:

You can also set some variables using the Data Protector GUI. On how to do this, see Figure 34 on page 92, Figure 35 on page 94, and "NDMP environment variables" on page 98.

The GUI setting overrides the setting in the omnirc file.

The NDMP specific omnirc file variables are:

- **OB2NDMPFH** (Y/N)

  Default value: Y

  When set to Y, the NDMP Server file history tracking is turned on, which is a prerequisite for browsing and restoring individual files. However, this impacts the time needed for such a backup.

  This setting overrides the file history setting on the NDMP Server every time a backup is started.

- **OB2NDMPDIRECT** (Y/N)

  Default value: Y

When set to Y, Data Protector uses the direct access restore functionality, provided that the NDMP Server file history tracking was turned on during the backup.

- **OB2NDMPMEMONLY** (0/1)

  Default value: 1

  This variable defines how the NDMP Media Agent uses system resources.

  When set to 1, the NDMP Media Agent uses system physical memory only.

  When set to 0, the NDMP Media Agent stores part of the catalog in file history swap files. Set the variable to 0 whenever the number of files in the backup specification exceeds 5 millions. Consequently, the NDMP Media Agent can handle backups of up to 20 million files (in one backup specification), provided the system has enough resources.

  For example, to back up 20 million files, where 10% of the total number of backed up files are directories, with the average directory name consisting of 25 characters, and average filename consisting of 10 characters, you need approximately 1.9 GB of system memory and 2.8 GB of disk space.

  For optimal performance, select 10 million files and directories for backup.

  For more information on file history swap files, see the OB2NDMPFHFILEOPT variable description.

- **OB2NDMPCATQUESIZE**

  Default value: 5

  This variable sets the number of internal buffers that hold catalog information before storing it to file history swap files. By fine tuning the value, you can increase, to a certain extent, NDMP backup performance.

  When set to 5, the NDMP Media Agent can process up to 20 million files (in one backup specification), provided that enough system resources are available (approximately 1.9 GB of system memory and 2.8 GB of disk space).

  Set the variable to higher values if the number of files in the backup specification is less than 20 millions and enough system memory is available.

  To calculate memory allocation overhead in kilobytes, multiply the variable value by 512.

- **OB2NDMPFHFILEOPT**

  Default values:

  Windows: `*Data_Protector_home*\tmp, 32, 1024`

  UNIX: `/var/opt/omni/tmp, 32, 1024`

  This variable fine tunes file history swap files usage. It has three parameters that define the following:

  1. Pathname of the directory where the file history swap files are stored.

2. Maximum number of file history swap files, created by Data Protector on the NDMP client's disk.
3. Maximum size of a file history swap file (in MB).

The parameters are separated by commas. You can specify several sets of parameters. Use a semicolon to separate them.

Example

Windows: `C:\tmp, 32, 1024; D:\tmp\tmp_1, 10, 1024`

UNIX: `/tmp, 10, 1024; /var/tmp, 5, 60`

When the files in the first directory are full, the integration writes data to the files in the next specified directory. If the allocated disk space is used up during the backup, the backup fails.

File history swap files can increase in size significantly. Use the following formula to calculate approximate disk consumption:

EstConsumption = (*NumOfFiles* + *NumOfDirs*) x (136 + *AverageFileNameSize*)

where `NumOfFiles` is the number of backed up files and `NumOfDirs` is the number of backed up directories.

See the calculations in Table 17 on page 101 that presume that the number of directories is up to 10% of the total number of files, the average directory name length is 25 characters, and the average file name length is 10 characters.

Table 17 Approximate disk consumption by file history swap files

Number of backed up files and directories	Approximate disk consumption by file history swap files
5 Millions	0.7 GB
10 Millions	1.4 GB
20 Millions	2.8 GB

Media management

Data Protector media management is limited because data is backed up by NDMP Server in its specific data format.

Data Protector supports the following media management functionalities:

- Import and export of media.
- Media scan.
- Media initialization.
- Dirty drive detection.

Data Protector does not support the following media management functionalities:

- Verification of backed up data.
- Media copy.

For more information, see online Help.

Troubleshooting

This section lists problems you might encounter when using the Data Protector NDMP Server integration.

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

Before you begin

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

Problems

Problem

End of media

At the end of the backup, Data Protector starts storing the catalog to the media. The catalog size increases with the number of files backed up. Since Data Protector has no control over how much free space is left on the media, the `End of Media` error may occur during the writing of the catalog. This has no impact on future restore because the catalog is still stored in the IDB. However, the medium cannot be imported anymore.

Problem

Import of NDMP media failed

Action

Ensure that the drive used for importing NMDP media is connected to an NDMP Server system.

Problem

A tape remains in the drive after a successful drive scan

Action

Eject the tape manually and set the `OB2SCTLMOVETIMEOUT` `omnirc` file variable on the NDMP client to a higher value (for example, 360000 or higher).

On how to set the `omnirc` file variables, see the online Help index: “omnirc options”.

Problem

Data Protector was unable to set NDMP record size

Data Protector reports:

```
DP was unable to set NDMP record size. Reason for this might
be that NDMP server doesn't support specified record size.
Please check the release notes in order to determine which
record size is supported for your NDMP server.
```

Action

See “Block size” on page 88.

4 Integrating VMware Virtual Infrastructure and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector VMware ESX Server 3.x integration (**VMware integration**). It describes concepts and methods you need to understand to back up and restore the following **VMware objects**:

- Complete virtual machines, including virtual machine snapshot trees
- Individual virtual machine disks
- Virtual machine filesystems
- Virtual machine memory files
- Virtual machine disk images

Data Protector integrates with VMware Virtual Infrastructure to back up virtual machines online or offline. During backup, virtual machines can be actively used.

Data Protector offers the following backup methods:

- Snapshot
- Suspend
- VCB file
- VCB image

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

- Full
- Incr
- Differential

For details, see Table 22 on page 132.

From a **Snapshot**, **Suspend**, and **VCBimage** backup session, you can restore virtual machines to the original or different datacenter.

From a **VCBfile** backup session, you can restore filesystems to the original or different Windows virtual machine. You can also restore it to a separate Windows system, which does not need to be part of the VMware Virtual Infrastructure environment.

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

Integration concepts

Data Protector integrates with VMware Virtual Infrastructure through the Data Protector integration agent (`vmware_bar.exe`), which channels communication between the Data Protector Session Manager and the clients in the VMware environment.

Data Protector supports environments where ESX Server systems are managed through a VirtualCenter system (**VirtualCenter environment**) as well as environments with standalone ESX Server systems (**standalone ESX Server environment**). Data Protector also supports mixed environments, in which some of the ESX Server systems are managed through a VirtualCenter and some are standalone. You can even have multiple VirtualCenter systems in your environment, each managing its own set of ESX Server systems.

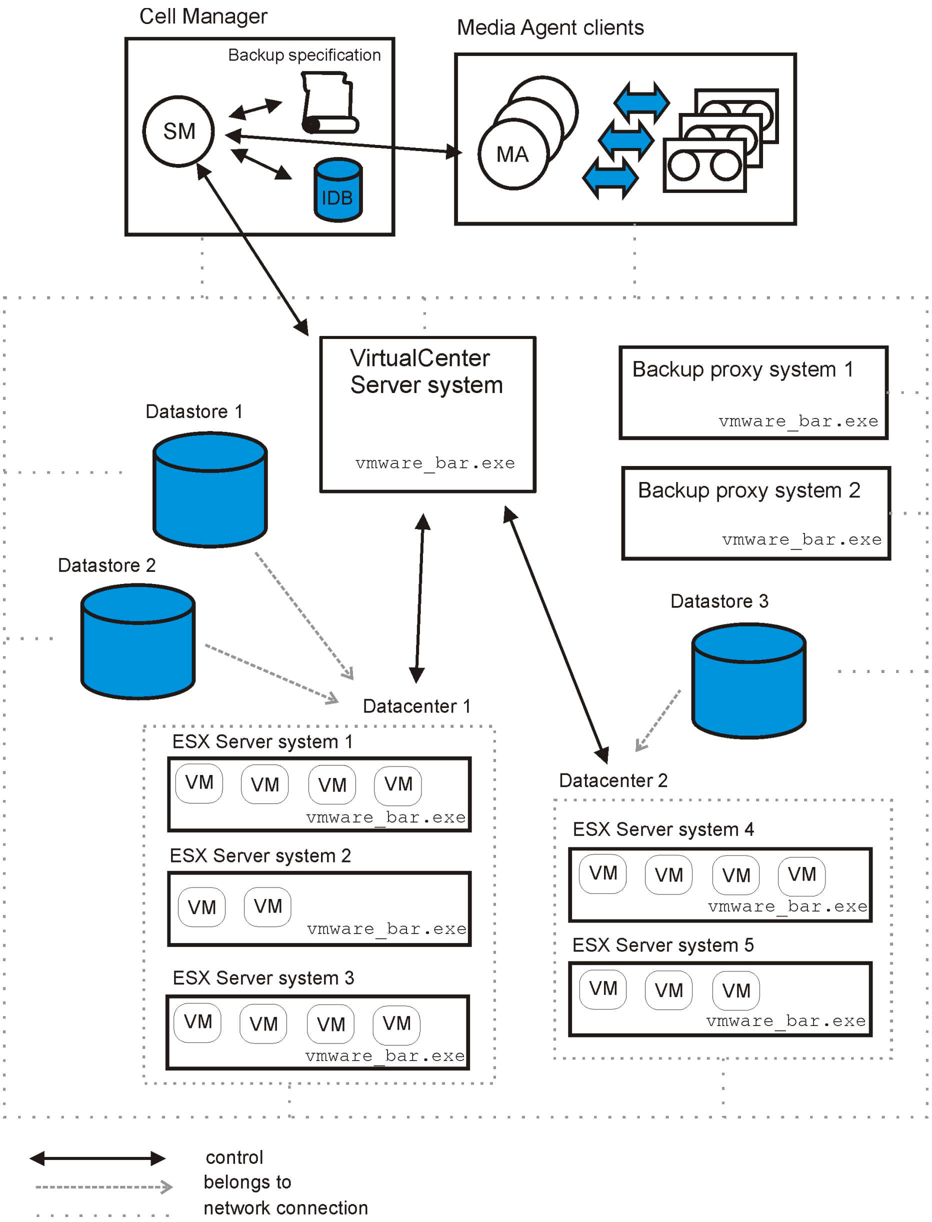

Figure 38 VirtualCenter environment

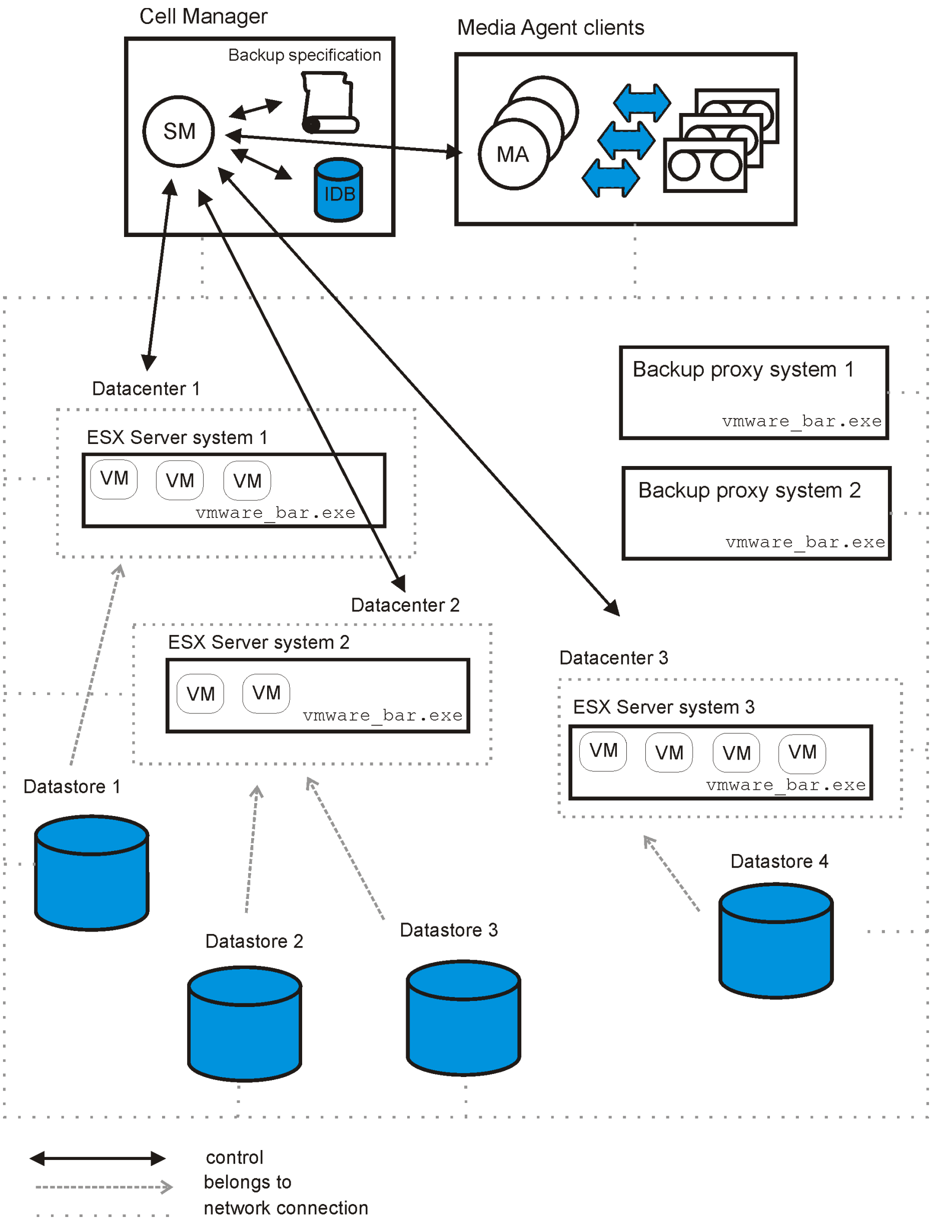

Figure 39 Standalone ESX Server environment

Table 18 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore
Backup Specification	A list of objects to be backed up, backup devices, and options to be used
IDB	The Data Protector Internal Database
vmware_bar.exe	Data Protector executable that enables data transfer between ESX Server and Data Protector media
ESX Server system	A modified Linux operating system that is capable of hosting multiple virtual machines.
VM	VMware specific term: Virtual machine. Virtualized x86 PC environment, in which a guest operating system and associated application software can run. Multiple virtual machines can run on the same ESX Server system concurrently.
Datastore	VMware specific term: Storage location for the virtual machine files. This can be a physical disk, a RAID, a SAN, or a partition on any of these.
Datacenter	VMware specific term: A logical unit consisting of one or more ESX Server systems. Each datacenter uses its own datastores.
MA	Data Protector General Media Agent

Table 19 briefly describes the VMware objects that you can back up and restore using the Data Protector VMware integration.

Table 19 VMware objects

VMware object	Description
Virtual machine	See description in the above legend. A virtual machine backup contains virtual machine configuration files, virtual machine disks, and, possibly, the virtual machine snapshot tree. For details on virtual machine snapshot trees, see "Snapshot method" on page 110.
Virtual machine disks	Each virtual machine has its own virtual machine disks, which are located on one or more datastores. You can back up individual virtual machine disks.

VMware object	Description
Virtual machine memory file	You can instruct Data Protector to save the memory of a running virtual machine to a virtual machine disk before the virtual machine is shut down and backed up. This virtual machine memory file can also be included in the backup.
Virtual machine filesystem	This is actually a NTFS filesystem from a Windows virtual machine. Filesystems from other operating systems cannot be backed up. You can choose to back up the complete filesystem tree or individual folders.

Backup methods

Data Protector offers four different backup methods:

- **Snapshot** method
- **Suspend** method
- **VCB file** and **VCB image** (consolidated backup methods)

Snapshot method

A virtual machine snapshot is an operation that memorizes the current state of a virtual machine, including the state of the data on the virtual machine disks as well as whether virtual machine is powered on, powered off, or suspended. Once a snapshot is created, you can revert the virtual machine to this point in time whenever needed.

During a **Snapshot** backup method, Data Protector first creates a snapshot to put the virtual machine into a consistent state and then the virtual machine data is backed up. The snapshot is named `_DP_SNAP_`. Data Protector created snapshots (**DP snapshots**) are distinguished from each other by a unique description that contains a timestamp. Avoid using this name for snapshots that you create yourself.

How DP snapshots are handled, whether they are kept or deleted after or during the backup, depends on the mode that you select. You can choose among the following snapshot handling modes:

- Disabled
- Single
- Mixed

The selected mode determines whether you can run only full backups or you can create backup chains consisting of a full backup followed by differential or/and incremental backups. The selected mode is directly connected with the number of DP snapshots that are left on the snapshot tree for backup purposes.

IMPORTANT:

Note that the following actions break your backup chain:

- Deletion of snapshots
- Reversion to a snapshot
- Creation of your own snapshots
- Change in the snapshot handling mode
- Restore of virtual machines

Consequently, you must run a full backup again to start a new backup chain. Otherwise, your subsequent incremental and/or differential sessions fail.

Figure 40 on page 112 shows how a virtual machine snapshot tree may look like at a certain point in time.

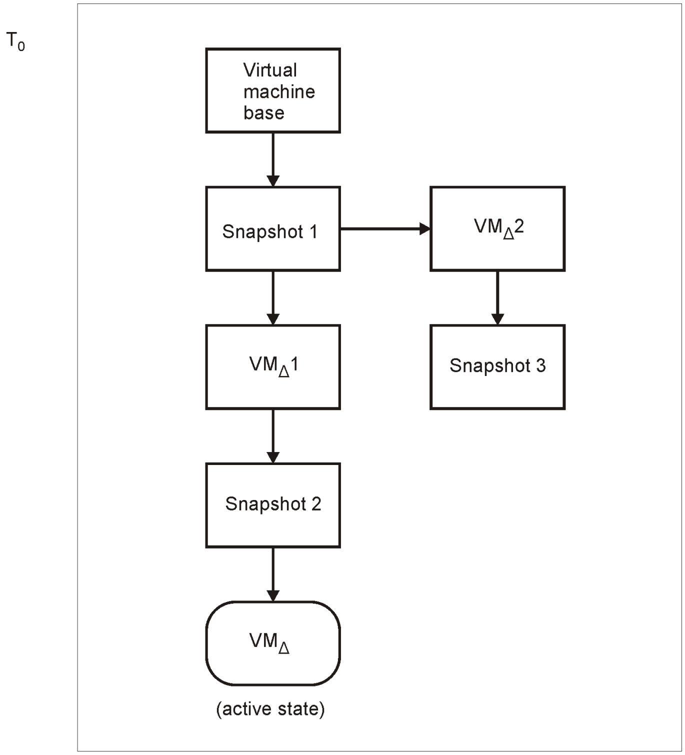

Figure 40 Branched snapshot tree

In its simplest form, a snapshot tree may look like shown below.

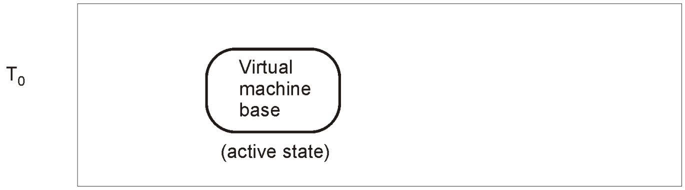

Figure 41 Simple snapshot tree

The complexity of your snapshot tree at the beginning of a backup chain does not impact how DP snapshots are handled during backup sessions. For this reason, this simple snapshot tree will be used as a starting point in the following illustrations.

Snapshot mode: disabled

In the **Disabled** mode, you can run only full backups. When the backup completes, the DP snapshot is deleted. See Figure 42 on page 113.

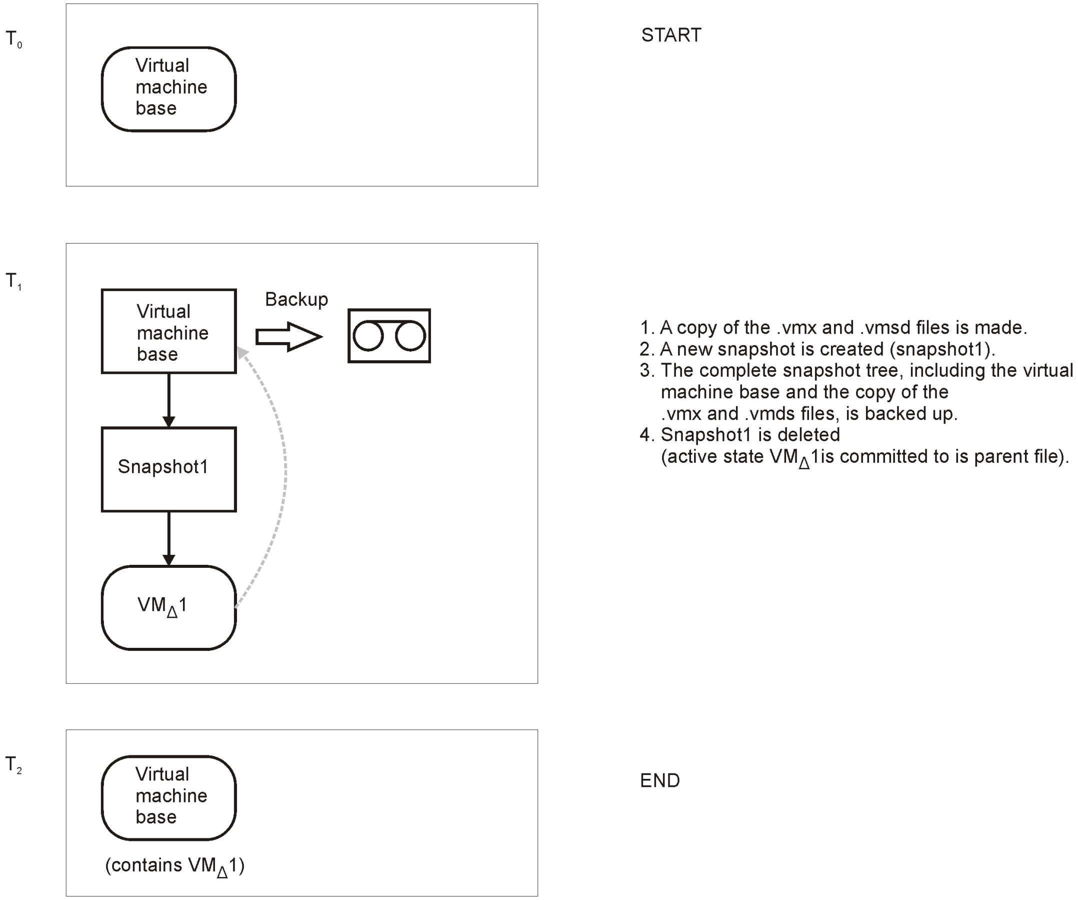

1. A copy of the .vmx and .vmsd files is made.
2. A new snapshot is created (snapshot1).
3. The complete snapshot tree, including the virtual machine base and the copy of the .vmx and .vmds files, is backed up.
4. Snapshot1 is deleted (active state VM_{Δ}1is committed to is parent file).

Figure 42 Full backup (disabled mode)

Snapshot mode: single

In the **Single** mode, you can run full backups in combination with either incremental or differential backups. However, you cannot mix incremental and differential backups in the same backup chain.

After the first full backup, you end up with one DP snapshot. A full backup in the Single mode progresses in the same way as a full backup in the Disabled mode, with the exception that the DP snapshot is not deleted at the end. A subsequent differential backup is shown in the following figure.

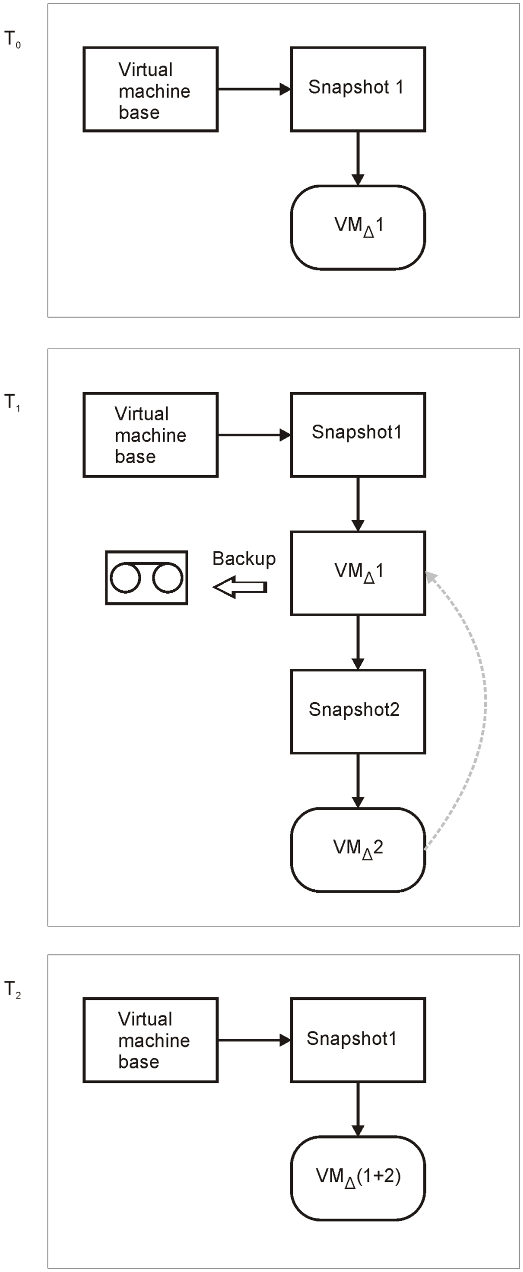

START

1. A copy of the .vmx and .vmsd file is made.
2. A new snapshot is created (snapshot2).
3. Changes made between snapshot1 and snapshot2, including the copy of the .vmx and .vmsd files, are backed up.
4. Snapshot2 is deleted ($\mathrm{VM}_{\Delta}2$ is committed to $\mathrm{VM}_{\Delta}1$).

END

Figure 43 Differential backup (single mode)

A backup chain consisting of a full backup followed by incremental sessions progresses in the same way, with the exception that `snapshot1` is deleted instead of `snapshot2` (see ???) at the end of an incremental session.

When you start a new full backup (a new backup chain), first the DP snapshot is deleted (active state is committed to its parent file) and then the session progresses in the same way as already described.

Snapshot mode: mixed

In the **Mixed** mode, you can run full, incremental, and differential backups, in all possible combinations. In this mode, up to two DP snapshots remain on a snapshot tree for backup purposes.

After the first full backup, you end up with one DP snapshot. The session progresses in the same way as a full backup in the **Single** mode. After a subsequent differential or incremental backup, you end up with two DP snapshots. The session progresses in the same way as a differential or incremental backup in the **Single** mode, with the exception that the DP snapshot is not deleted at the end. The progress of an additional differential and incremental session is shown below.

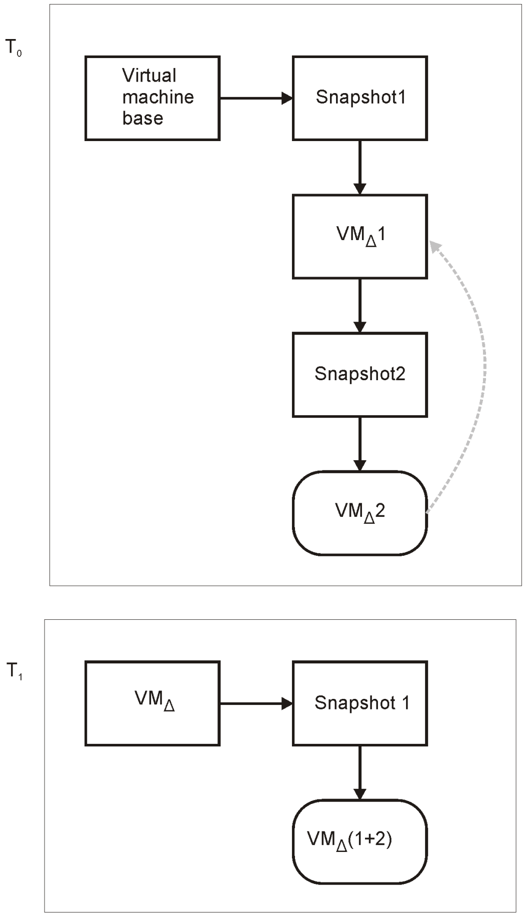

START

1. A copy of the .vmx and .vmsd files is made.
2. Snapshot2 is deleted (VM$_{\Delta}$2 is committed VM$_{\Delta}$1).

Figure 44 Differential backup – 1st part (mixed mode)

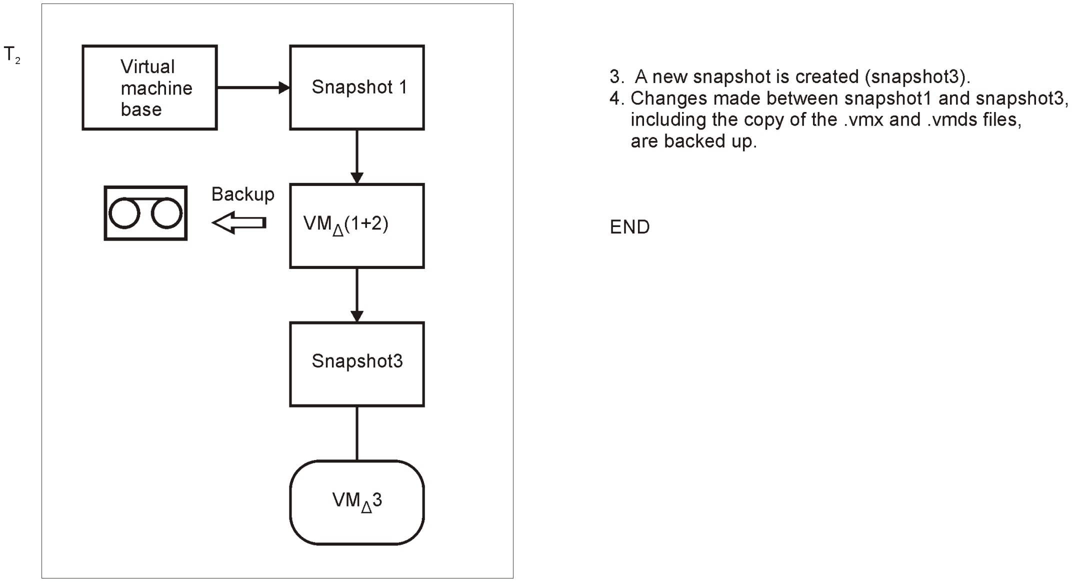

3. A new snapshot is created (snapshot3).
4. Changes made between snapshot1 and snapshot3, including the copy of the .vmx and .vmds files, are backed up.

END

Figure 45 Differential backup – 2nd part (mixed mode)

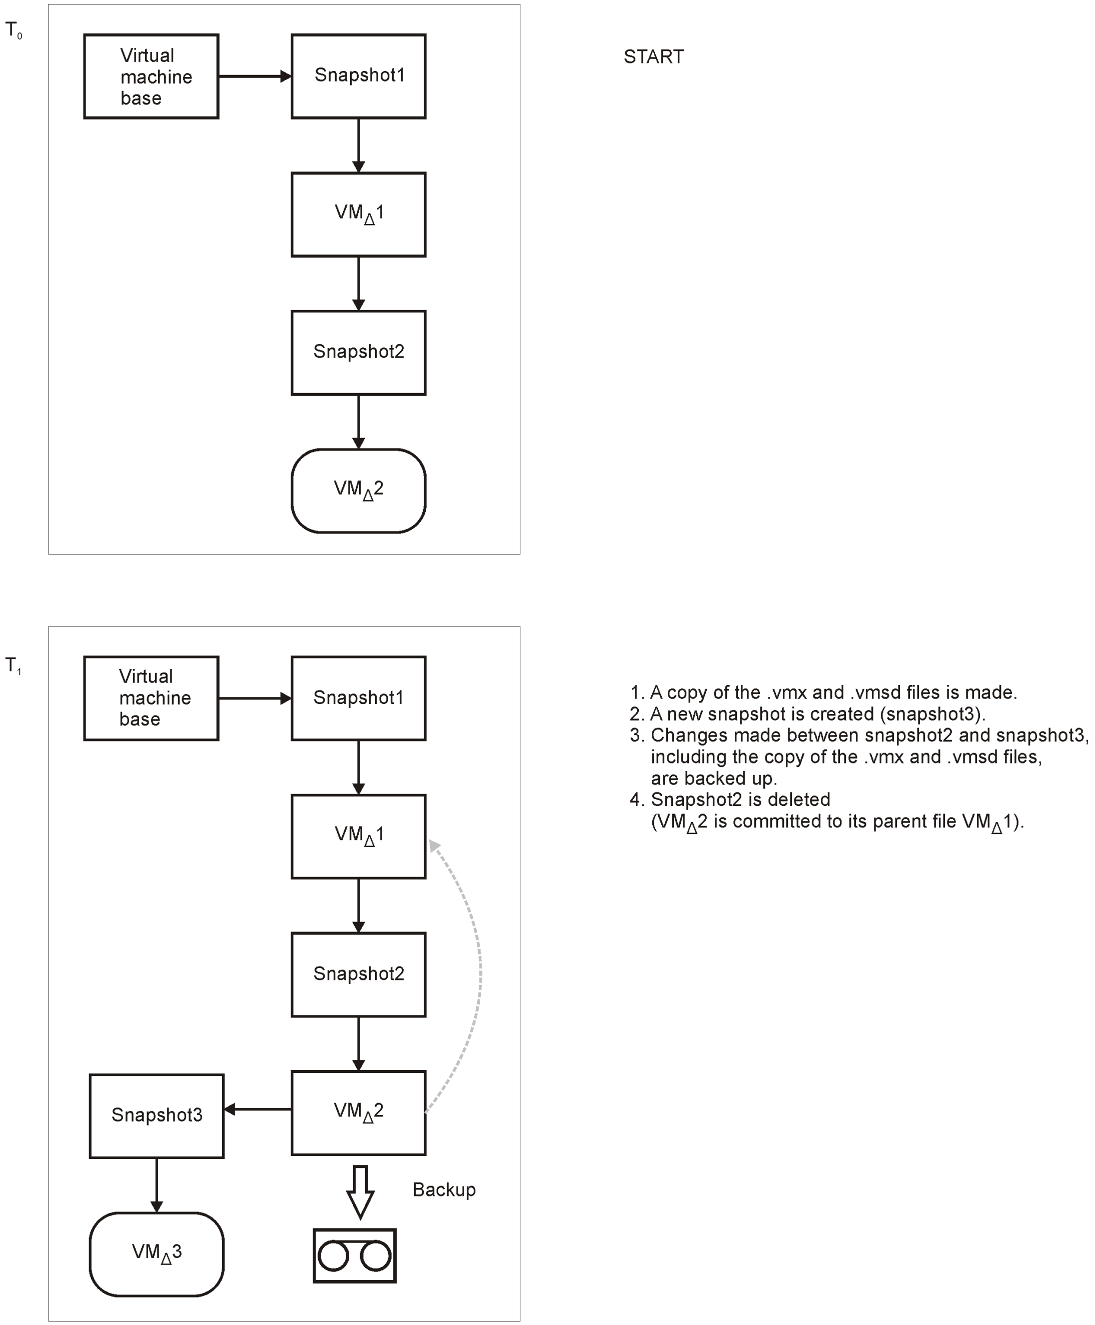

START

1. A copy of the .vmx and .vmsd files is made.
2. A new snapshot is created (snapshot3).
3. Changes made between snapshot2 and snapshot3, including the copy of the .vmx and .vmsd files, are backed up.
4. Snapshot2 is deleted ($VM_{\Delta}2$ is committed to its parent file $VM_{\Delta}1$).

Figure 46 Incremental backup – 1st part (mixed mode)

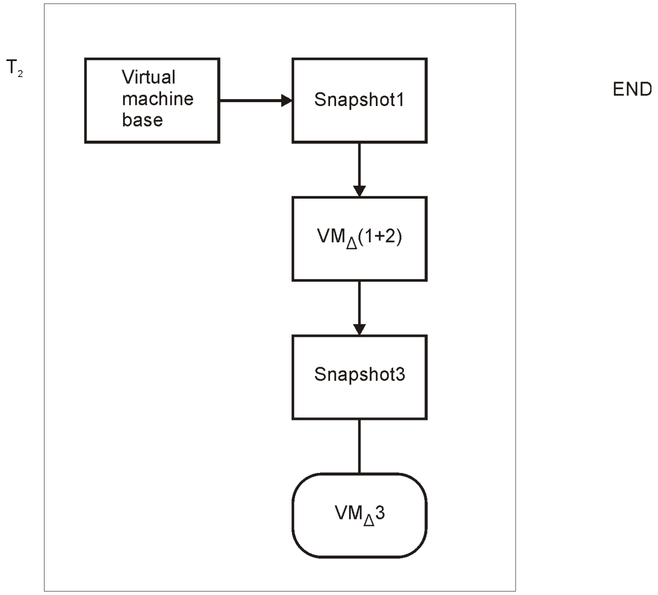

Figure 47 Incremental backup– 2nd part (mixed mode)

When you start a full backup again, first all DP snapshots are deleted (changes are committed to parent files) and then the session progresses in the same way as already described.

Restore

When you start a restore from a **Snapshot** backup session, Data Protector automatically restores the complete restore chain, consisting the last full backup, differential backup (if it exists) and all incremental backups since the last full or differential backup up to the selected backup session. Of course, all these sessions are **Snapshot** backup sessions.

Whenever a chain of sessions is restored, some of the restored snapshots are automatically deleted (changes are committed to parent files), since, at the end, the virtual machine snapshot tree must be in exactly the same state as at the time of the last incremental or differential backup. Otherwise, the virtual machine cannot be registered and powered-on. Restore session messages will inform you of the automatic commit operations. Note that the commit operations also reduce disk space usage.

If you want to restore individual virtual machine disks, you need to do some additional steps to be able to register and power-on the restored virtual machine.

Suspend method

During a **Suspend** backup method, the memory of a running virtual machine is saved to a virtual machine disk and then the virtual machine is shut down. Once the virtual machine is shut down, its files are backed up and, after the backup, the virtual machine resumes its original state (it is put online again).

This method supports full, incremental, and differential sessions. During an incremental or differential backup, only those files are backed up whose modification time has changed since the last backup or last full backup.

Restore

When you start a restore from a **Suspend** backup session, Data Protector automatically restores the complete restore chain, consisting of the last full backup, differential backup (if it exists) and all incremental backups since the last full or differential backup up to the selected backup session. Of course, all these sessions are **Suspend** backup sessions.

When you restore individual virtual machine disks, you need to do some additional steps to be able to register and power-on the restored virtual machine.

VCB backup method

For a **VCBimage** or **VCBfile** backup method, you need at least one backup proxy system in your environment and your virtual machines must be located on SAN storage. The **VCB** backup method enables you to back up only the current state of the virtual machine. Information about the snapshot tree and the changes made on non-active snapshot branches are not included in the backup. The **VCBfile** backup method enables you to back up Windows NTFS filesystems (Windows virtual machines). Other filesystem types cannot be backed up. The **VCBimage** backup method enables you to back up virtual machine disk images.

Backup flow

1. The Data Protector Cell Manager starts the `vmware_bar.exe` agent on the VirtualCenter system (VirtualCenter environment) or ESX Server system (standalone ESX Server environement), providing a list of virtual machines to be backed up.
2. The `vmware_bar.exe` agent starts the `vmware_bar.exe` agent on the corresponding backup proxy system. The VMware `vcbMounter` tool, which is activated by the Data Protector agent, creates a virtual machine snapshot and mounts the virtual machine disks to the backup proxy system.

3. **VCBfile**: On the backup proxy local disk, links to the virtual machines disks are created and the virtual machine filesystems are backed up to Data Protector media.

   VCBimage: Virtual machine disks are copied to the backup proxy local disk (virtual machine disks are **exported**). From there, the virtual machine disk images are backed up to Data Protector media.

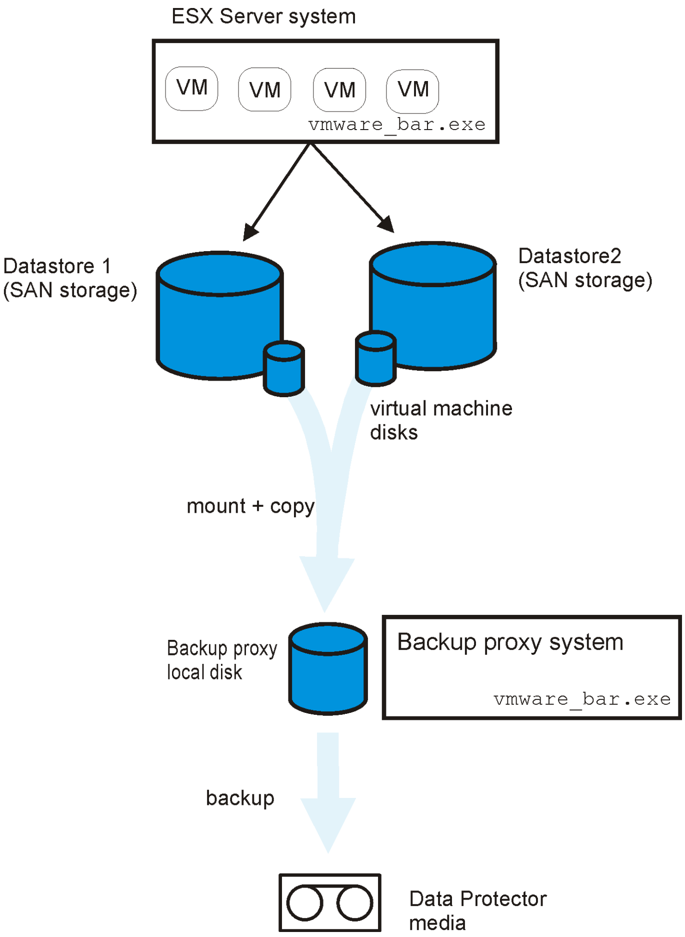

Figure 48 VCB image method

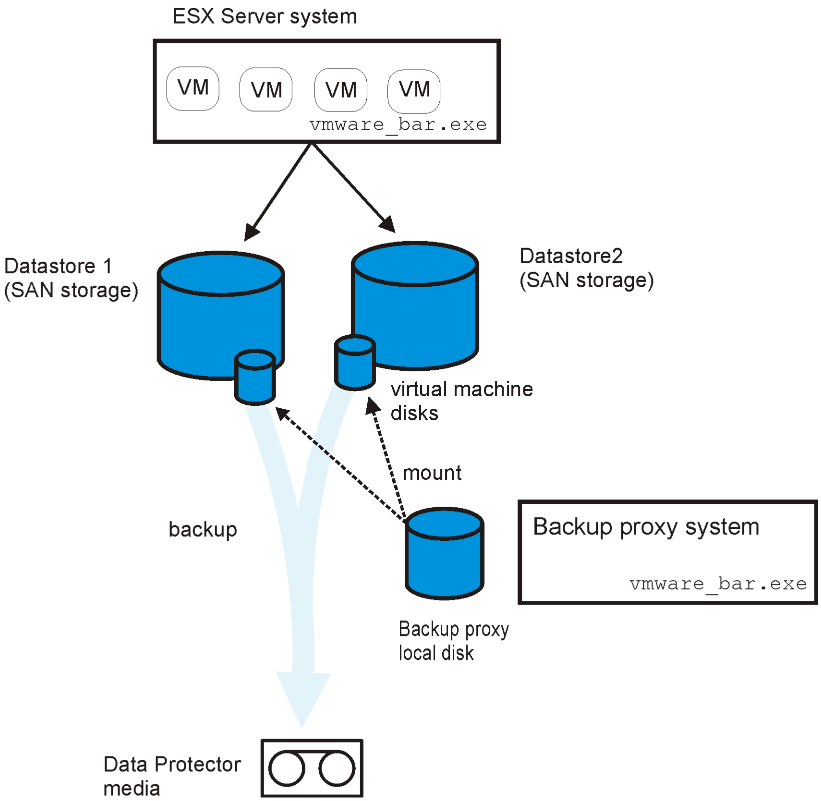

Figure 49 VCB file method

Restore

When you start a restore from a **VCB file** backup session, Data Protector automatically restores the complete restore chain consisting of the last full backup, differential backup (if it exists) and all incremental backups since the last full or differential backup up to the selected backup session. These sessions are **VCB file** backup sessions. Before the actual restore starts, the virtual machines to be restored are put offline. For a **VCB file** restore, you must specify for each virtual machine filesystem separately whether you want to restore it to the original or a different virtual machine. You can also restore it to a completely different Windows system. In the latter case, the destination system must be part of the Data Protector cell and have the `VMware Integration` component installed. If you restore to the virtual machine directly, the `VMware Integration` must be installed also inside the virtual machine itself.

When you restore from a **VCB image** backup session, you can restore the virtual machine to the original or different datacenter. If you want to restore individual virtual machine disks, you need to do some additional steps to be able to register and power-on the restored virtual machine.

Method overview

Backup method	Needs a backup proxy system	Supported backup types			Backs up all snapshot branches
		Full	Differential	Incremental	
Suspend		✓	✓	✓	✓
Snapshot		✓	✓	✓	✓
VCB image	✓	✓			
VCB file	✓	✓	✓	✓	

Clusters

Data Protector supports environments in which ESX Server systems or VirtualCenter systems are in a cluster.

ESX Server systems in a cluster

ESX Server systems can be in a cluster only in VirtualCenter environments. There are two different cluster types:

- **VMware load balancing cluster**

  VMware Distributed Resource Scheduler (DRS) monitors workload on ESX Server systems that belong to the same VMware load balancing cluster. As a result, virtual machines running on ESX Server systems with higher load are moved to those with lighter load with the help of the VMware VMotion component. For example, if you add a new ESX Server system to the cluster, some of the running virtual machines migrate there.

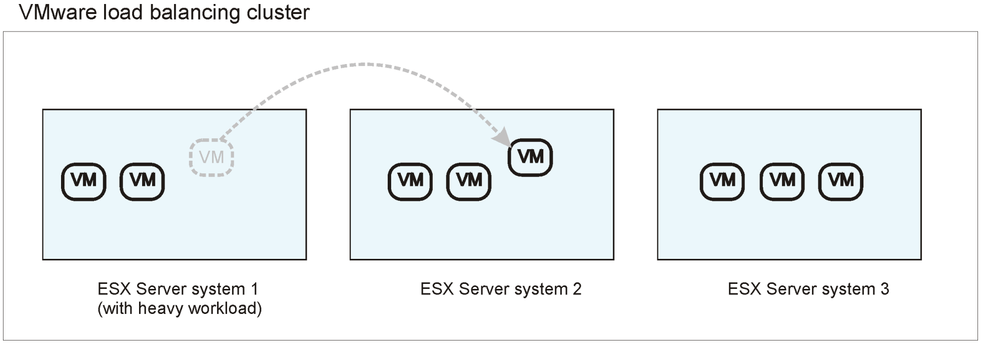

Figure 50 VMware load balancing cluster

- **VMware high availability cluster**

  VMware high availability cluster is a VMware service that monitors heartbeats of ESX Server systems. If needed, it restarts virtual machines on another ESX Server systems within the same cluster.

  Installation requirements: The Data Protector `VMware Integration` component must be installed on all ESX Server systems configured in the cluster.

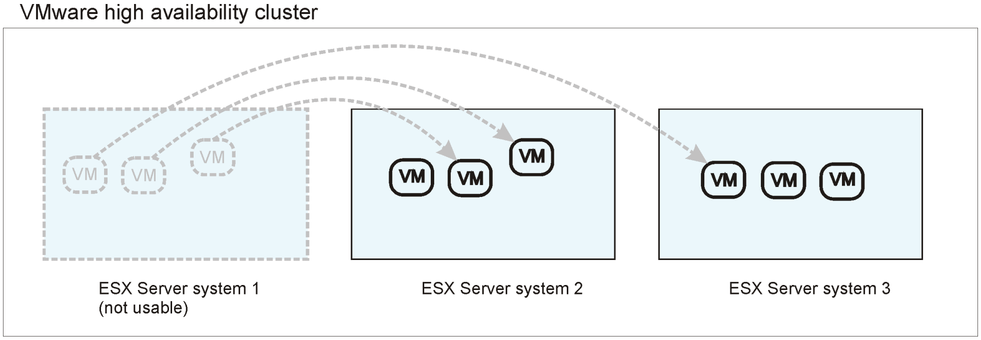

Figure 51 VMware high availability cluster

Installation requirements: The Data Protector `VMware Integration` component must be installed on all the ESX Server systems in the cluster. Note that the migration cannot be successful if the virtual machine files are stored on local ESX Server disks. The virtual machines files must be stored on shared SAN datastores.

VirtualCenter systems in a cluster

VirtualCenter systems can be in a cluster with the help of the Microsoft Cluster Service functionality.

Installation requirements: The Data Protector `VMware Integration` component must be installed on both VirtualCenter cluster nodes.

Configuring the integration

You need to configure VirtualCenter systems (VirtualCenter environment), ESX Server systems (standalone ESX Server environment), and VMware users.

Prerequisites

- Ensure that you have correctly installed and configured VMware environment.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on installing, configuring, and using VMware Virtual Infrastructure, see VMware documentation. For the **VCBfile** and **VCBimage** backup methods, ensure that you have at least one backup proxy system configured in your environment. For these methods, your virtual machines must be configured on SAN storage.
- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HP Data Protector installation and licensing guide*.

  The `VMware Integration` component must be installed on all ESX Server systems from which you plan to back up virtual machines, as well as on VirtualCenter systems if they exists. For consolidated backup methods, the integration component must be installed also on the backup proxy systems. To be able to restore filesystems from a **VCBfile** backup to virtual machines, install the integration component also inside the virtual machines.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the VMware Virtual Infrastructure and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every ESX Server system and VirtualCenter system in the VMware environment.

Configuring clusters

ESX Server systems in a cluster

If your ESX Server systems are in a high availability cluster, open the cluster settings dialog box in the Virtual Infrastructure client and set the **Allow virtual machines to be powered on even if they violate availability constrains** option ON.

Configuring VMware users

1. Identify or configure the following users:

   Table 20 VMware users

VirtualCenter users (VirtualCenter environment)	For each VirtualCenter system, identify a Windows operating system user that installed the VirtualCenter.
ESX Server users (standalone ESX Server environment)	For each standalone ESX Server system, identify an operating system user that has read, write, and execute permissions on the related datastores.

2. Add all the users to the Data Protector `admin` or `operator` user group.

   For details on adding users to Data Protector groups, see the online Help index: "adding users".

NOTE:

If your VirtualCenter system is in a cluster, configure the users on both nodes.

Configuring VirtualCenter or ESX Server systems

You need to provide the following login information for each VirtualCenter system and standalone ESX Server system:

- Username
- Password
- Web root (optional)

- Port (optional)

Once the information is provided, the connection is first tested. If it is successful, the login information is saved in a VirtualCenter specific or ESX Server specific configuration file on the Cell Manager.

NOTE:

If the connection with new login information fails, it is not saved in the configuration file. Note that you are not informed about this failure if the configuration file already contains login information. Consequently, this old login information will be used in future sessions. To check whether Data Protector has accepted new login information, ...TBD.

To configure VirtualCenter systems or standalone ESX Server systems, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You configure VirtualCenter or ESX Server system when you create the first backup specification for this VirtualCenter or ESX Server system. Specify the information in the Configure VMware dialog box.

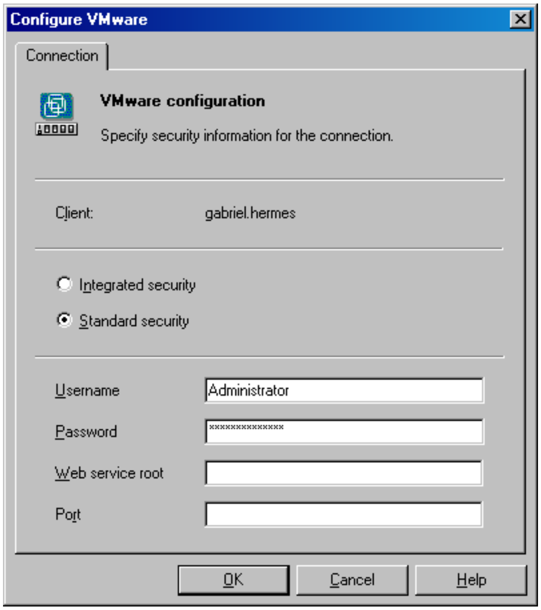

Configure VMware

Connection

VMware configuration

Specify security information for the connection.

Client: gabriel.hermes

○ Integrated security

● Standard security

Username: Administrator

Password: ***************

Web service root:

Port:

Figure 52 Configuring a VirtualCenter system

Select either **Integrated security** or **Standard security**. If **Integrated security** is selected, Data Protector connects to the VirtualCenter system or ESX Server system using the login information from the file `/etc/vmware/backuptools.conf`. If you want to specify the login information manually, select **Standard security**.

- **Username** and **Password**: Specify a VirtualCenter or ESX Server operating system user account that has TBD.
- **Web service**: Optionally, change the web service entry point URI. Default: `/sdk`
- **Port**: Optionally, change the TCP/IP port number on which the web service server should be listening.

  Default: `443` (SSL encrypted HTTP), `80` (unencrypted HTTP).

  By default, the SSL encrypted HTTP is used. To use unencrypted HTTP, set the Data Protector `omnirc` variable `OB2_VMWARE_HTTP` to `1`. On how to set the `omnirc` variable, see the online Help index: “omnirc options”.

These parameters are then saved in the VirtualCenter or ESX Server specific configuration file on the Cell Manager.

Using the Data Protector CLI

From the system:

VirtualCenter system: `Data_Protector_home\bin`

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

run:

For **Integrated security**:

```
util_vmware -config -security 1
```

For **Standard security**:

```
util_vmware -config -security 0 —user VMware_username
—password VMware_password [-webroot web_service_root] [-port
web_service_port]
```

Checking the configuration

To verify the connection, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can verify the connection to the VirtualCenter system or ESX Server system after you have created at least one backup specification.

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **VMware**. Click the backup specification for the VirtualCenter or ESX Server system to be checked.
3. Right-click the VirtualCenter or ESX Server system and click **Check configuration**.

Using the Data Protector CLI

From the system:

VirtualCenter system: `Data_Protector_home\bin`

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

run:

```
util_vmware -chkconf [-instance datacenter]
```

Backup

You can back up the following VMware objects:

- Complete virtual machines, including virtual machine snapshot trees
- Individual virtual machine disks
- Virtual machine filesystems
- Virtual machine memory files

The integration provides four backup methods:

Table 21 Backup methods

Backup method	How is backup consistency achieved?	What can be backed up?
Snapshot	Data Protector creates a virtual machine snapshot.	• Complete virtual machines • Virtual machine memory files • Individual virtual machine disks
Suspend	Data Protector suspends the virtual machine.	
VCB file	A backup proxy system, which is involved in the backup, creates a virtual machine snapshot.	Virtual machine filesystems
VCB image		Virtual machine disk images

For details, see the .

The integration provides backups of the following types:

Table 22 Backup types

Full	Backs up complete virtual machines.
Incr	Backs up changes made to virtual machines since the last backup of any type. Not available for the **VCBimage** backup method.
Differential	Backs up changes made to virtual machines since the last full backup. Not available for the **VCBimage** backup method.

Considerations

- **Datacenter pathnames:** In VirtualCenter environments, the length of a datacenter pathname should not exceed 79 characters. Suppose you created a folder `My datacenters` in the root folder of the `Hosts & Clusters` inventory view mode, using the Virtual Infrastructure Client interface, and then you created a datacenter `Datacenter1` within that folder. In this case, the datacenter pathname is `/My datacenters/Datacenter1`, which is acceptable because it consists of only 27 characters. The same logic applies when you create folders and datacenters in any other inventory view mode.

In standalone ESX Server environments, datacenter pathnames cannot exceed 79 characters because they are always `/ha-datacenter`.

- **Backup specification names:** The name under which a backup specification is saved should not contain double quotes. Otherwise, you encounter problems when you try to open the backup specification.
- **Concurrent sessions:** Backup sessions that use the same devices or back up the same datacenter cannot run concurrently.

Configuring virtual machines

Configure virtual machines to provide additional information on how to perform various backup methods.

For the **Snapshot** method, specify how to handle virtual machine snapshots that are created during backup. Note that not all snapshot handling modes support incremental and differential backups.

For the **VCBfile** and **VCBimage** backup methods, specify which backup proxy system and mount points should be used to back up filesystems and virtual machine disk images.

The **Suspend** backup method has no specifics.

You can configure each virtual machine separately or all together. Configuration settings for virtual machines of the same datacenter are saved in a separate configuration file on the Cell Manager. The file is used in all backup sessions involving this particular datacenter, even if different backup specifications are used.

To configure virtual machines, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You configure virtual machines when you create or modify a backup specification. In the Source page of a backup specification, right-click the client system at the top or any of the virtual machines listed below and click **Configure Virtual Machines**.

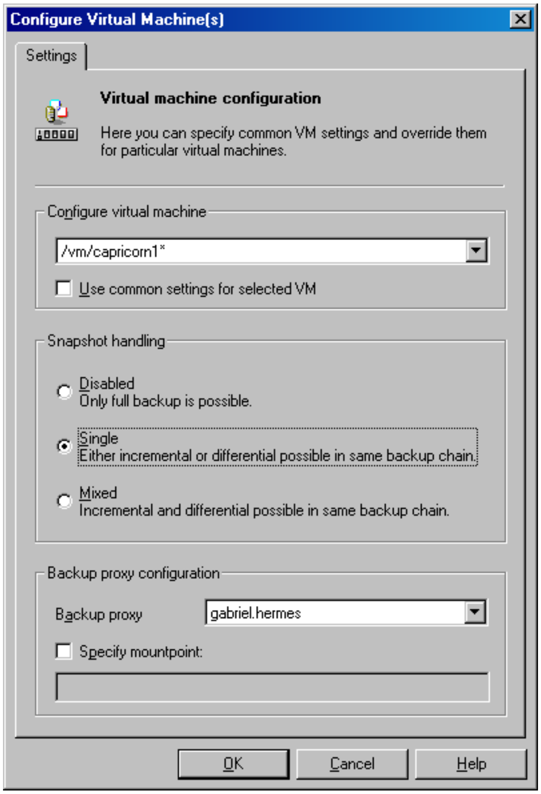

Configure Virtual Machine(s)

Settings

Virtual machine configuration

Here you can specify common VM settings and override them for particular virtual machines.

Configure virtual machine

/vm/capricorn1*

☐ Use common settings for selected VM

Snapshot handling

○ Disabled
Only full backup is possible.

● Single
Either incremental or differential possible in same backup chain.

○ Mixed
Incremental and differential possible in same backup chain.

Backup proxy configuration

Backup proxy: gabriel.hermes

☐ Specify mountpoint:

Figure 53 Configuring virtual machines

In the Configure Virtual Machine(s) dialog box, specify the following options:

Table 23 Virtual machine options

Configure virtual machine	Select whether you want to specify common virtual machine settings (**Common VM Settings**) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.	
	Use common settings for selected VM	Available only if a virtual machine is selected. Set this option ON to apply the common virtual machine settings for the selected virtual machine. Default: ON
	Use default settings	Available only if **Common VM Settings** is selected. Select this option to set default values for the common virtual machine settings.

		Default: ON
Snapshot handling (applicable for the **Snapshot** backup method)	**Disabled** (default)	This mode supports only full backups. The virtual machine snapshot that is created during backup is deleted at the end of the session. For details, see “Snapshot mode: disabled” on page 113.
	Single	This mode supports full, incremental, and differential backups. However, you cannot mix incremental and differential backups in the same backup chain. A backup chain may consist of a full backup followed by differential sessions, or of a full backup followed by incremental sessions. Data Protector keeps one DP snapshot for backup purposes. For details, see “Snapshot mode: single” on page 113.
	Mixed	This mode supports full, incremental, and differential backups, in all combinations. Data Protector keeps up to two DP snapshots for backup purposes. For details, see “Snapshot mode: mixed” on page 116.
Backup proxy configuration (applicable for the **VCBfile** and **VCBimage** backup methods)	**Backup proxy**	Select a backup proxy system to be used for VCB backup methods. Note that Data Protector lists all systems that have the `VMware Integration` component installed, including those that are not backup proxy systems.
	Specify mountpoint	During a VCB backup session, virtual machine disks are mounted to a backup proxy local disk. Select this option to specify a different mount point directory on the backup proxy system. This is particularly useful for the **VCBimage** backup method, during which the virtual machine disks are copied to the backup proxy local disk. This option enables you to specify a mount point where you have enough disk space. Default: `Data_Protector_home\tmp`

Using the Data Protector CLI

From the system:

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

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

```
util_vmware -configure –instance datacenter -vm vm1_path
-chain { 0 | 1 | 2} -proxy backup_proxy1 [-proxymount
proxy_mount_point1] [-vm vm2_path -chain { 0 | 1 | 2} -proxy
backup_proxy2 [-proxymount proxy_mount_point2]...]
```

The values `{ 0 | 1 | 2}` represent the **Disabled**, **Single**, and **Mixed** snapshot handling mode respectively.

To change virtual machine specific settings back to default, run:

```
util_vmware -config -instance datacenter -vm vm1_path -default
[-vm vm2_path -default ...]
```

TIP:

You can join the options for configuring virtual machines and the options for configuring VirtualCenter system c
Suppose you want to set **Integrated security** for the VirtualCenter system `virtualcenter2.company.com`
mode and usage of the backup proxy system `proxy2.company.com` for the virtual machine vm_path that be
Then, on the VirtualCenter system`virtualcenter2.company.com`, from the directory `Data_Protector`

```
util_vmware.exe -config -security 1 -instance /My datacenters/Datacenter1 -vm
```

Creating backup specifications

Create a backup specification using the Data Protector Manager.

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

4. In **Client**, select a VirtualCenter sytem (VirtualCenter environment) or ESX Server system (standalone ESX Server environment). If the VirtualCenter system is in a cluster, specify the virtual VirtualCenter name.

NOTE:

If you select an ESX Server system that is managed by a VirtualCenter system, you may encounter problems because the ESX Server system is not aware about what is happening with the ESX Server systems that are in the same datacenter.

The **Client** drop-down list shows all clients that have the `VMware integration` component installed.

If the VirtualCenter or ESX Server system is not configured yet, a warning is displayed that the configuration check failed. Click **OK** to open the Configure VMware dialog box and provide the connection parameters as described in "Configuring virtual machines" on page 133.

In **Application database**, select a datacenter that you want to back up. A standalone ESX Server system has only one datacenter (`/ha-datacenter`).

If you have selected a standalone ESX Server system, you also need to provide an ESX Server user (**Username** and **Group name**). This user must be configured as described in "Configuring VMware users" on page 128. This user will be the backup owner.

Click **Next**.

5. For the **Snapshot**, **Suspend**, and **VCBimage** backup methods, select virtual machines or individual virtual machine disks that you want to back up.

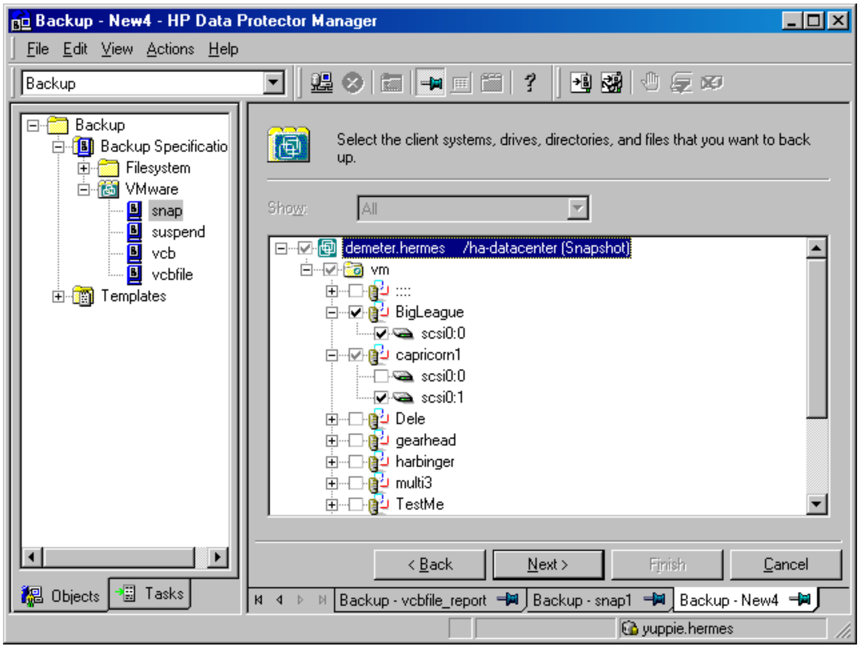

Select the client systems, drives, directories, and files that you want to back up.

- ☑ demeter.hermes /ha-datacenter (Snapshot)
  - ☑ vm
    - □ :::
    - ☑ BigLeague
      - ☑ scsi0:0
    - ☑ capricorn1
      - □ scsi0:0
      - ☑ scsi0:1
    - □ Dele
    - □ gearhead
    - □ harbinger
    - □ multi3
    - □ TestMe

Figure 54 Selecting VMware objects (Snapshot, Suspend, VCB image)

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

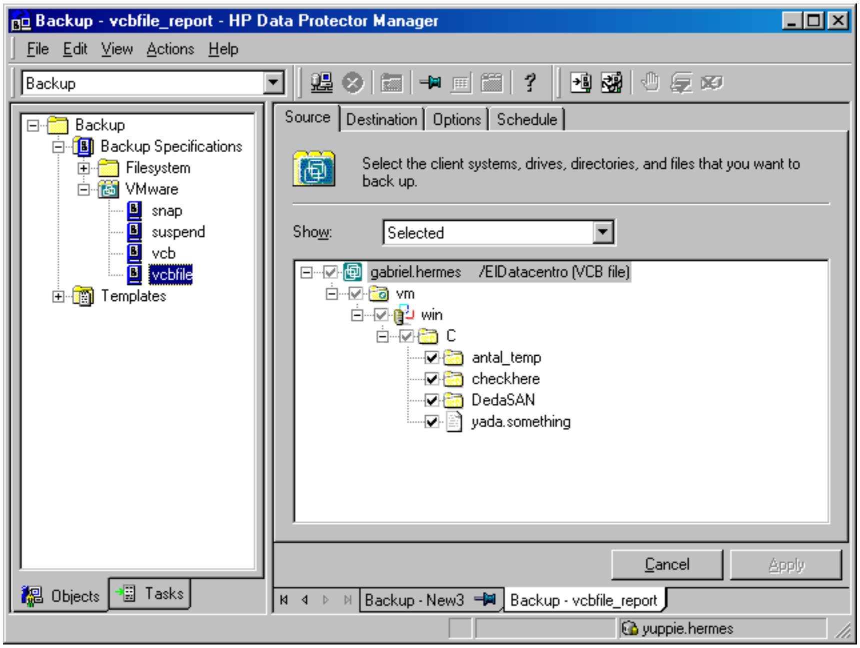

Figure 55 Selecting VMware objects (VCB file)

NOTE:

If you select the client system at the top, all existing virtual machines in the datacenter will be backed up, even if you create new virtual machines in the mean time.

If your virtual machines are not configured yet, right-click the client system at the top or any of the virtual machines listed below, and click **Configure Virtual Machines**. For details, see "Configuring VirtualCenter or ESX Server systems" on page 128.

Click **Next**.

6. Select devices to use for the backup.

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

   Click **Next**.

7. Set backup options. For information on application specific backup options, see Table 24 on page 141.

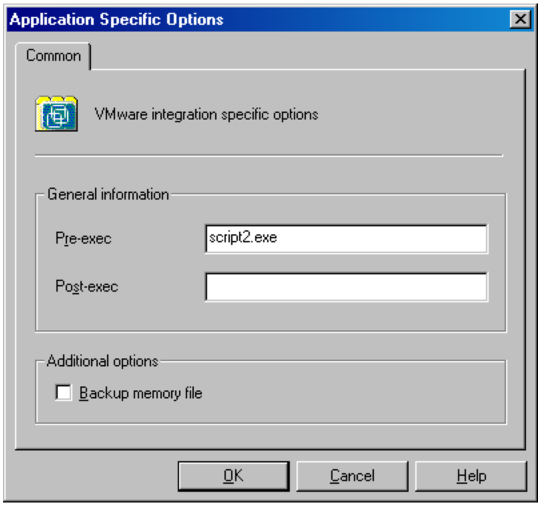

Application Specific Options

Common

VMware integration specific options

General information

Pre-exec: script2.exe

Post-exec:

Additional options

□ Backup memory file

Figure 56 Application specific options

Click **Next**.

8. Optionally, schedule the backup. See "Scheduling backup specifications" on page 142.

   Click **Next**.

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

TIP:

Preview your backup specification before using it for real. See "Previewing backup sessions" on page 143.

Table 24 VMware backup options

Options	Description
Pre-exec, Post-exec	The command specified here is run by `vmware_bar.exe` on the VirtualCenter system or standalone ESX Server system before the backup (`pre-exec`) or after it (`post-exec`). Do not use double quotes. Type only the name of the command and ensure that the command resides in the following directory: ***VirtualCenter system:*** `Data_Protector_home\bin` ***Standalone ESX Server system:*** `/opt/omni/lbin`
Backup memory file	Available only for the **Snapshot** and **Suspend** backup methods. Set this option ON to back up virtual machines together with their memory files. If you restore from such a backup, you do not need to boot the virtual machines after the restore. They are automatically put online. However, the backup lasts considerably longer if this option is ON.

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.

NOTE:

Suppose you created a backup specification in such a way that you selected the client system in the Source page. When you open such a backup specification A backup using such a backup specification will always include all virtual machines, even if new virtual machines were created in the mean time.

Scheduling backup specifications

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

Scheduling example

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

1. In the **Schedule** property page of the backup specification, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See Figure 57 on page 143. Under **Session options**, select **Differential** from the **Backup type** drop-down list.

   Click **OK**.
3. Repeat Step 1 on page 142 and Step 2 on page 142 to schedule differential backups also at 13:00 and 18:00.
4. Click **Apply** to save the changes.

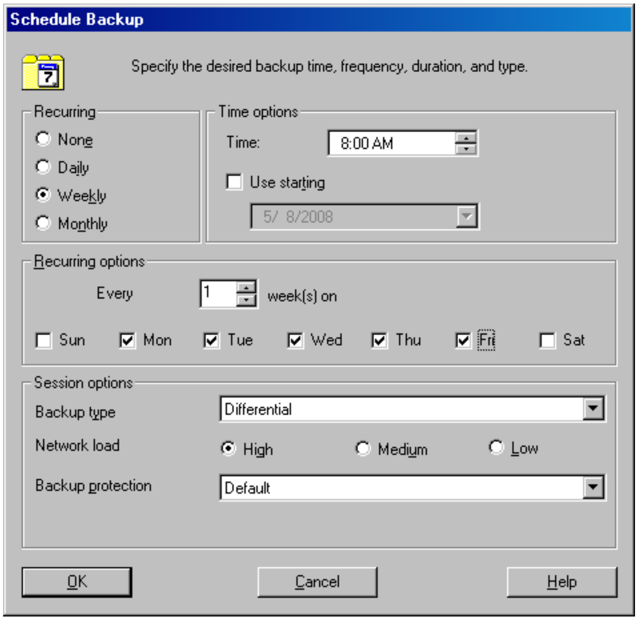

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

○ None
○ Daily
● Weekly
○ Monthly

Time options

Time: 8:00 AM

□ Use starting

5/ 8/2008

Recurring options

Every 1 week(s) on

□ Sun ☑ Mon ☑ Tue ☑ Wed ☑ Thu ☑ Fri □ Sat

Session options

Backup type: Differential

Network load: ● High ○ Medium ○ Low

Backup protection: Default

Figure 57 Scheduling a backup specification

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the system:

VirtualCenter system: `Data_Protector_home\bin`

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

`Data_Protector_home``\bin`, run:

```
omnib -vmware_list backup_specification_name -test_bar
```

What happens during the preview?

The following is tested:

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

Starting backup sessions

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

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the system:

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

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

run:

```
omnib -vmware_list backup_specification_name [-barmode
VMware_mode] [List_options]
```

where `VMware_mode` is one of the following backup types:

```
full|incr|diff
```

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

Examples

To start a full backup using the backup specification `snapshot2`, run:

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

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

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

Preparing for disaster recovery

To be able to do a disaster recovery, you need backups of the following objects:

Table 25 What must be backed up

Objects	How to back up
ESX Server console	**1.** Ensure that the Data Protector `Disk Agent` component is installed on all the ESX Server systems. **2.** In the Backup context of the Data Protector GUI, right-click **Filesystem** and click **Add backup** to create a backup specification of the filesystem type. In the Source page of the backup specification, select the ESX Server consoles of all the ESX Server systems. **3.** Start a backup using the newly created backup specification.
VirtualCenter configuration database (applicable only for VirtualCenter environments)	**1.** Ensure that the Data Protector `Oracle Integration` component is installed on the VirtualCenter system. **2.** In the Backup context of the Data Protector GUI, right-click **Oracle Server** and click **Add backup** to create a backup specification of the Oracle type. In **Application database**, type the name of the VirtualCenter configuration database. Continue with the backup specification creation as described in the *HP Data Protector integration guide for Oracle and SAP*.

Objects	How to back up
	3. Start a backup using the newly created backup specification.
VMware virtual machines	Back up the virtual machines as described in this chapter.

Restore

You can restore virtual machines using the Data Protector GUI or CLI.

Whenever you restore virtual machines that already exist in the destination datacenter, they are first shut down and, if needed, also unregistered. Then, the files are restored and the virtual machines return to a state that you specify in the restore options. For example, they can be registered again and put online.

Before you begin

- If you plan to restore virtual machines to another location, ensure that the destination client is part of the Data Protector cell and has the `VMware Integration` component installed.

Considerations

- **Concurrent sessions:** Restore sessions that use the same devices or restore the same datacenter cannot run concurrently.
- **Orphaned virtual machines:** If you open the Virtual Infrastructure client while the restore session is in progress, sometimes, you will find that the virtual machines that are being restored are shaded and the note `(orphaned)` is added next to them. This happens in the following situations:
  - A virtual machine that is being restored is still registered in the destination datacenter, although the virtual machine files have been deleted.

Finding information for restore

You can find information about backup objects in the Data Protector IDB. For example, you may want to know which backup type and media were used, and which messages were displayed during the backup. To retrieve this information, use the Data Protector GUI or CLI.

Using the Data Protector GUI

In the Internal Database context, expand **Objects** or **Sessions**.

If you expand **Objects**, backup objects are sorted according to virtual machines for which they were created. For example, backup objects for the virtual machine `mach1` from the datacenter `ELDatacentro` are listed under `/%2FEldatacentro/0/%2Fvm%2Fmach1`.

If you expand **Sessions**, backup objects are sorted according to sessions in which they were created. For example, backup objects created in the session `2008/08/15-7` are listed under `2008/08/15-7`.

To view details on a backup object, right-click the backup object and click **Properties**.

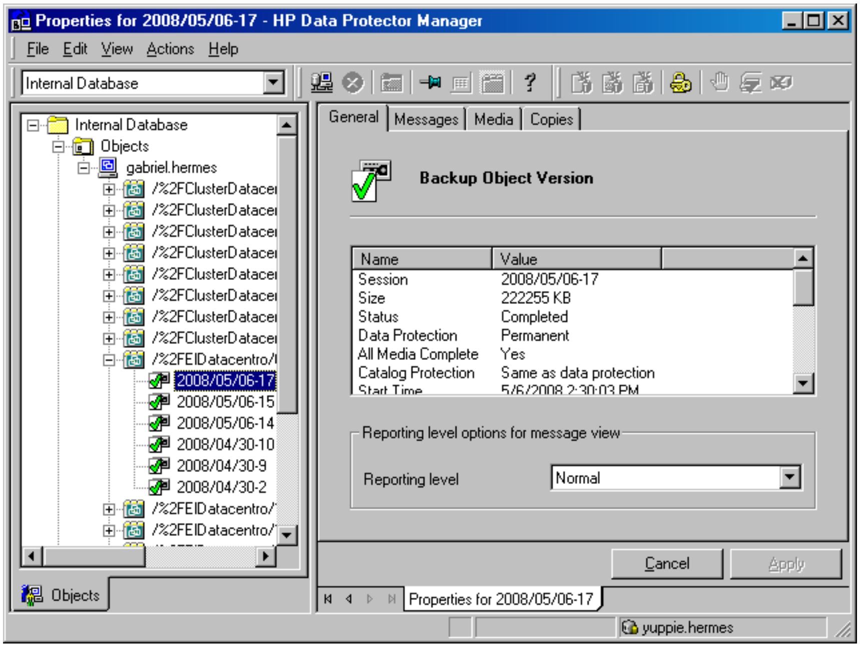

Name	Value
Session	2008/05/06-17
Size	222255 KB
Status	Completed
Data Protection	Permanent
All Media Complete	Yes
Catalog Protection	Same as data protection
Start Time	5/6/2008 2:30:03 PM

Reporting level options for message view

Reporting level: Normal

Figure 58 Backup object information

TIP:

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

Using the Data Protector CLI

1. Go to the following directory:

   VirtualCenter system: `Data_Protector_home\bin`

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

2. Get a list of VMware backup objects created in a particular backup session:

   ```
   omnidb -session session_id
   ```

3. Get details on a particular backup object:

   ```
   omnidb -vmware backup_object_name -session session_id
   -catalog
   ```

   Here is one example of a backup object name:

   ```
   gabriel.hermes.com::/%2FElDatacentro/0/%2Fvm%2Fharbour
   ```

For details, see the omnidb man page.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **VMware**, expand a VirtualCenter system (VirtualCenter environment) or a standalone ESX Server system (standalone ESX Server environment) and click the datacenter you would like to restore.

3. In the **Source** page, virtual machines backed up with the backup method specified in**Backup method** are displayed. To display virtual machines backed up with a different method, change **Backup method**. By using the **From** and **To** options, you narrow the scope of displayed virtual machines to those backed up within the specified time interval.

   Select VMware objects you would like to restore. You cannot restore virtual machines backed up with different backup methods in the same session. Each selected VMware object will be restored from the last backup created during the specified time interval.

IMPORTANT:

If you restore from a **Snapshot**, **Suspend**, or **VCB image** backup, check if the complete virtual machines were backed up. If only individual virtual machine disks were backed up, you cannot automatically register and power-on such virtual machine when the restore completes. In this case, you have to restore the virtual machine disks with the
Register virtual machines if needed and
Power-on virtual machines after restore options set OFF.

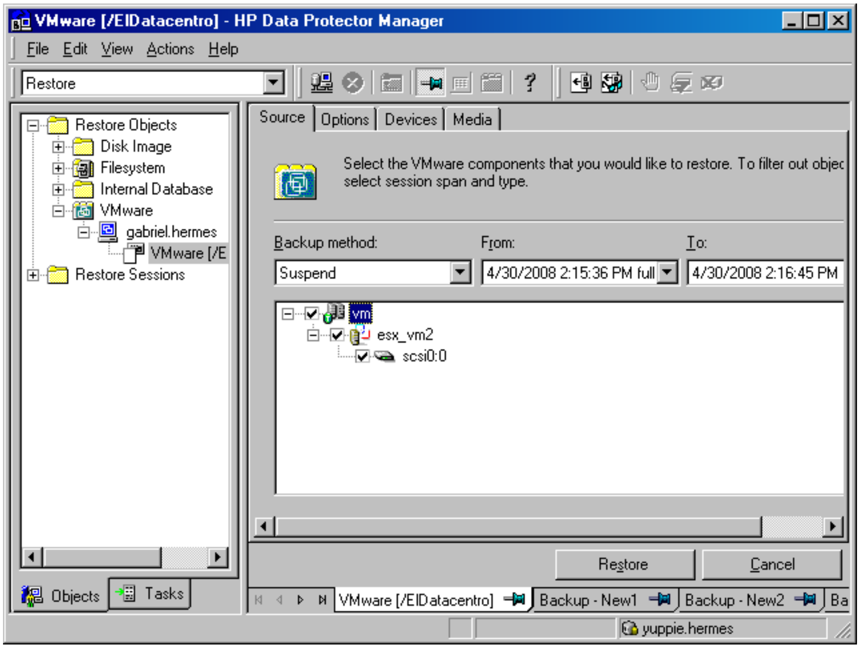

Select the VMware components that you would like to restore. To filter out objec select session span and type.

Backup method: Suspend

From: 4/30/2008 2:15:36 PM full

To: 4/30/2008 2:16:45 PM

Figure 59 Selecting VMware objects for restore (Snapshot, Suspend, VCB image)

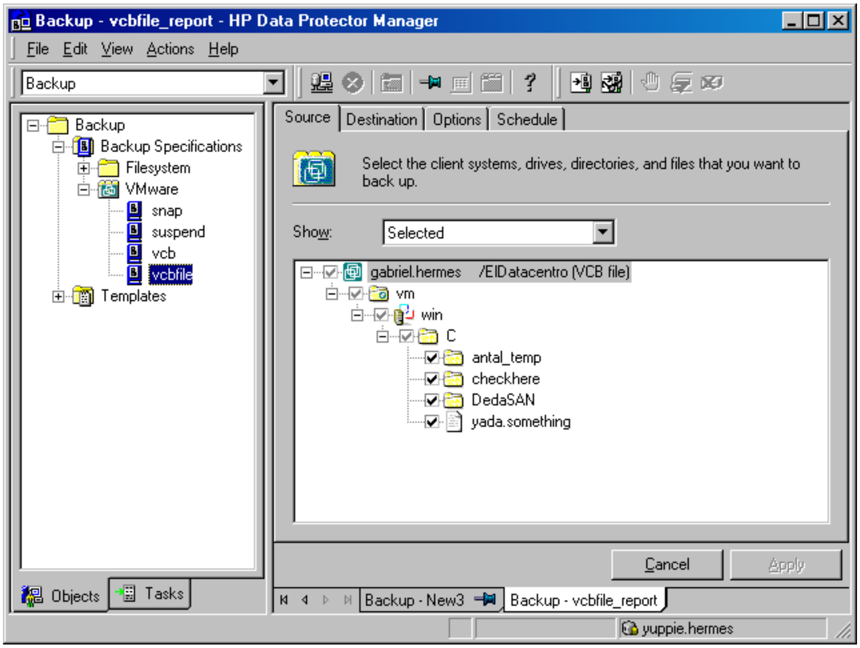

Backup - vcbfile_report - HP Data Protector Manager

File Edit View Actions Help

Backup

- Backup
  - Backup Specifications
    - Filesystem
    - VMware
      - snap
      - suspend
      - vcb
      - vcbfile
  - Templates

Objects | Tasks

Source | Destination | Options | Schedule

Select the client systems, drives, directories, and files that you want to back up.

Show: Selected

- ☑ gabriel.hermes /EIDatacentro (VCB file)
  - ☑ vm
    - ☑ win
      - ☑ C
        - ☑ antal_temp
        - ☑ checkhere
        - ☑ DedaSAN
        - ☑ yada.something

Cancel | Apply

Backup - New3 | Backup - vcbfile_report

yuppie.hermes

Figure 60 Selecting VMware objects for restore (VCB file)

For each selected VMware object, Data Protector automatically restores the complete restore chain. For example, if the backup to restore from is:

Full backup	Data Protector restores only this full backup.
Differential backup	Data Protector first restores the latest full backup and then the differential backup.
Incremental backup	Data Protector first restores the latest full backup, then the latest differential backup (if one exists) and all the incremental backups from the latest differential or full backup up to the selected version.

4. In the **Options** page, specify the VMware restore options. For details, see Table 26 on page 153.

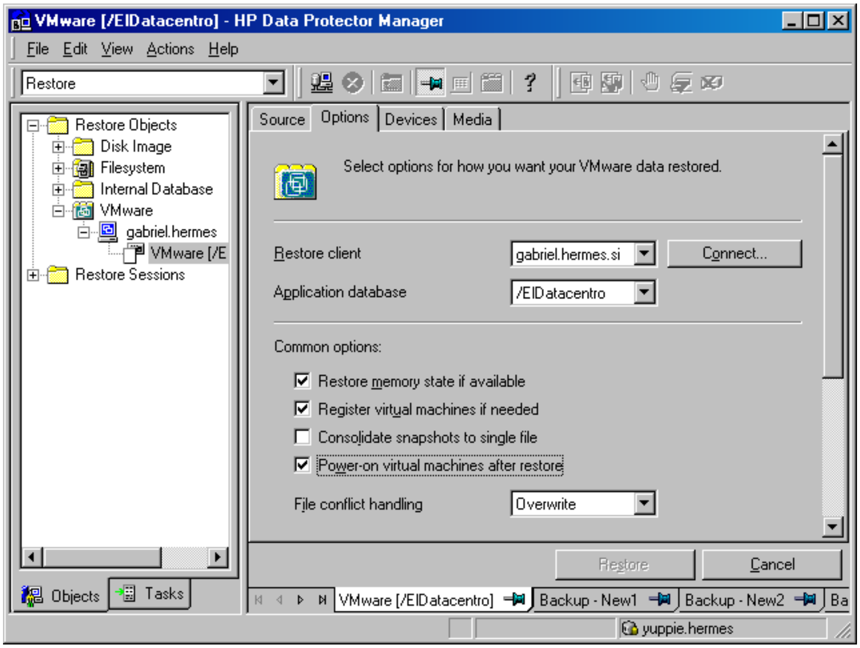

Select options for how you want your VMware data restored.

Restore client: gabriel.hermes.si

Application database: /EIDatacentro

Common options:

☑ Restore memory state if available

☑ Register virtual machines if needed

☐ Consolidate snapshots to single file

☑ Power-on virtual machines after restore

File conflict handling: Overwrite

Figure 61 Restore options (Snapshot, Suspend, VCB image)

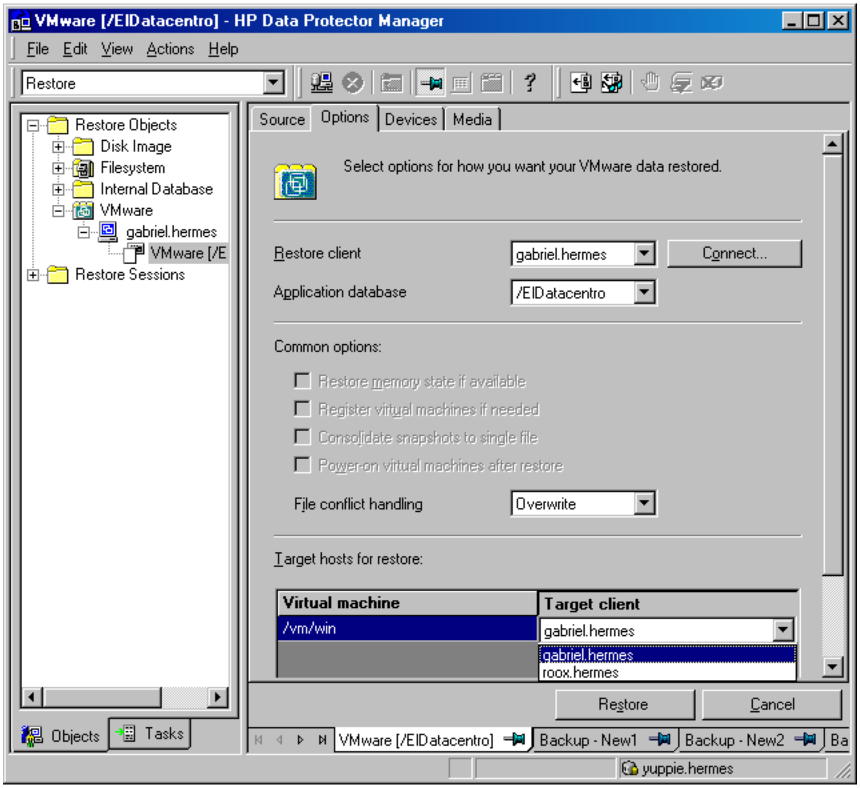

VMware [/EIDatacentro] - HP Data Protector Manager

Select options for how you want your VMware data restored.

Restore client: gabriel.hermes

Application database: /EIDatacentro

Common options:

- □ Restore memory state if available
- □ Register virtual machines if needed
- □ Consolidate snapshots to single file
- □ Power-on virtual machines after restore

File conflict handling: Overwrite

Target hosts for restore:

Virtual machine	Target client
/vm/win	gabriel.hermes

Figure 62 Restore options (VCB file)

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

   Click **Finish** to start the restore.

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

If you restored virtual machines with the **Register virtual machines if needed** and **Power-on virtual machines after restore** options set OFF, you need to do some additional steps as described in the "Recovering virtual machines" on page 156.

Table 26 VMware restore options

Option	Description
Restore client	Specifies a VirtualCenter system or a standalone ESX Server system on which the restore session is started. If you specify a client that is not configured yet, a warning is displayed. Click **OK** to open the Configure VMware dialog box and provide the connection parameters as described in "Configuring VirtualCenter or ESX Server systems" on page 128. If you restore from a **Snapshot**, **Suspend**, or **VCBimage** backup, this client should be the one that manages the datacenter to which you want to restore the virtual machines. Default: The client on which the backup was started.
Application database	Specifies in which datacenter to restore the virtual machines. The destination datacenter should have the same name and configuration as the datacenter from which the virtual machines were backed up. This option is not applicable for the **VCBfile** method. For this method, you specify the restore destination in **Target hosts for restore** option at the bottom. Default: The datacenter from which the virtual machines were backed up.
Restore memory state if available	Set this option ON to restore also the virtual machine memory file if it was backed up. This option is not available if you restore from a **VCBfile** backup. Default: OFF
Register virtual machines if needed	Set this option ON to be able to restore virtual machines to datacenters in which virtual machines with such names do not exist. If this option is OFF, the restore of such virtual machines is skipped. This option is not available if you restore from a **VCBfile** backup. Default: OFF

<table>
<tr><th colspan="2">Option</th><th>Description</th></tr>
<tr><td colspan="2"><b>Consolidate snapshots to single file</b></td><td>Set this option ON to delete all existing snapshots when the restore completes. It means that all the changes on the active branch are committed to the base. Consequently, changes made on non-active branches are lost.
This option is not available if you restore from a <b>VCBfile</b> backup.
Default: OFF</td></tr>
<tr><td colspan="2"><b>Power-on virtual machines after restore</b></td><td>Select this option to put the newly restored virtual machines online when the session completes. If virtual machines are restored from a backup that includes memory file and the memory file is also restored (<b>Restore memory state if available</b> is ON), the virtual machines are automatically put online after the restore.
This option is not available if you restore from a <b>VCBfile</b> backup.
Default: OFF</td></tr>
<tr><td rowspan="4"><b>File conflict handling</b></td><td colspan="2">Specifies what to do if virtual machines to be restored already exist in the destination datacenter. You can choose among the following:</td></tr>
<tr><td><b>Overwrite</b> (default)</td><td>The existing virtual machine files are overwritten with those from the backup.</td></tr>
<tr><td><b>Keep latest</b></td><td>The restore of virtual machine files is skipped if the files in the destination datacenter are more recent than those from the backup. Otherwise, the files are overwritten with those from the backup.</td></tr>
<tr><td><b>Skip</b></td><td>The restore of existing virtual machine files is skipped.</td></tr>
<tr><td colspan="2"><b>Target hosts for restore</b></td><td>Specifies the restore destination for the virtual machine filesystems. The target system must be a Windows virtual machine or a separate Windows system, having the VMFS filesystem.
This option is available only if you restore from a <b>VCBfile</b> backup.</td></tr>
</table>

Restoring using the Data Protector CLI

From the system:

VirtualCenter system: *Data_Protector_home*\bin

Standalone ESX Server system: /opt/omni/lbin

run:

```
omnir -vmware -barhost Original_VirtualCenter_or
ESX_Server_client -instance Original_datacenter [-destination
Target_client] [-newinstance Target_datacenter] [-session
Session_ID] -method Backup_method [-register] [-poweron]
[-overwrite [older]] -vm Virtual_machine1 [-vm
Virtual_machine2 ...]
```

Example

Suppose you want to restore the virtual machines MachineA and MachineB. At the time of backup, the virtual machines were running on the ESX Server systems that belonged to the datacenter MyDatacenter managed by the VirtualCenter system Virtualcenter1.company.com. The virtual machines were backed up with the Snapshot backup method. Suppose you want to restore the virtual machines to the datacenter NewDatacenter that is managed by a different VirtualCenter system (Virtualcenter2.company.com. In the datacenter NewDatacenter, the virtual machines MachineA and MachineB are not registered. You want to restore from the backup session 2008/07/14-1. If, in this session, virtual machine memory files were also backed up, you want to restore them as well. You also want to ensure that the newly restored virtual machines are put online when the session completes. If virtual machine files to be restored already exist in the destination datacenter, preserve them in case they are more recent than those from the backup. To achieve all this, run:

Example

```
omnir -vmware -barhost Virtualcenter1.company.com -datacenter
MyDatacenter -destination Virtualcenter2.company.com
-newinstance NewDatacenter -session 2006/2/7-31 -vm MachineA
-vm MachineB -overwrite older -memory -register -poweron
```

Recovering virtual machines

Whenever you restore virtual machines with the **Register virtual machines if needed** and **Power-on virtual machines after restore** options set OFF, you need to do additional steps to fully recover the virtual machines:

1. For each virtual machine, list the current virtual machine disk files. See examples below.
2. In case a chain of sessions (full, differential, incremental) was restored, the files from different sessions must be correctly linked: each virtual disk metadata file must be updated to point to the virtual disk metadata file from the previous backup session. The same must be done for the virtual disk extent files. See examples below.
3. Export the latest virtual disk metadata file.
4. In the Virtual Infrastructure Client, select the virtual machine:

   Edit settings > **Add disk** > **Use existing virtual disk and latest disk metadata file**

   Select the latest virtual disk metadatafile and click **OK**.

Restoring using another device

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

Using the Data Protector GUI

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

Using the Data Protector CLI

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

Data_Protector_home`\Config\Server\cell\restoredev`

Disaster recovery

Disaster recovery is a very complex process, involving products from different vendors. Therefore, you need to check the instructions from operating system and VMware vendor on how to prepare for disaster recovery.

The following are the main steps needed to recover a virtual machine after a disaster:

1. Reinstall the VMware environment. The configuration should be the same as during the backup (VirtualCenter systems, ESX Server systems and datacenters should have the same names). Install Data Protector in the newly configured environment.
2. Restore the console of the ESX Server system on which the virtual machine was running to the newly configured ESX Server system from a Data Protector filesystem backup. For details, see the online Help.
3. Restore the original VirtualCenter database from a Data Protector Oracle backup (if needed). For details, see the *HP Data Protector integration guide for Oracle and SAP*.
4. Restore the virtual machine from a Data Protector VMware backup as described in this chapter.

Monitoring sessions

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

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

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

Troubleshooting

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

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

Before you begin

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

Checks and verifications

If your configuration, backup, or restore failed:

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

  VirtualCenter system: `*Data_Protector_home*\log`.

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

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

Additionally, if your backup failed:

- Check the configuration of the VirtualCenter system or standalone ESX Server system as described in “Configuring VirtualCenter or ESX Server systems” on page 128.

Problems

Problem

Backup fails with “insufficient resources to satisfy failover level”

If your ESX Server systems are in a high availability cluster and you start a backup of virtual machines that have migrated to another ESX Server systems, the backup fails with an error similar to the following:

```
[Critical] From: OB2BAR_VMWARE_BAR@gabriel.hermes.si
"/ClusterDatacenter" Time: 7.4.2008 16:13:50 Virtual machine
'/vm/vmsan1': operation failed: Error: {
localizedMessage='Insufficient resources to satisfy configured
failover level for HA.';
```

Action

1. Open the cluster settings dialog box in the Virtual Infrastructure client and set the **Allow virtual machines to be powered on even if they violate availability constrains** option ON.
2. Restart the backup.

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 (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.

Audit Logs

Data files to which auditing information is stored.

Audit Report

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

Auditing Information

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

autochanger

See library.

autoloader

See library.

Automatic Storage Management

(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.

automigration

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

See also Virtual Library System (VLS) and virtual tape.

BACKINT

(SAP R/3 specific term) SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API

The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain

See restore chain.

backup device

A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation

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

backup ID

An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object

A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk).
A backup object is defined by:

- Client name: Hostname of the Data Protector client where the backup object resides.
- Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is located (drive on Windows and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.
- Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration

objects — displays the integration type (for example, SAP or Lotus).

- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — “Bar”.

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

backup session A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also backup specification, incremental backup, and full backup.

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

backup set *(Oracle specific term)* A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)* A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.
See also application system, target volume, and replica.

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

backup to IAP

A Data Protector based backup to the HP Integrated Archiving Platform (IAP) appliance. It takes advantage of the IAP capability to eliminate redundancies in the stored data at a block (or chunk) level, by creating a unique content address for each data chunk. Only changed chunks are transmitted over the network and added to the store.

backup view

Data Protector provides different views for backup specifications:
By Type - according to the type of data available for backups/templates. Default view.
By Group - according to the group to which backup specifications/templates belong.
By Name - according to the name of backup specifications/templates.
By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC

(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 should be connected to the backup system.
See also HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.

BC EVA

(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.

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 manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE

(SAP R/3 specific term) An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also BRBACKUP, and BRRESTORE.

BRBACKUP

(SAP R/3 specific term) An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data

files, or of all tablespaces and, if necessary, of the online redo log files.
See also BRARCHIVE, and BRRESTORE.

BRRESTORE

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

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

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also BRBACKUP, and BRARCHIVE.

BSM

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

CA

(HP StorageWorks Disk Array XP specific term) Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
See also BC *(HP StorageWorks Disk Array XP specific term)*, Main Control Unit and HP StorageWorks Disk Array XP LDEV.

CA+BC EVA

(HP StorageWorks EVA specific term) The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.
See also BC EVA, replica, and source volume.

CAP

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

catalog protection

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

CDB

The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions,, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also MMDB.

CDF file

(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.

Change Journal

(Windows specific term) A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider

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

channel

(Oracle specific term) An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type 'disk'
- type ‘sbt_tape’

If the specified channel is of type ‘sbt_tape’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

chunking

(IAP specific term) The process of dividing data into blocks (chunks), where each chunk gets a unique content address. This address is then used to determine whether a particular chunk is already backed up to the IAP appliance. If the duplicate data is identified (two addresses are identical, that is the address is the same as for another data chunk already stored into IAP), it is not backed up. This way, the data redundancy is eliminated and the optimal data storage is achieved.
See also backup to IAP.

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 filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client

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

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

Cluster Continuous Replication

(Microsoft Exchange Server specific term) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.
A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.
See also Exchange Replication Service and Local Continuous Replication.

CMD Script for Informix Server

(Informix Server specific term) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Registration Database

(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 (CLI)

A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA

(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 and HP StorageWorks SMI-S EVA provider.

Command View VLS

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

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), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

CSM
The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

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 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.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

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

DC directory
The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located 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 use a custom location. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

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

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

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

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group
(EMC Symmetrix specific term) A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

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

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

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

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

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 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.

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller

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

DR image

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

DR OS

A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See client backup with disk discovery.

EMC Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*

See Symmetrix Agent (SYMA).

emergency boot file

(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 Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the `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.

Event Logs

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service

(Microsoft Exchange Server specific term) The Microsoft Exchange Server service that represents storage groups that were replicated using either Local Continuous Replication (LCR) or Cluster Continuous Replication (CCR) technology.
See also Cluster Continuous Replication and Local Continuous Replication.

exchanger

Also referred to as SCSI Exchanger.
See also library.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also importing media.

Extensible Storage Engine (ESE)

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

failover

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

failover

(HP 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 bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge

A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot

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

file jukebox device

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

file library device

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

File Replication Service (FRS)

A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file tree walk

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

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror

(HP 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 number.

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 (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup

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

full database backup

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup

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

full ZDB

A ZDB to tape or ZDB to disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.
See also incremental ZDB.

global options file

A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the `/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, Linux, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI, Data Protector also provides a Java-based graphical user interface with the same look and feel. As Java can run on numerous platforms, the Data Protector Java GUI is supported on a larger number of platforms than the original Data Protector GUI.

hard recovery

(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.

HP ITO

See OM.

HP OpC

See OM.

HP OpenView SMART Plug-In (SPI)

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

HP OM

See **OM**

HP StorageWorks Disk Array XP LDEV

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also BC, CA *(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.

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also BC VA and replica.

HP VPO

See OM.

ICDA

(EMC Symmetrix specific term) EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

IDB recovery file

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

importing media A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
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.

incremental1 mailbox backup An incremental1 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 two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also Key Management Service and Site Replication Service.

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

initializing *See* formatting.

Installation Server A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery *(ZDB specific term)* A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

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

Internet Information Services (IIS) *(Windows specific term)* Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

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

ITO *See* OM.

Java GUI Client The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities and requires connection to the Java GUI Server to function.

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

jukebox *See* library.

jukebox device A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

Key Management Service *(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.

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

load balancing By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

Local Continuous Replication

(Microsoft Exchange Server specific term) Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.

An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal.

A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.

See also Cluster Continuous Replication and Exchange Replication Service.

lock name

You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script

(Informix Server UNIX specific term) A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server `ALARMPROGRAM` configuration parameter defaults to the *`INFORMIXDIR`*`/etc/log_full.sh`, where *`INFORMIXDIR`* is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the `ALARMPROGRAM` configuration parameter to *`INFORMIXDIR`*`/etc/no_log.sh`.

logging level

The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files
This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID
(Microsoft SQL Server specific term) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

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

- `user_name` is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- `password` must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.
- `service` is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database
(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is `user_name/password@service`, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, `service` is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API
(Lotus Domino Server specific term) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX

systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet

See Wake ONLAN.

mailbox

(Microsoft Exchange Server specific term) The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store

(Microsoft Exchange Server specific term) A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

Main Control Unit (MCU)

(HP StorageWorks Disk Array XP specific term) An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. *See also* BC (HP StorageWorks Disk Array XP specific term), CA *(HP StorageWorks Disk Array XP specific term)*, and HP StorageWorks Disk Array XP LDEV.

Manager-of-Managers (MoM)

See MoM.

MAPI

(Microsoft Exchange Server specific term) The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU

See Main Control Unit (MCU).

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy

Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

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

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

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

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

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

media pool A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

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

media usage policy The media usage policy controls how new backups are added to the already used media. It can be `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 modification date is kept. Files not present on the disk are always restored.
See also overwrite.

Microsoft Exchange Server

A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC)

(Windows specific term) An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server

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

Microsoft Volume Shadow Copy Service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets. *See also* shadow copy, shadow copy provider, replica, and writer.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*

See target volume.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*

See replica set rotation.

MMD

The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices,

libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also CMMDB, CDB.

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount request

A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point

The access point in a directory structure for a disk or logical volume, for example`/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) Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror.
See also first-level mirror.

multi-drive server

A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat

See IDB recovery file.

OBDR capable device

A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object

See backup object.

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

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

object copy session A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

object ID *(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 (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.
- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

See also zero downtime backup (ZDB) and online backup.

offline recovery

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log

See archived redo log.

On-Bar

(Informix Server specific term) A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

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

ONCONFIG

(Informix Server specific term) An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file in the directory `*INFORMIXDIR*\etc` (on Windows) or `*INFORMIXDIR*/etc/` (on UNIX).

online backup

A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period 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 (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished.
- For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
See also zero downtime backup (ZDB), and offline backup.

online 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* OM.

OpenSSH A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard *(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.

OM

HP Operations Manager software for UNIX provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OM management servers on HP-UX, Solaris, and Linux. Earlier versions of OM were called IT/Operation, Operations Center and Vantage Point Operations.
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\Se ver\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 re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

parallel restore Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism The concept of reading multiple data streams from an online database.

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

post-exec A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also pre-exec.

pre- and post-exec commands Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

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

pre-exec

A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also post-exec.

primary volume (P-VOL)

(HP StorageWorks Disk Array XP specific term) Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and also catalog protection.

public folder store

(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 *See* Remote Control Unit (RCU).

RDBMS Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)* A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

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

Recovery Catalog *(Oracle specific term)* A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

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

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

recovery files *(Oracle specific term)* Recovery files are Oracle 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.

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 (RCU) *(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 or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects

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

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

replica set rotation
(ZDB specific term) The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also replica and replica set.

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

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

resync mode
(HP StorageWorks Disk Array XP VSS provider specific term) One of two XP VSS hardware provider operation modes. When the XP provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.
See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

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

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM

(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.

scan

A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler

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

secondary volume (S-VOL)

(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) and Main Control Unit (MCU)

session

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

session ID

An identifier of a backup, restore, object copy, object consolidation, or media management session, consisting of the date when the session ran and a unique number.

session key

This environment variable for the pre-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the `omnimnt`, `omnistat`, and `omniabort` commands.

shadow copy

(Microsoft VSS specific term) A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.
See also Microsoft Volume Shadow Copy Service and replica.

shadow copy provider

(Microsoft VSS specific term) An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also shadow copy.

shadow copy set

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

shared disks

A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF

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

single instancing

(IAP specific term) The process of recognizing redundancy of data, at both a whole object and a chunk level. It computes a strong hash for each data chunk and uses it as a unique content address needed to determine whether attempts to store duplicates are being made.
See also backup to IAP.

Site Replication Service

(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.

smart copy

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

smart copy pool

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

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, and media management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot

(HP StorageWorks VA and HP StorageWorks EVA specific term) A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also replica and snapshot creation.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)*

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation

(HP StorageWorks VA and HP StorageWorks EVA specific term) A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for

use. However background copying processes normally continue after creation.
See also snapshot.

source (R1) device *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also target (R2) device.

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

sparse file A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents of the source volumes.
See also replica and split mirror creation.

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

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

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also split mirror.

split mirror restore

(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.

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.

Storage Group *(Microsoft Exchange Server specific term)* A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

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 conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

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

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

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

System State *(Windows specific term)* The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

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

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 *(Microsoft SQL Server specific term)* An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)* A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

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

transaction log backup Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

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

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

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

transportable snapshot *(Microsoft VSS specific term)* A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed.
See also Microsoft Volume Shadow Copy Service (VSS).

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

UIProxy

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

unattended operation

See lights-out operation.

user account (Data Protector user account)

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user 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 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

(Microsoft 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 Library System (VLS)

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

virtual server

A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape

(VLS specific term) An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs.
See also Virtual Library System (VLS) and Virtual Tape Library.

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

volser
(ADIC and STK specific term) A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mount point
(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 OM.

VSS
See Microsoft Volume Shadow Copy Service.

VSS compliant mode
(HP StorageWorks Disk Array XP VSS provider specific term) One of two XP VSS hardware provider operation modes. When the XP provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching of the disks.
See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

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

Web reporting The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows CONFIGURATION backup Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

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

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer *(Microsoft VSS specific term)* A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

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, instant recovery, and split mirror restore.
See also zero downtime backup (ZDB).

ZDB to disk
(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape
(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape
(ZDB specific term) A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.
See also zero downtime backup (ZDB), ZDB to disk, instant recovery, ZDB to disk+tape, and replica.

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index