

HP Data Protector

Software Version: 8.10

Integration guide for Sybase and Network Data Management Protocol Server

Document Release Date: November 2016

Software Release Date: November 2016



Legal Notices

Warranty

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.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2016 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: <http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: <http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at: <http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is <http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Contents	3
Chapter 1: Data Protector Sybase Server integration	7
Introduction	7
Integration concepts	7
Data Protector CLI commands	8
Configuring the integration	9
Prerequisites	9
Before you begin	9
Cluster-aware clients	10
Configuring Sybase users	10
Configuring Sybase instances	10
Before you begin	10
Using the Data Protector GUI	10
Using the Data Protector CLI	13
Checking the configuration	13
Using the Data Protector GUI	14
Using the Data Protector CLI	14
Backup	14
Creating backup specifications	15
Modifying backup specifications	20
Scheduling backup sessions	20
Previewing backup sessions	21
Using the Data Protector GUI	21
Using the Data Protector CLI	21
What happens during the preview?	22
Starting backup sessions	22
Using the Data Protector GUI	23
Using the Data Protector CLI	23
Using Sybase commands	23
Restore	24

Localized database names	24
Finding information for restore	24
Using the Data Protector GUI	25
Using the Data Protector CLI	25
Using the Data Protector syb_tool command	25
Using the standard Data Protector CLI commands	28
Restoring using the Sybase isql command	30
Restore examples	32
Restoring using another device	33
Monitoring sessions	34
Troubleshooting	34
Before you begin	34
Checks and verifications	34
Problems	36
Chapter 2: Data Protector Network Data Management Protocol Server integration	37
Introduction	37
Integration concept	38
Configuring the integration	40
Prerequisites	40
Importing NDMP Server systems	40
Creating media pools	42
Configuring NDMP devices	42
Configuring tape libraries	44
Configuring standalone devices	46
Network Appliance configuration	48
Standalone tape devices and drives in a tape library	48
Library robotics	48
EMC Celerra configuration	49
SCSI devices	49
Hitachi BlueArc or Hitachi configuration	50
Library robotics	50

Standalone tape devices and drives in a tape library	51
Block size	51
Backup	52
Before you begin	53
Creating backup specifications	53
Backing up large data set	57
Modifying backup specifications	57
Starting backup sessions	57
Restore	57
Restoring using the Data Protector GUI	58
Direct access restore	60
Recommendations for single-file restore from a large backup set	61
Restoring using another device	61
NDMP environment variables	61
The NDMP specific omnirc options	64
Media management	66
Troubleshooting	67
Before you begin	67
Problems	67
Appendix A: Data Protector NetApp SnapManager solution	69
Introduction	69
Concepts	69
Configuration	69
Prerequisites	69
Configuration procedure	69
Backup	70
Limitations	70
Creating a backup specification	71
Restore	73
omnisnapmgr.pl reference page	76
SYNOPSIS	76
DESCRIPTION	76

OPTIONS	76
NOTES	77
EXAMPLES	77
Glossary	79
Index	120
We appreciate your feedback!	124

Chapter 1: Data Protector Sybase Server integration

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 1: 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 the *HP Data Protector 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. [Sybase integration architecture on the next page](#) shows the architecture of the Data Protector Sybase integration.

Figure 1: Sybase integration architecture

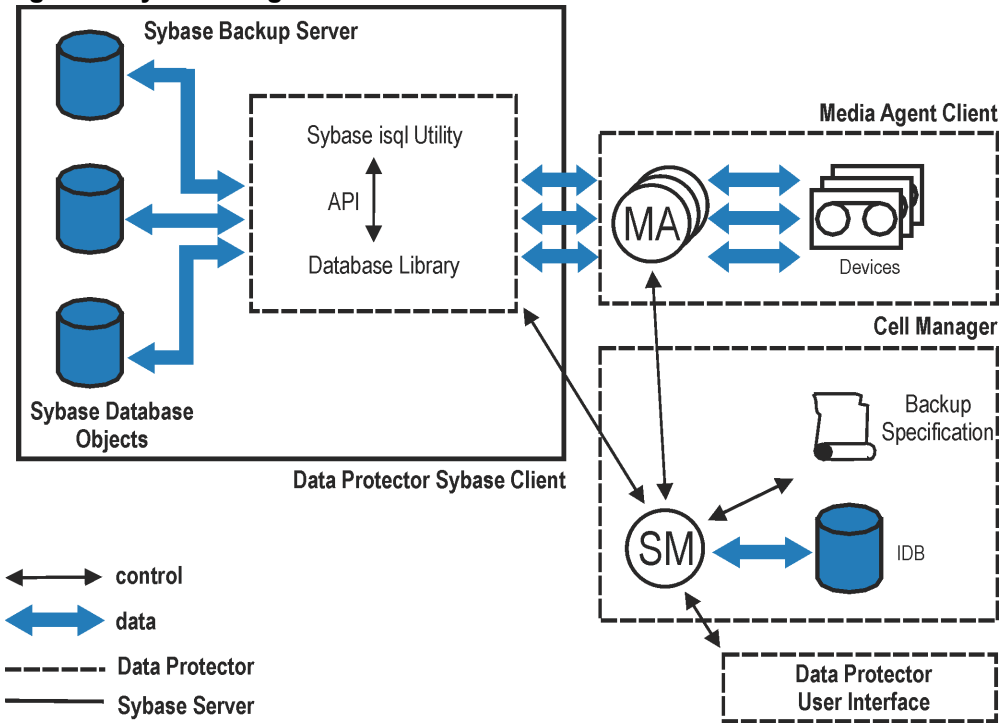


Table 2: Legend

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

Execute the Data Protector CLI commands from the following directories:

Windows systems: `Data_Protector_home\bin`

UNIX systems:

Command	Directory
testbar	opt/omni/bin
omnigetmsg	opt/omni/lbin
util_sybase.pl	

For other command locations, see the `omniintro` reference page in the *HP Data Protector Command Line Interface Reference* or the `omniintro` man page.

To execute the commands, you must have appropriate Data Protector user rights. For information, see the *HP Data Protector 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://support.openview.hp.com/selfsolve/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 systems, 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 *HP Data Protector 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 either the Data Protector GUI or 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

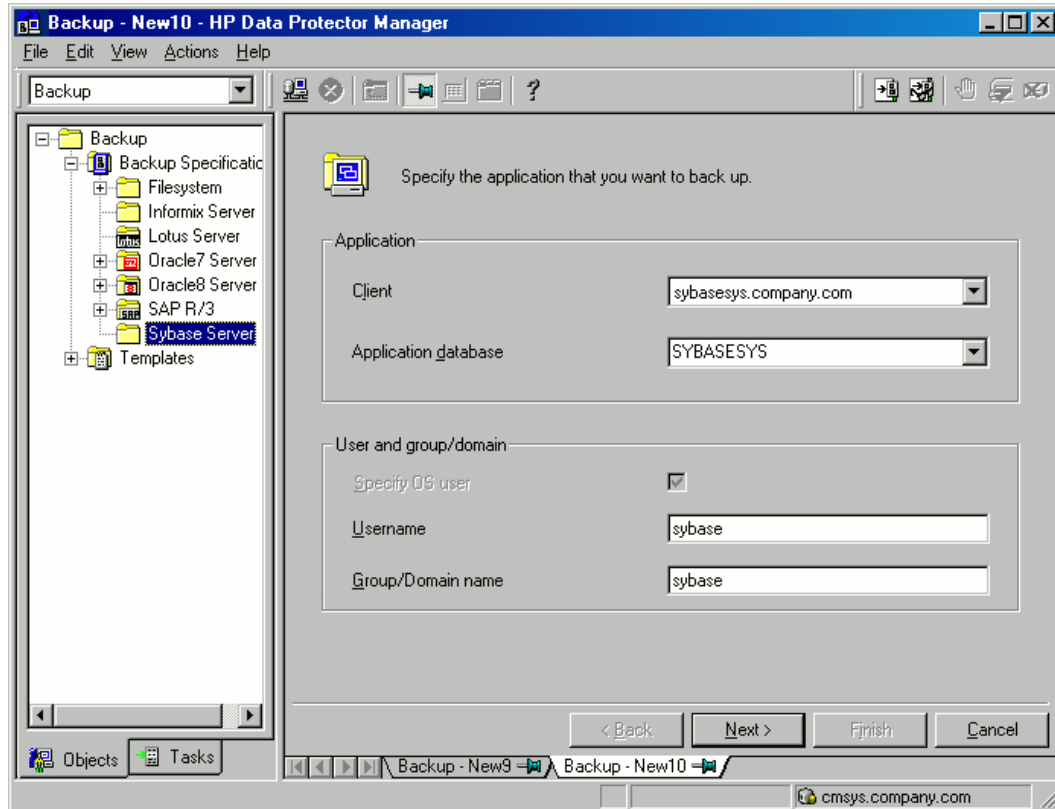
1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Sybase Server**, and click **Add Backup**.

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 systems: Type sybase in both **Username** and **Group/Domain name**. This user becomes the backup owner.

Figure 2: 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 systems, you need to set the Sybase Server home directory, and username and password of the Sybase instance user that has the Sybase rights to back up and restore databases. See [Configuring a Sybase instance \(Windows systems\) on the next page](#) and [Configuring a Sybase instance \(UNIX systems\) on the next page](#).

Figure 3: Configuring a Sybase instance (Windows systems)

The screenshot shows the 'Configure Sybase' dialog box for Windows systems. The title bar reads 'Configure Sybase'. Below the title bar, there is a Windows logo and the text 'Sybase configuration'. The dialog contains several input fields: 'Client' with the value 'sybasesys.company.com', 'Sybase Server name' with the value 'SYBASESYS', 'Sybase Server home directory' with the value 'C:\Sybase', 'Full pathname for Sybase jsq command' with the value 'C:\Sybase\OCS-12_0\bin\isql.exe', 'Sybase User' with the value 'sa', 'Password of the Sybase User' with a masked password 'xxxxxxx', 'SYBASE_ASE (only for Sybase 12.x)' with the value 'ASE-12_0', and 'SYBASE_OCS (only for Sybase 12.x)' with the value 'OCS-12_0'. At the bottom, there are three buttons: 'OK', 'Cancel', and 'Help'.

Figure 4: Configuring a Sybase instance (UNIX systems)

The screenshot shows the 'Configure Sybase' dialog box for UNIX systems. The title bar reads 'Configure Sybase'. Below the title bar, there is a Windows logo and the text 'Sybase configuration'. The dialog contains several input fields: 'Client' with the value 'sybasesys.company.com', 'Sybase Server name' with the value 'SYBASESYS', 'Sybase Server home directory' with the value '/applications/sybase', 'Full pathname for Sybase jsq command' with the value '/applications/sybase/OCS-12_0/bin/isql', 'Sybase User' with the value 'sa', 'Password of the Sybase User' with a masked password 'xxxxxxx', 'SYBASE_ASE (only for Sybase 12.x)' with the value 'ASE-12_0', and 'SYBASE_OCS (only for Sybase 12.x)' with the value 'OCS-12_0'. At the bottom, there are three buttons: 'OK', 'Cancel', and 'Help'.

Click **OK**.

6. The Sybase instance is configured. Exit the GUI or proceed with creating the backup specification at [Select the databases you want to back up. on page 15.](#)

Using the Data Protector CLI

Execute:

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

UNIX systems: util_sybase.pl -CONFIG Sybase_instance Sybase_home \
isql_path Sybase_user Sybase_password Sybase_ASE Sybase_OCS

Parameter description

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

The message *RETVAL*0 indicates successful configuration. Otherwise, you receive *RETVAL**error_number*. To get the error description, execute:

```
omnigetmsg 12 error_number.
```

Example 1

To configure the Sybase instance *mysybase*, execute:

```
util_sybase.pl -CONFIG mysybase /applications/sybase.15/ \  
/applications/sybase.15/OCS-15_0/bin/isql sa " " ASE-15_0 OCS-15_0
```

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 either the Data Protector GUI or 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

Execute:

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

UNIX systems: util_sybase.pl -CHKCONF *Sybase_instance_name*

Backup

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

Table 3: Backup types

Full	Backs up all selected Sybase databases and transaction logs.
Trans ¹	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`

¹ For this backup type, the transaction logs must be placed on a separate Sybase database device. Otherwise, the backup fails. For details of how to place transaction logs on a separate Sybase database device, see the Sybase documentation.

- 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.

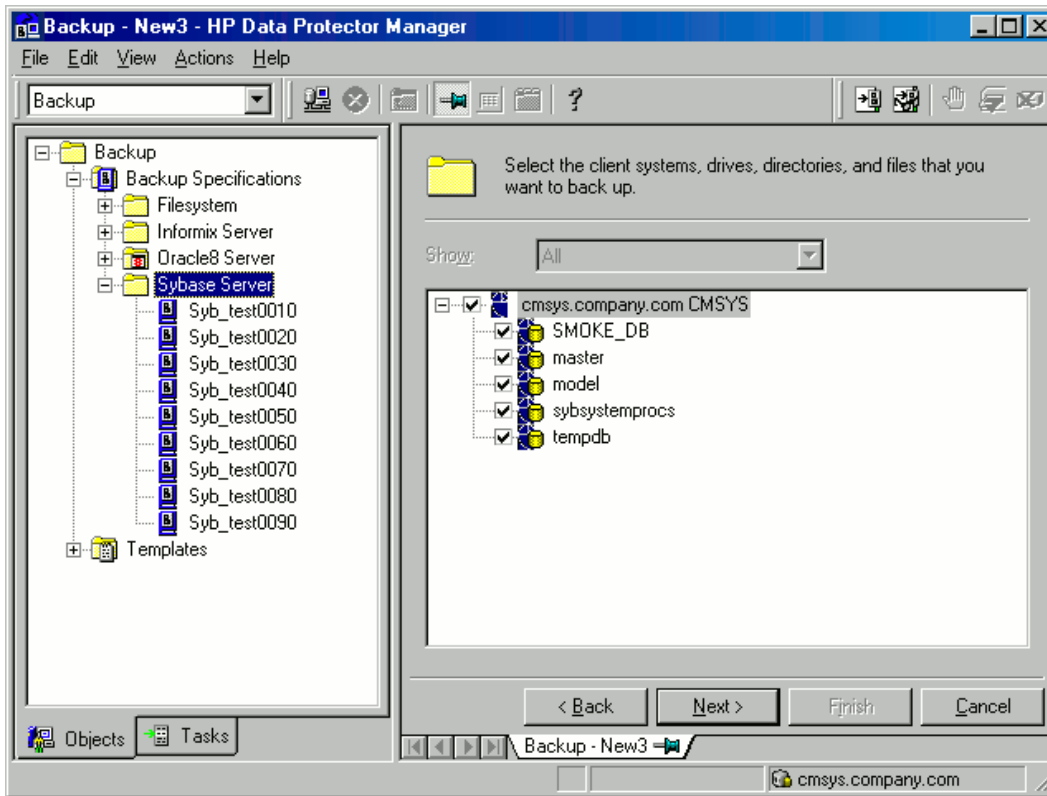
In **Application database**, type the Sybase instance name.

UNIX systems: Type `sybase` in both **Username** and **Group/Domain name**. This user becomes 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 10](#).
6. Select the databases you want to back up.

Figure 5: 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 [Sybase backup options on page 20](#).

Figure 6: Pre- and post-exec commands (Windows systems)

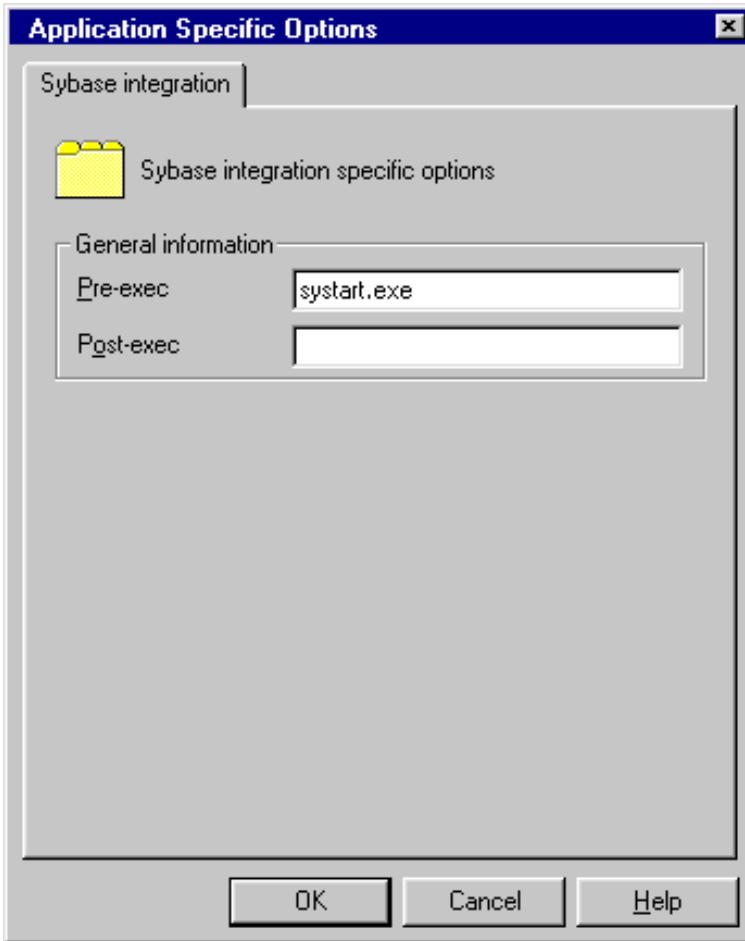
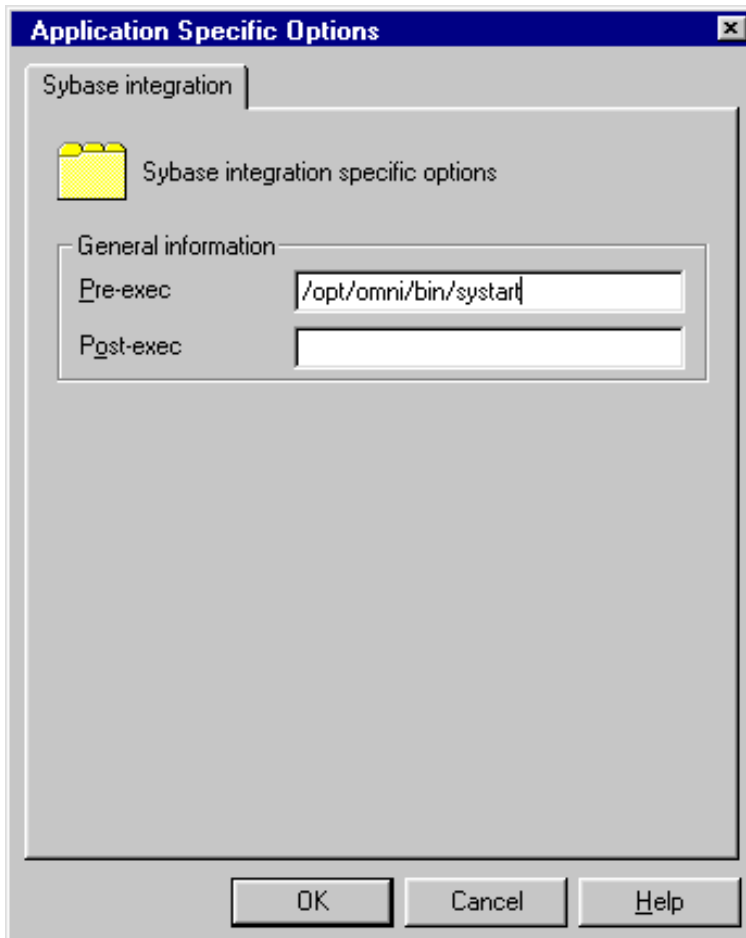


Figure 7: Pre- and post-exec commands (UNIX systems)



Click **Next**.

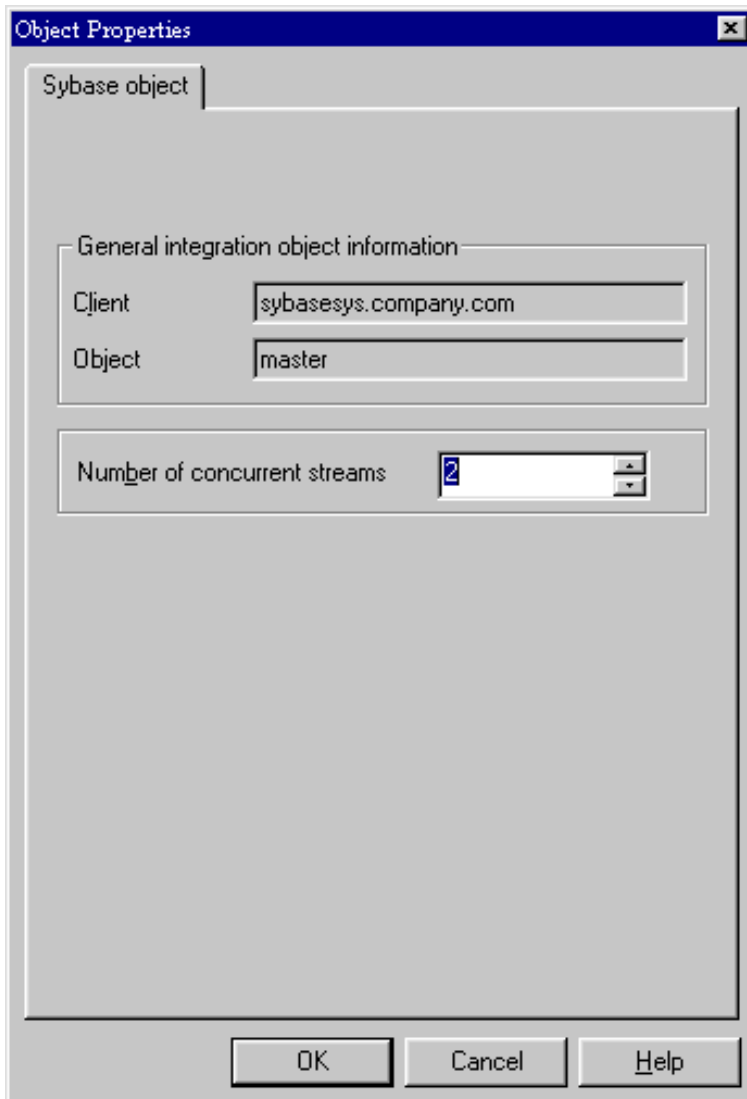
9. Optionally, schedule the backup. For more information, see [Scheduling backup sessions on page 20](#).

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*.

Figure 8: 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 21](#).

Table 4: Sybase backup options

Pre-exec, Post-exec	<p>Specify a command that will be started by <code>ob2sybase.exe</code> (Windows systems) or <code>ob2sybase.pl</code> (UNIX systems) on the Sybase Server system before the backup of every selected database (pre-exec) or after it (post-exec). Do not use double quotes.</p> <p>Windows systems: Provide only the name of the command. The command must reside in the default Data Protector commands directory. See Pre- and post-exec commands (Windows systems) on page 17.</p> <p>UNIX systems: Provide the pathname of the command. See Pre- and post-exec commands (UNIX systems) on page 18.</p>
---------------------	---

Modifying backup specifications

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

Scheduling backup sessions

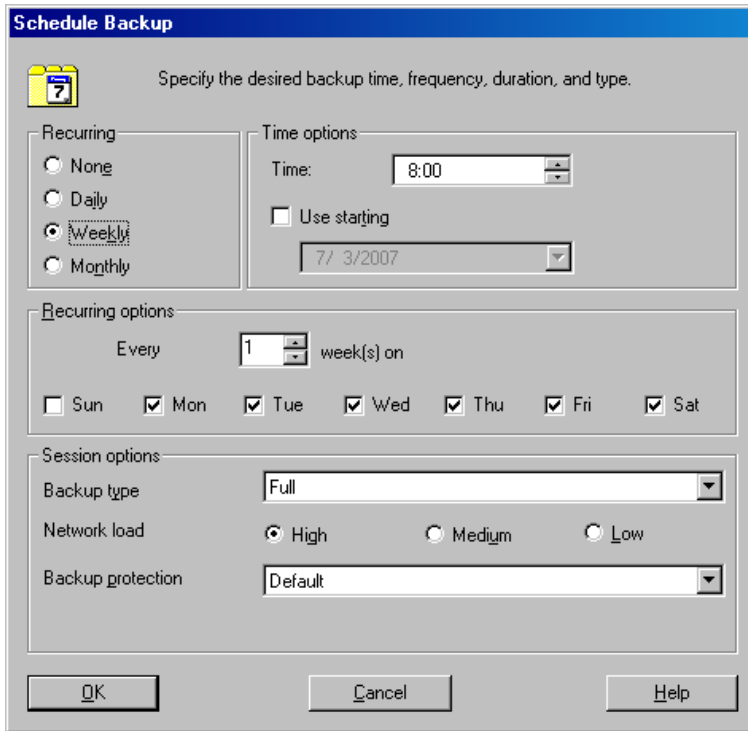
You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: “scheduled backups”.

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 [Scheduling a backup session on the next page](#). Click **OK**.
3. Repeat [In the Schedule property page, select the starting date in the calendar and click Add to open the Schedule Backup dialog box.](#) above and 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 [Scheduling a backup session on the next page](#). Click **OK.** above to schedule another backup at 13:00, and another one at 18:00.
4. Click **Apply** to save the changes.

Figure 9: Scheduling a backup session



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**.

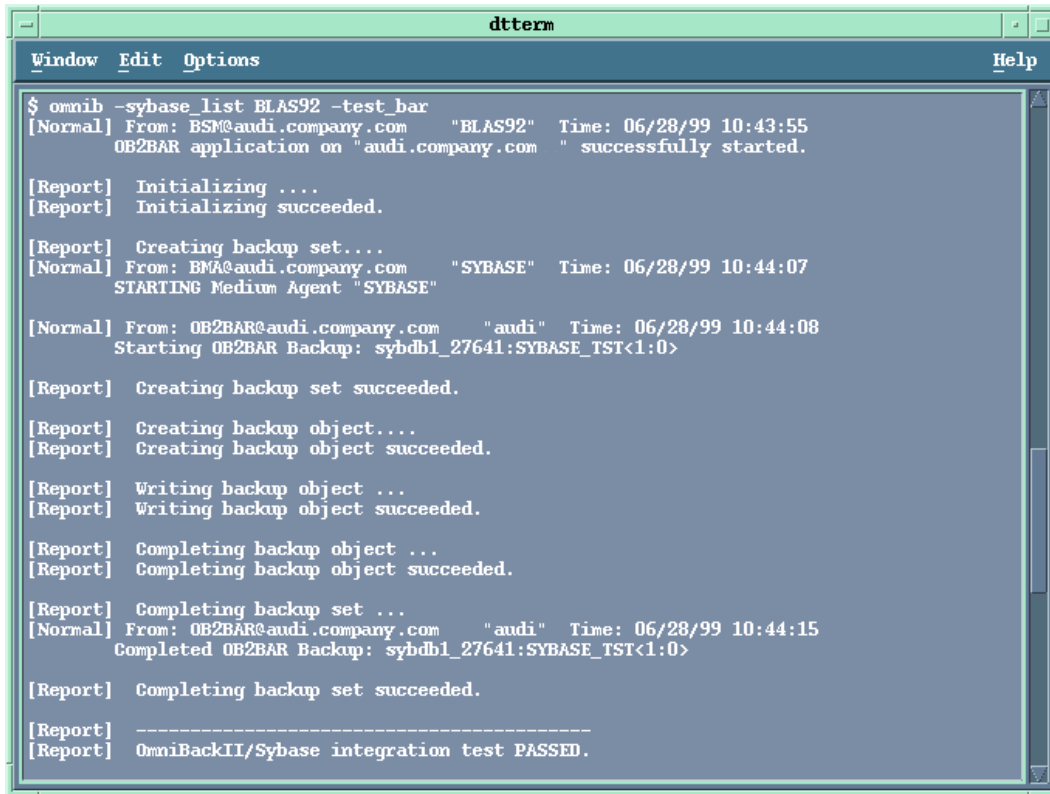
The message `Session completed successfully` is displayed at the end of a successful preview.

Using the Data Protector CLI

Execute:

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

Figure 10: Example of previewing a backup



```
dtterm
Window Edit Options Help
$ omnib -sybase_list BLAS92 -test_bar
[Normal] From: BSM@audi.company.com "BLAS92" Time: 06/28/99 10:43:55
OB2BAR application on "audi.company.com" successfully started.

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

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

[Normal] From: OB2BAR@audi.company.com "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.company.com "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.
```

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 use 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

Execute:

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

Parameter description

<code>backup_specification</code>	Name of the Data Protector Sybase backup specification.
<code>sybase_mode</code>	Backup type. Select between <code>full</code> and <code>trans</code> .
<code>options</code>	For information, see the <code>omnib</code> man page.

Example

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

```
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. Execute the Sybase `isql` command:

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

Parameter description

<i>Sybase_instance</i>	Sybase instance name.
<i>Sybase_user</i>	Sybase instance user.
<i>Sybase_password</i>	Password of the Sybase instance user.
<i>database</i>	Name of the database to be backed up.
<i>backup_specification</i>	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 systems) or code page (on UNIX systems):

1. Set the encoding used on the terminal to UTF-8.
2. **Windows systems:** 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 below](#).

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 30](#).

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

Parameter description

<i>database</i>	Database to be restored.
<i>Sybase_instance</i>	Sybase instance from which the database to be restored was backed up.
<i>date</i>	Point in time. The first backup version created after this point in time is restored. Use the 0-24h time format.
<i>new_database</i>	Target database to which to restore.
<i>new_Sybase_instance</i>	Target Sybase instance to which to restore.
<i>file</i>	Pathname of a file to which the load command or command sequence is recorded.
<i>-media</i>	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 option `OB2SybaseTransLogDelay`. The default value is 20 seconds.

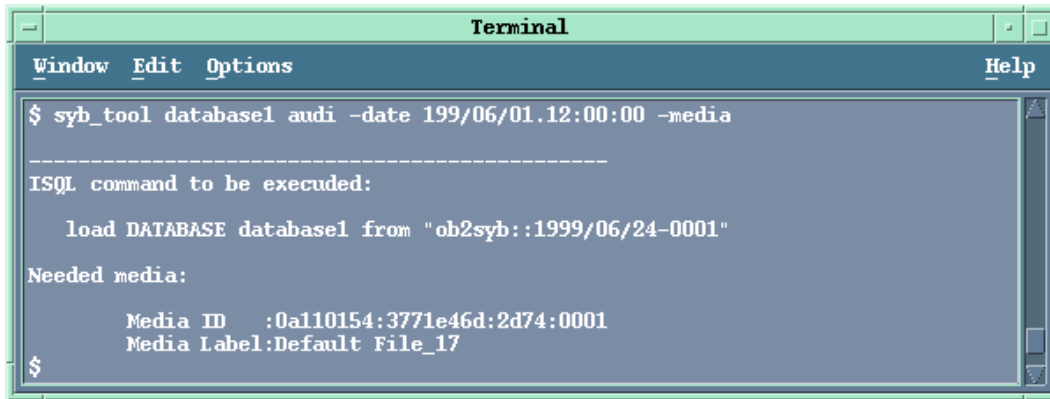
Example 1

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, execute:

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

See [Running the syb_tool command on the next page](#).

Figure 11: Running the syb_tool command



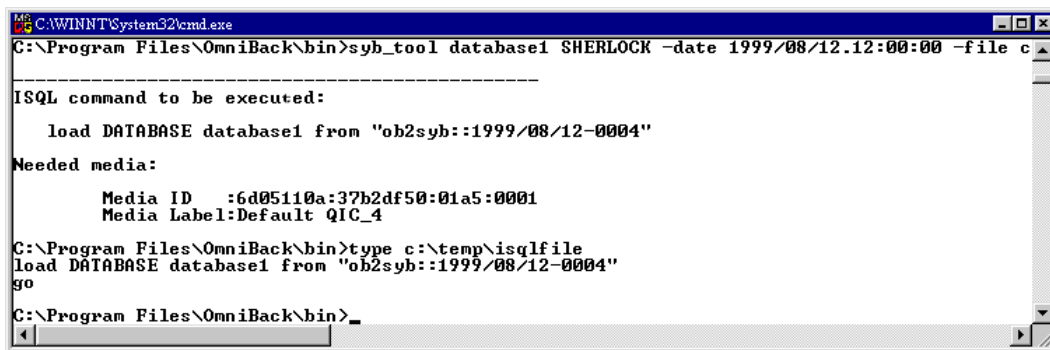
```
Terminal
Window Edit Options Help
$ syb_tool database1 audi -date 199/06/01.12:00:00 -media
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/06/24-0001"
Needed media:
    Media ID   :0a110154:3771e46d:2d74:0001
    Media Label:Default File_17
$
```

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), execute:

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

Figure 12: Running the syb_tool command with the -file and -media options



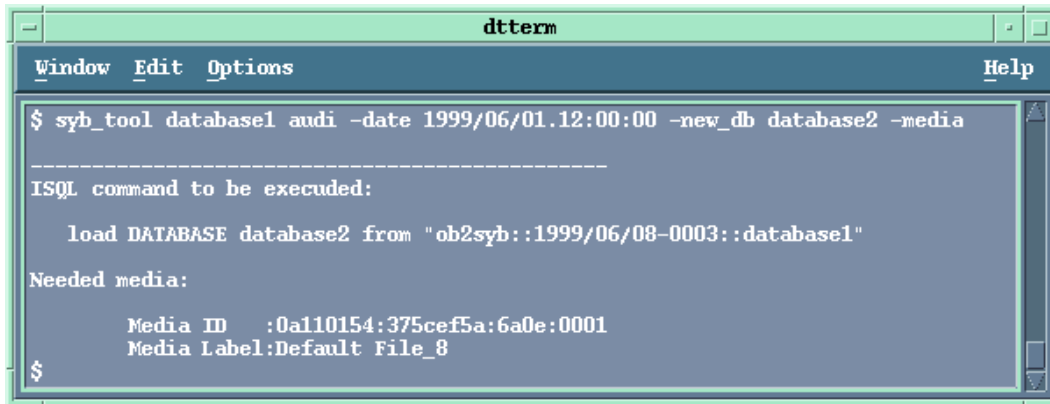
```
C:\WINNT\System32\cmd.exe
C:\Program Files\OmniBack\bin>syb_tool database1 SHERLOCK -date 1999/08/12.12:00:00 -file c
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/08/12-0004"
Needed media:
    Media ID   :6d05110a:37b2df50:01a5:0001
    Media Label:Default QIC_4
C:\Program Files\OmniBack\bin>type c:\temp\isqlfile
load DATABASE database1 from "ob2syb::1999/08/12-0004"
go
C:\Program Files\OmniBack\bin>_
```

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, execute:

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

Figure 13: The load command for restore to a different database



```
dtterm
Window Edit Options Help
$ 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:

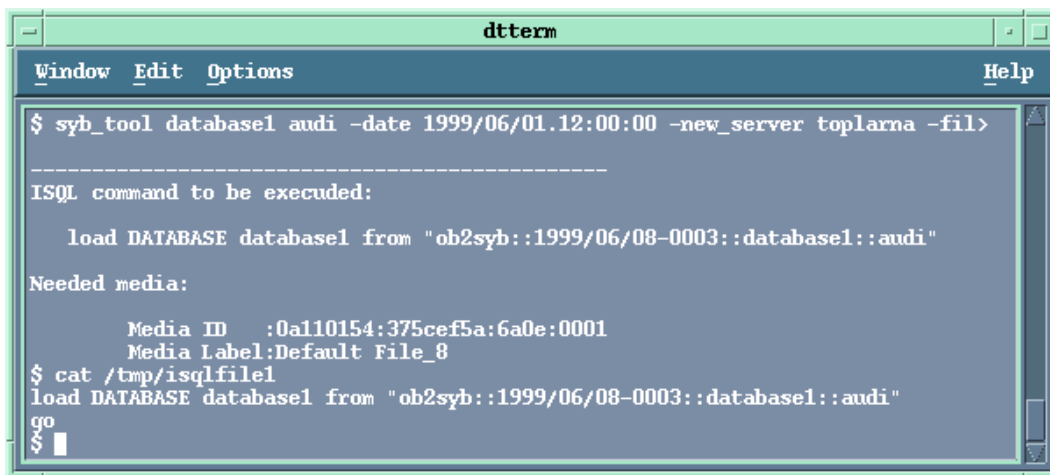
  Media ID   :0a110154:375cef5a:6a0e:0001
  Media Label:Default File_8
$
```

Example 4

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

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

Figure 14: The load command for restore to a different server



```
dtterm
Window Edit Options Help
$ 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
$
```

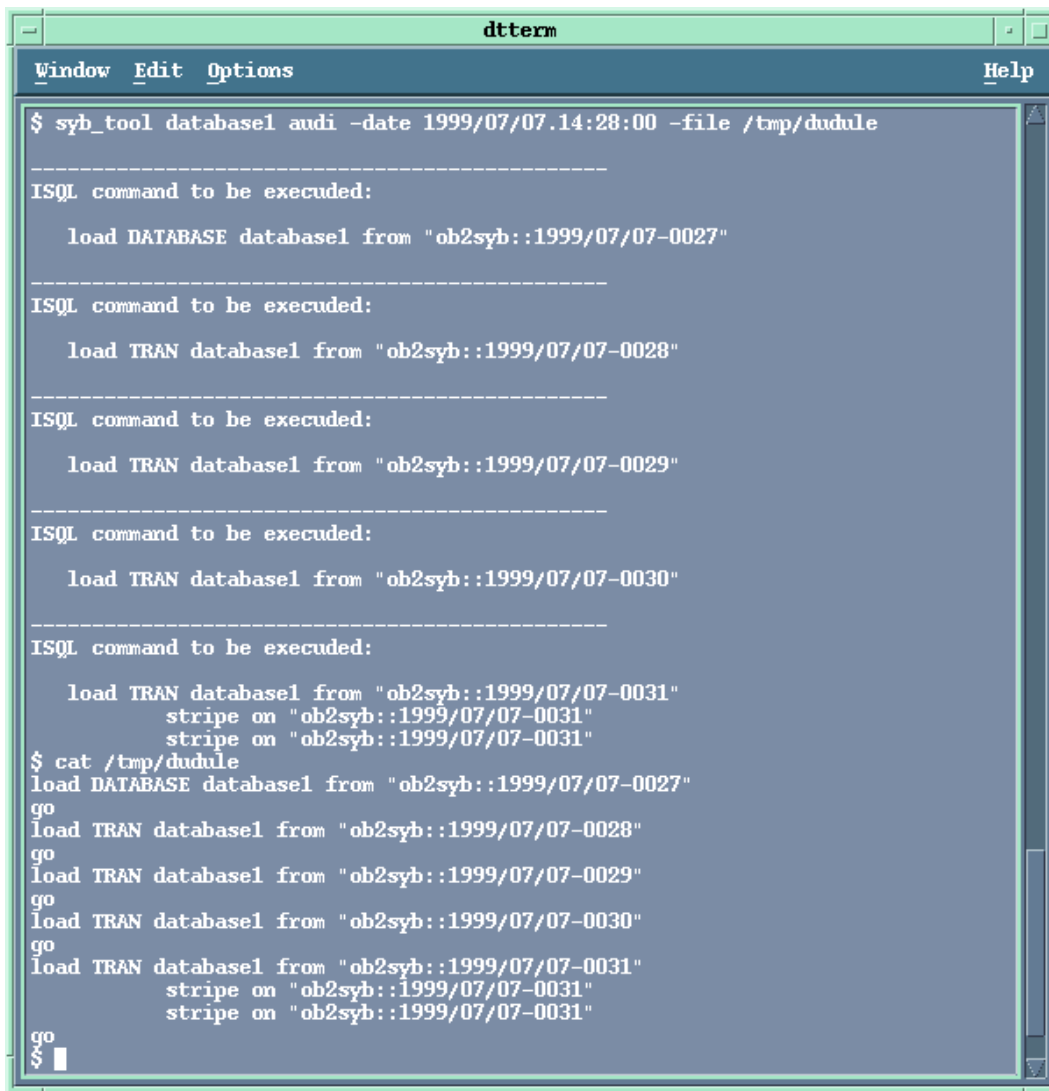
Example 5

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, execute:

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

You see in [Loading transaction logs from multiple backups on the next page](#) that you need to restore one full backup and four transaction log backups, the last one backed up with concurrency 3.

Figure 15: Loading transaction logs from multiple backups



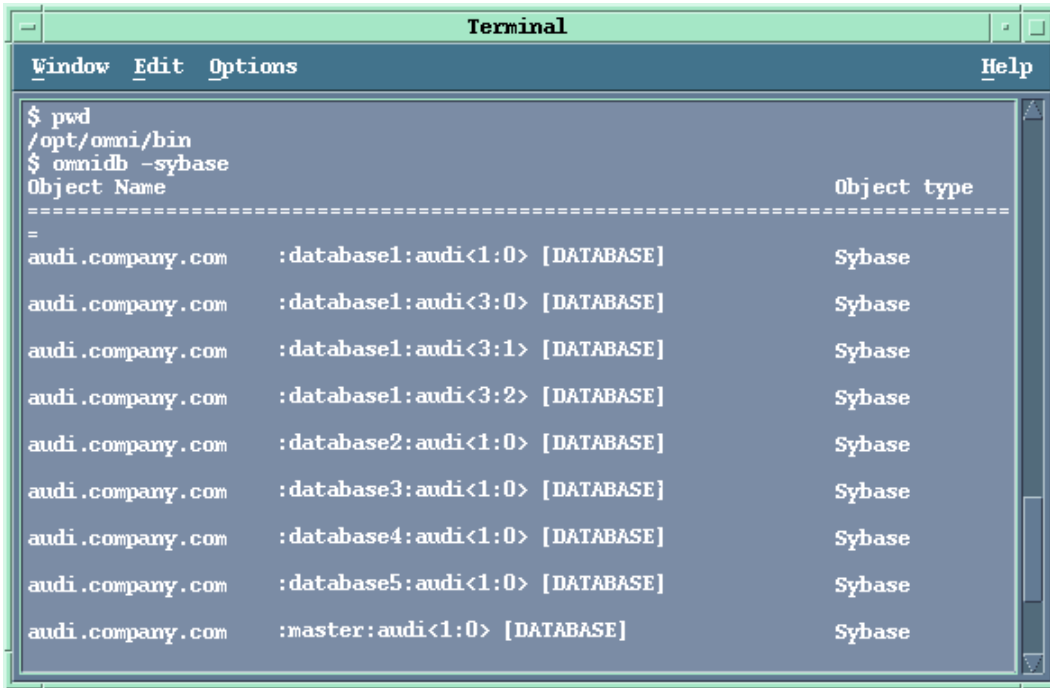
```
dtterm
Window Edit Options Help
$ syb_tool database1 audi -date 1999/07/07.14:28:00 -file /tmp/dudule
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/07/07-0027"
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0028"
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0029"
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0030"
-----
ISQL command to be executed:
    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
$
```

Using the standard Data Protector CLI commands

1. Get a list of backed up Sybase databases:

```
omnidb -sybase
```

Figure 16: Example of a list of backed up Sybase databases

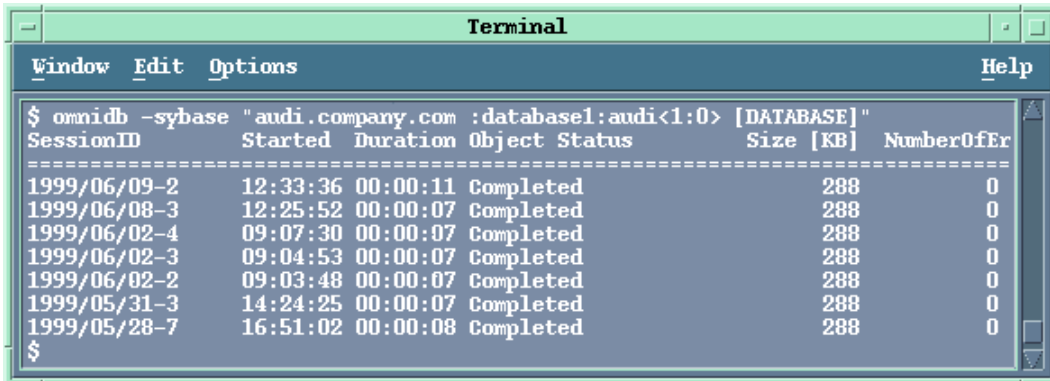


```
Terminal
Window Edit Options Help
$ pwd
/opt/omni/bin
$ omnidb -sybase
Object Name                                     Object type
-----
audi.company.com :database1:audi<1:0> [DATABASE] Sybase
audi.company.com :database1:audi<3:0> [DATABASE] Sybase
audi.company.com :database1:audi<3:1> [DATABASE] Sybase
audi.company.com :database1:audi<3:2> [DATABASE] Sybase
audi.company.com :database2:audi<1:0> [DATABASE] Sybase
audi.company.com :database3:audi<1:0> [DATABASE] Sybase
audi.company.com :database4:audi<1:0> [DATABASE] Sybase
audi.company.com :database5:audi<1:0> [DATABASE] Sybase
audi.company.com :master:audi<1:0> [DATABASE] Sybase
```

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

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

Figure 17: Example of a list of backup sessions for a specific object



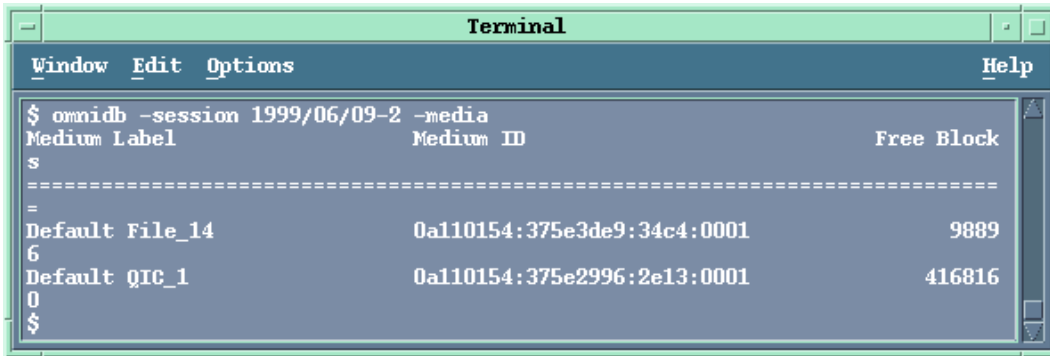
```
Terminal
Window Edit Options Help
$ omnidb -sybase "audi.company.com :database1:audi<1:0> [DATABASE]"
SessionID      Started Duration Object Status      Size [KB] NumberOfEr
-----
1999/06/09-2   12:33:36 00:00:11 Completed      288      0
1999/06/08-3   12:25:52 00:00:07 Completed      288      0
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
$
```

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
```

Figure 18: 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 systems, 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

<i>Sybase_instance</i>	Sybase instance name.
<i>Sybase_user</i>	Sybase instance user.
<i>Sybase_password</i>	Password of the Sybase instance user.
<i>input_file</i>	The file to which the load parameter was saved. See also Localized database names on page 24 .

3. If you did not provide the load command in a file, type the desired load command in the first line. To execute 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.
<i>version</i>	Session ID of the backup version to restore from. You can also type <code>latest</code> to restore from the latest backup.
<i>new_database</i>	Target database to which to restore.
<i>database</i>	Database to be restored.
<i>Sybase_instance</i>	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.

To restore a database to a different Sybase instance on another client system, set the `OB2HOSTNAME` variable on the target client: add the `OB2HOSTNAME=BackupClient.company.com` variable entry to the `Sybase_TargetInstance.cfg` configuration file, located in the default Data Protector temporary files directory.

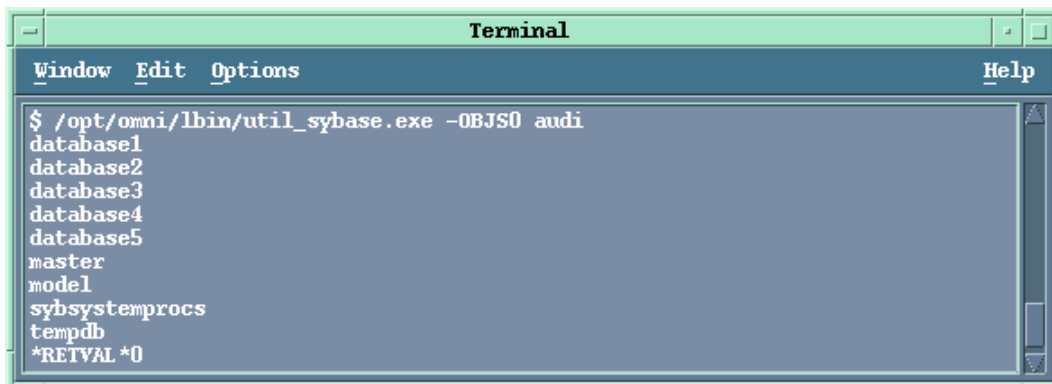
For details on the Sybase `load` command, see the *Adaptive Server Enterprise System Administration Guide*.

Tip: To list all Sybase databases of a particular Sybase instance, execute:

Windows systems: `perl -I..\lib\perl util_sybase.pl -OBS0 \`
`Sybase_instance_name`

UNIX systems: `util_sybase.pl -OBS0 Sybase_instance_name`

Figure 19: Example of a list of Sybase databases



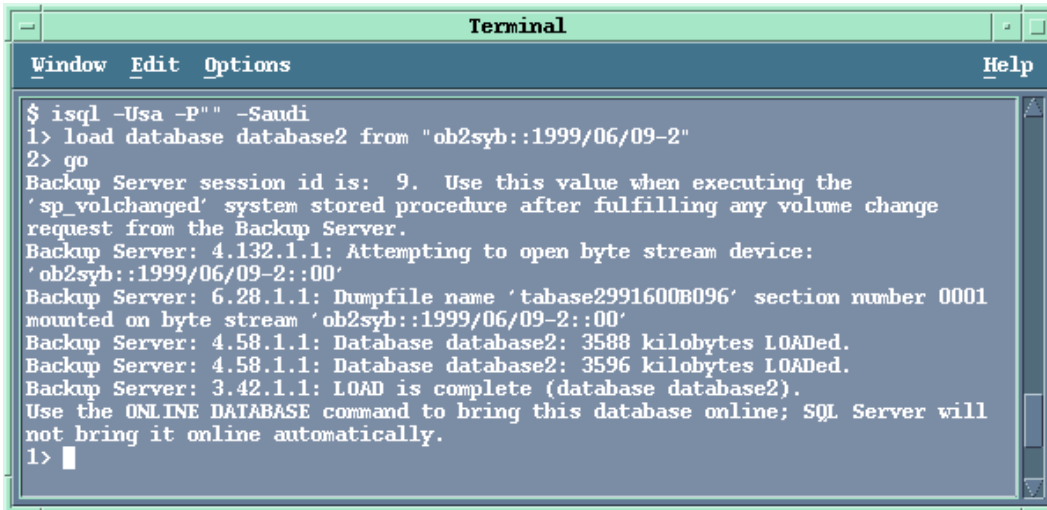
Restore examples

Example 1

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

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

Figure 20: Restoring a database from a specific session



```
Terminal  
Window Edit Options Help  
$ 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> █
```

Example 2

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

1. Create a database device. See [Creating a database device below](#).


Figure 21: Creating a database device



```
Terminal  
Window Edit Options Help  
$ 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> █
```

2. Create an empty database, named Sybdata1. See [Creating an empty database on the next page](#).

Figure 22: Creating an empty database



3. Restore Sybdata to Sybdata1 by executing:

```
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, execute:

```
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 Cyrillic and Latin characters, and for which the load command was saved in the file restore_20100609-2.txt, execute :

```
isql -S instance1 -U admin -PSybase_password -J utf8 -i restore_20100609-2.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 systems: *Data_Protector_program_data\Config\server\Cell\restoredev*

UNIX systems: */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 *HP Data Protector 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 *HP Data Protector 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://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors written to `debug.log`, located on the Sybase Server system in the default Data Protector log files directory.
- Make a test backup and restore of any filesystem on the problematic client. For information, see the *HP Data Protector 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 systems:** 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 systems: The System Properties dialog box, which you access by double-clicking System in the Control Panel.

UNIX systems: 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 13](#).
- Test the backup specification as described in [Previewing backup sessions on page 21](#).

If the Data Protector part of the test fails:

- a. **UNIX systems:** 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.
- b. 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 23](#).

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 execute:
 - 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 systems: *Data_Protector_home\help\enu*

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

- On the Sybase Server system, examine system errors, reported in the debug.log file located in the default Data Protector log files directory.

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 *HP Data Protector Help* index: “changing user rights”.

Problems

Problem

Restore to another client system fails

When you start a restore of a database to the original Sybase instance, the session finishes successfully. However, when you start a restore of the database to a different Sybase instance on another client, your restore session fails with a message similar to the following:

```
Mar 11 18:16:13 2010: Backup Server: 4.124.2.1: Archive API error
for device='ob2syb::2010/03/11-19::test_db:
:incprod::00': Vendor application name=Data Protector A.06.10,
Library version=221, API routine=syb_read(), Message=Object version
not found.ar 11 18:16:13 2010: Backup Server: 6.32.2.3: ob2syb::
2010/03/11-19::test_db::incprod::00: volume not valid or not requested
(server: , session id: 62.) Mar 11 18:20:07 2010: Backup Server:
4.132.1.1: Attempting to open byte stream device: 'ob2syb::
2010/03/11-19::test_db::incprod::00'
```

The problem is that the IDB uses the name of the destination client instead of the name of the client from which the database was backed up.

Action

1. Set the `OB2HOSTNAME` variable on the target client: add the `OB2HOSTNAME=BackupClient.company.com` variable entry to the `Sybase_TargetInstance.cfg` configuration file, located in the default Data Protector temporary files directory.
2. Restart the restore of the database.

Chapter 2: Data Protector Network Data Management Protocol Server integration

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 (**NAS**) 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
- Incr1

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

The Data Protector NDMP Server integration supports the following two types of backup:

- for EMC Celerra (**Celerra**):
 - Dump
The default backup type, that backs up data at a file level.
 - NDMP volume backup (**NVB**)
An EMC-specific NDMP backup type, that backs up data blocks at a volume level.
- for Network Appliance (**NetApp**):
 - Dump
The default backup type, that backs up data at a file level.
 - Snap mirror to tape backup (**SMTape backup**)


A NetApp-specific NDMP backup type, that creates a snapshot of the source volume and backs up the current and all previous snapshot copies.

This chapter provides information specific to the Data Protector NDMP Server integration. For general Data Protector procedures and options, see the *HP Data Protector 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, through the NDMP interfaces, the NDMP Server. [Data Protector NDMP Server integration architecture below](#) shows the architecture of the integration.

Figure 23: 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 the following NDMP Server types:

- NetApp NAS device (**NetApp**)
- Celerra NAS device (**Celerra**)
- Hitachi BlueArc NAS device (**BlueArc**)
- Hitachi NAS device (**Hitachi**)
- HP-X9000 NAS device (**HP-X9000**)

In a typical environment ([The NDMP environment configuration below](#)), 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 24: 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 default Data Protector temporary files directory.

For more information on file history swap files, see [The NDMP specific omnirc options on page 64](#).

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://support.openview.hp.com/selfsolve/manuals>.
 - For information on installing, configuring, and using NDMP Server, see the NDMP Server documentation.
- To use the Data Protector media copy or automated media copy functionality, make sure that the NetApp Server and EMC Celerra Server support NDMP v4 protocol.
- Ensure that you have correctly installed Data Protector. For information of 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.

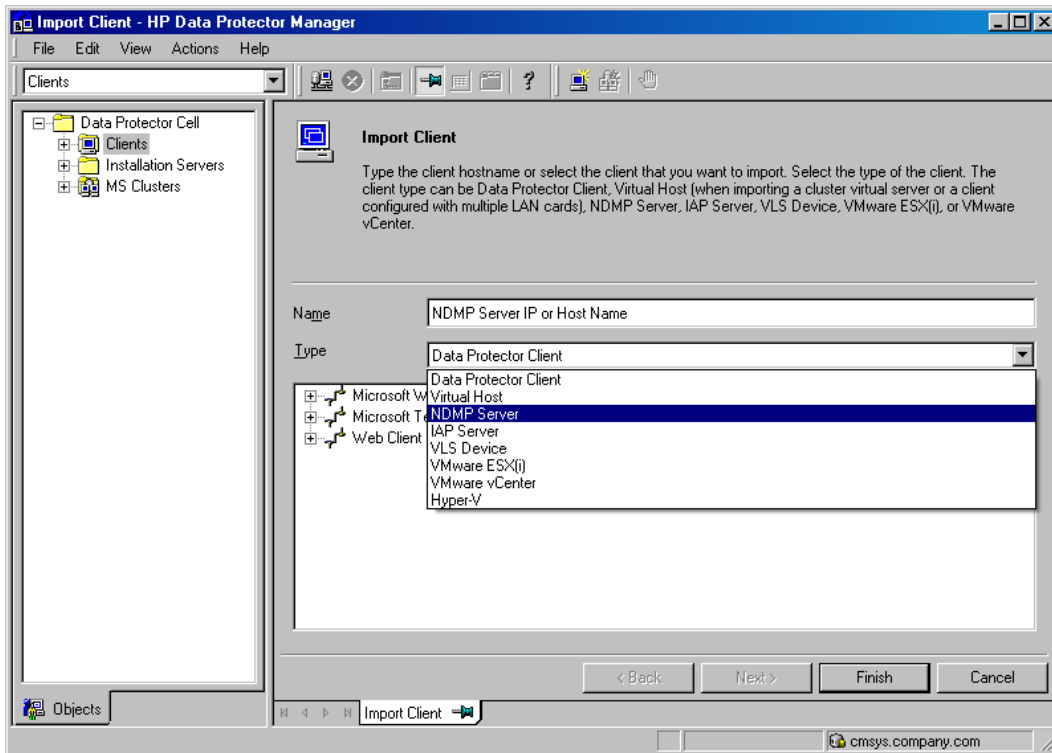
Importing NDMP Server systems

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 the **Name** text box, type the name of the NDMP Server system you want to import.

In the **Type** drop-down list, select **NDMP Server**.

Figure 25: Specifying an NDMP Server system



Click **Next**.

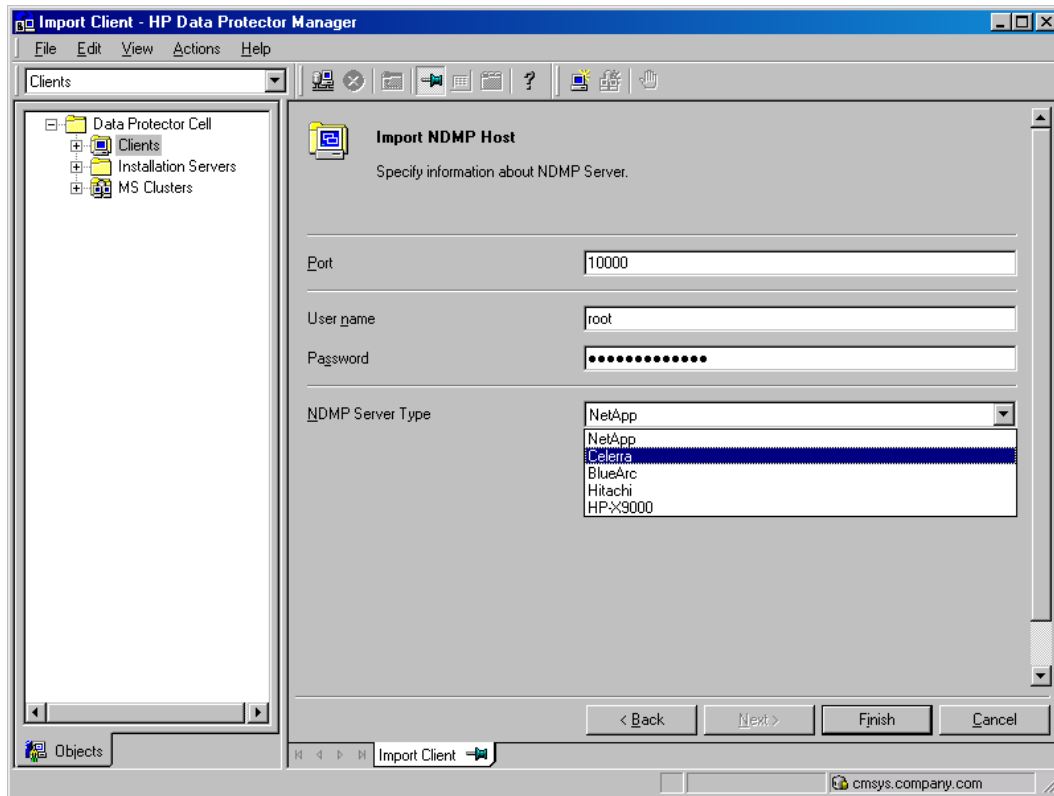
4. In the **Port** text box, specify the TCP/IP port number of the NDMP Server. The default number is 10000.

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.

In the **NDMP Server Type** drop-down list, select the NAS device type.

Figure 26: Specifying an NDMP Server system



Click **Finish**.

Creating media pools

Create a special media pool for NDMP media. For information, see the *HP Data Protector Help* index: "creating media pools".

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

Prerequisites

- The source medium and the target medium used for media copying should be of the same media type.

Limitations

- A medium cannot be used by different NDMP Server types. Consequently, the data that was backed up from an NDMP Server of a particular type (for example, NDMP-NetApp) cannot be restored to an NDMP Server of another type (for example, NDMP-Celerra).

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.

- The source and the target drive used for media copy should be connected to the same NDMP Server on which the backup was performed.

Library robotics can be connected to:

- NDMP Server system ([Library configuration 1 below](#)).
- NDMP client ([Library configuration 2 on the next page](#)).
- Data Protector client with the general Media Agent installed (**general Media Agent client**) ([Library configuration 2 on the next page](#)).

If the robotics is connected to the NDMP Server system, it must be supported by both NDMP Server and Data Protector. The robotics can only be controlled by a Data Protector NDMP client.

Figure 27: Library configuration 1





Figure 28: Library configuration 2



Multiple library drives can be connected to one NDMP Server system. If the library robotics is connected to a Data Protector NDMP client, and the drives are connected to the NDMP Server system, the drives can be shared between multiple NDMP Server systems and general Media Agent clients as well as between Data Protector and other applications. NDMP media on the other hand cannot be shared. 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 [Configuring a library on the next page](#).

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

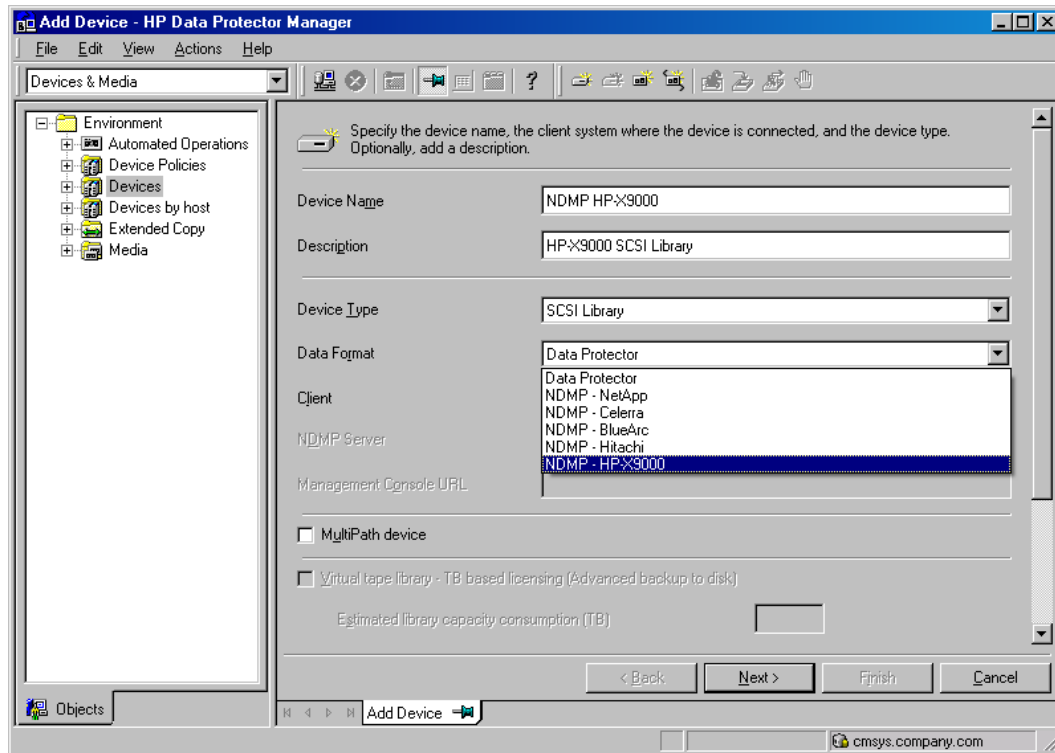
In **Interface Type**, select the NAS device used.

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.

Figure 29: Configuring a library



Click **Next**.

4. Specify the library robotics SCSI address and the drive handling. For information, see [Network Appliance configuration on page 48](#), [EMC Celerra configuration on page 49](#), and [Hitachi BlueArc or Hitachi configuration on page 50](#)

Click **Next**.

5. Specify the slots to be used by Data Protector.

Click **Next**.

6. In the **Media Type** drop-down list, select the media type used in the library.
7. Click **Finish** and then click **Yes** to configure the drives in the library.
8. Type a name for the drive. Optionally, describe the drive.

In **Data Format**, 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.

Click **Next**.

9. Specify the SCSI address of the drive. For information, see [Network Appliance configuration on page 48](#), [EMC Celerra configuration on page 49](#), and [Hitachi BlueArc or Hitachi configuration on page 50](#).

Do not change the drive index number.

Click **Next**.

10. Specify the media pool for the NDMP media.

To specify advanced device options, click **Advanced**. For information on supported block sizes, see [Block size on page 51](#).

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

Click **Next**.

11. Select the device policies for the new drive and specify the device tag.

Click **Finish**.

12. Click **Yes** to create another drive or **NO** to finish the configuration.

For information of 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 *HP Data Protector Help* index: "configuring SCSI libraries". Then configure the drives as described in [Type a name for the drive. Optionally, describe the drive. on the previous page](#) through [Click Yes to create another drive or NO to finish the configuration. above](#).

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.

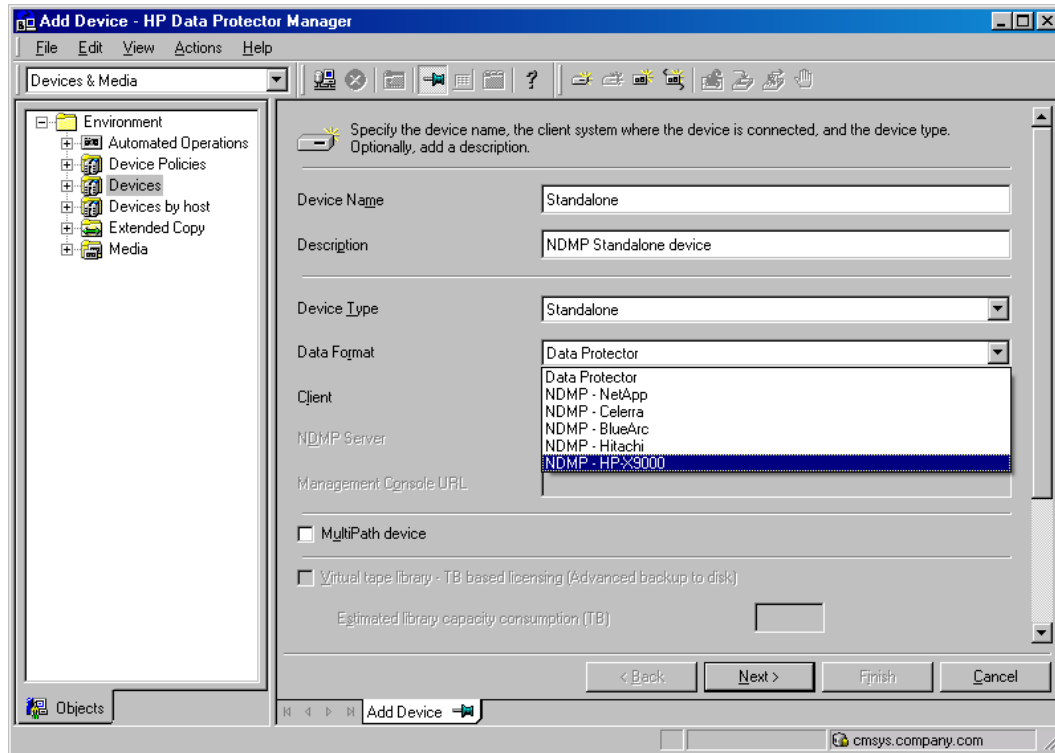
In **Device Type**, select **Standalone**.

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.

Figure 30: Configuring a standalone device



Click **Next**.

4. Provide the SCSI address of the device. For information, see [Network Appliance configuration on the next page](#), [EMC Celerra configuration on page 49](#), and [Hitachi BlueArc or Hitachi configuration on page 50](#).

Click **Next**.

5. Specify the media pool.

To specify advanced device options, click **Advanced**. For information on supported block sizes, see [Block size on page 51](#).

Note: Multiplexing data streams 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, execute:

```
sysconfig -t
```

on the NDMP Server system. The SCSI address is written at the beginning of the output and consists of four parts. See [Analyzing the drive's SCSI address below](#).

Table 5: 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.
{1 m h a}	Data density and compression.

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, execute:

```
sysconfig -m
```

on the NDMP Server system. The SCSI address consists of two parts. See [Analyzing the library Robotics' SCSI address on the next page](#).

¹ Data Protector supports only the no rewind devices.

Table 6: 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. execute:

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

Example

See [Example of a list of SCSI devices below](#) 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 7: 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\$

Hitachi BlueArc or Hitachi configuration

Before you begin

- Ensure that the NDMP Server is online.

Important: To properly configure a SCSI device connected to a Hitachi BlueArc server, two clients must be connected to this device and configured in Data Protector. NDMP Media Agent client is used for data transfer between a SCSI device and Hitachi BlueArc, while a separate client with the General Media Agent component installed should be configured as a SCSI device (not an NDMP SCSI device), to provide the SCSI information about the device.

The following sections explain how to obtain information for configuring a Hitachi or a Hitachi BlueArc NDMP Server. For configuration of your device, see [Configuring tape libraries on page 44](#) or [Configuring standalone devices on page 46](#).

Library robotics

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

1. Log in to the NDMP Server system.
2. Get the Enterprise Virtual Server ID (EVS ID) of the NDMP Server system for which you need to configure a device. Execute the command:

```
evs list
```

3. Get the SCSI address of the library robotics connected to the NDMP Server system. Execute the command:

```
ndmp-devices-list -t changer -v EvsID
```

The SCSI address is written at the beginning of the `ndmp-device-list` output and consists of the LUN, Target and device ID number. See [Analyzing the library robotics' SCSI address below](#).

Table 8: Analyzing the library robotics' SCSI address

Parts	Description
dev/mc	Media changer device
d2 1 2 3...	Device ID number.

Example

The output for a media changer device of library robotics is:

```
5/dev/mc_d2|0 01YFPdba00 0x2001f29ccd2ca000:0 N/A
```

Where `/dev/mc_d2|0` is the SCSI address of the media changer device.

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:

1. Log in to the NDMP Server system.
2. Get the Enterprise Virtual Server ID (EVS ID) of the NDMP Server system for which you need to configure a device. Execute the command:

```
evs list
```

3. Get the SCSI address of standalone tape devices or drives in a tape library connected to the NDMP Server system. Execute the command:

```
ndmp-devices-list -t tape -v EvsID
```

The SCSI address is written at the beginning of the `ndmp-device-list` output and consists of the LUN, Target and device ID number.

Table 9: Analyzing the drive's SCSI address

Parts	Description
dev/mt	Tape device
d2 1 2 3...	Device ID number.

Example

The output for a standalone tape device (or a drive in a tape library) is:

```
16/dev/mt_d2|1 01YFPdba01 0x2001f29ccd2ca000:1
```

Where `/dev/mt_d2|1` is the SCSI address of the tape device.

Block size

The integration supports variable tape block sizes. Before selecting the block size for each NAS device, you should consider the following:

- Ensure that the NDMP Server is configured to support variable block size.
- The device used for restore must have the same or greater block size than the one that was used for backup.
- By NetApp, for a SMTape backup on Data ONTAP version earlier than version 8.0, you must set the **Block size** option to 240 kB. The required block size for the Data ONTAP version 8.0 is between 4 kB to 256 kB.

Note that Data Protector supports block sizes between 8 kB to 1024 kB. The default block size is 256 kB.

- By Celerra, the block size value should not be greater than the Celerra `readWriteBlockSizeInKB` parameter.

Tip: To get the current value of the `readWriteBlockSizeInKB` parameter, execute:

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

Note:

- 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.
- 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.

Backup

Limitations

- Only filesystem backup is supported.
- You cannot store an NDMP backup and a standard Data Protector backup on the same medium.
- Device concurrency is limited to 1.
- You cannot browse devices and filesystems.
- Object copying and object mirroring are not supported.
- Media copying and automated media copying are not supported for NDMP-Celerra backup sessions.
- Backing up data using the Data Protector **Reconnect broken connections** functionality is not supported.
- The NVB backup type enables you to only back up entire file systems. For example, you can back up `/ufs1`, but not `/ufs1/dir1`.
- The NVB backup type is supported on EMC Celerra DART version 5.6.46.11 or later.
- Directory direct access restore (DDAR) cannot be used with backup images created with the NDMP volume backup (NVB) option selected.

- The NVB backup type and file or directory filtering cannot be used together. If both are used, NVB takes precedence and the filters have no effect.
- With the SMTape backup type, a backup image of a volume in a particular aggregate type cannot be used for restore to a volume in a different aggregate type.
- With the SMTape backup type, a backup image of a volume in a regular aggregate cannot be used for restore to a volume in a larger aggregate, and the other way round.
- The SMTape backup type offers only full backup (level-0 backup).
- The SMTape backup type enables you to only back up entire file systems. For example, you can back up /ufs1, but not /ufs1/dir1.
- 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

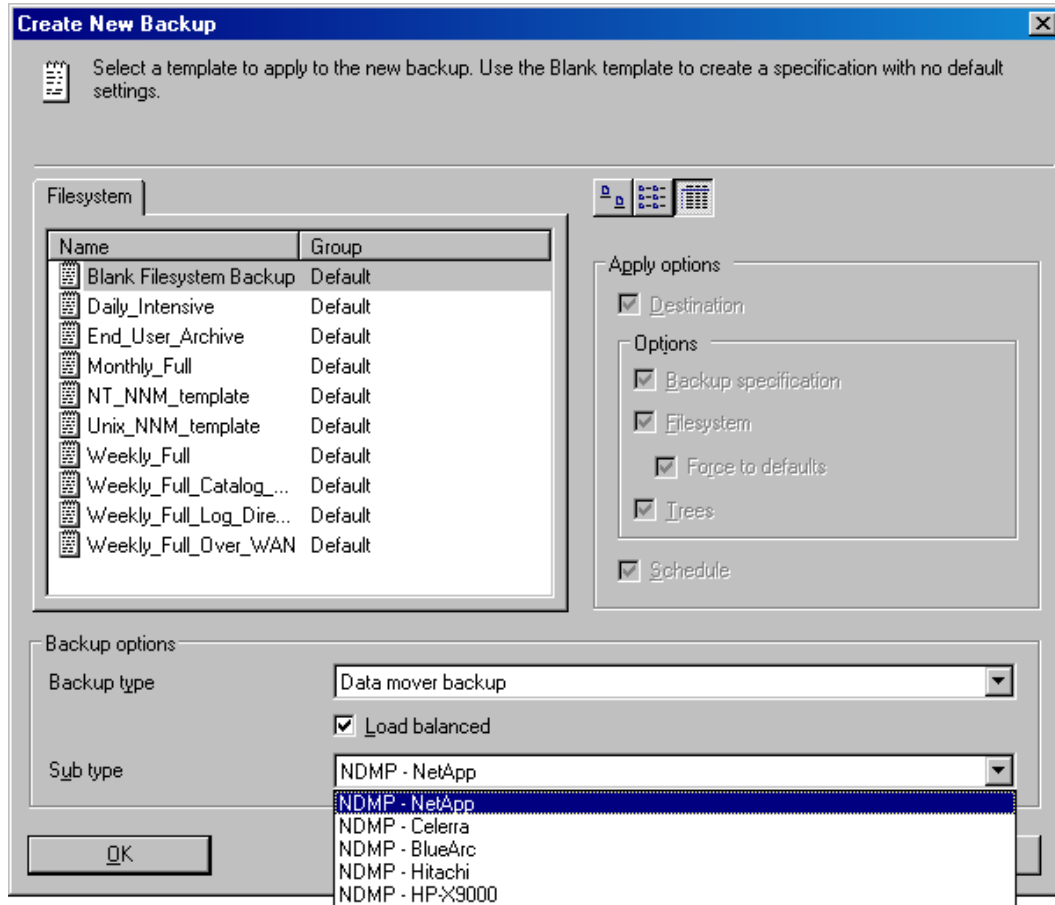
- Ensure that media to be used are formatted.
- **NetApp systems:** Get information about filesystems exported from the NDMP Server system by executing `exportfs`.

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 the NDMP Server type (for example, **NDMP-NetApp**). Optionally, select the **Load balanced** option. See [Selecting a backup template on the next page](#).

Figure 31: Selecting a backup template



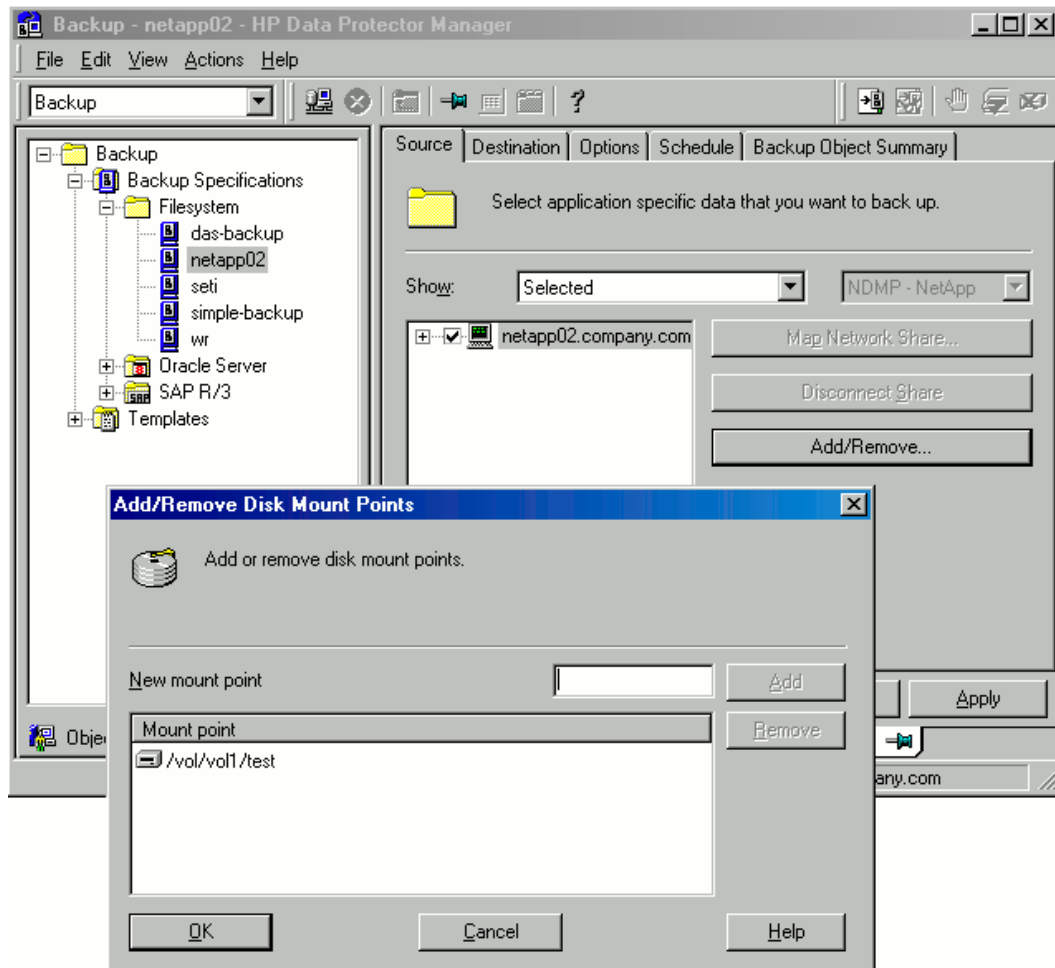
Click **OK**.

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 [Specifying the NDMP Server mountpoints for backup \(UNIX systems\) on the next page](#).

Click **OK**.

Figure 32: Specifying the NDMP Server mountpoints for backup (UNIX systems)



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**.

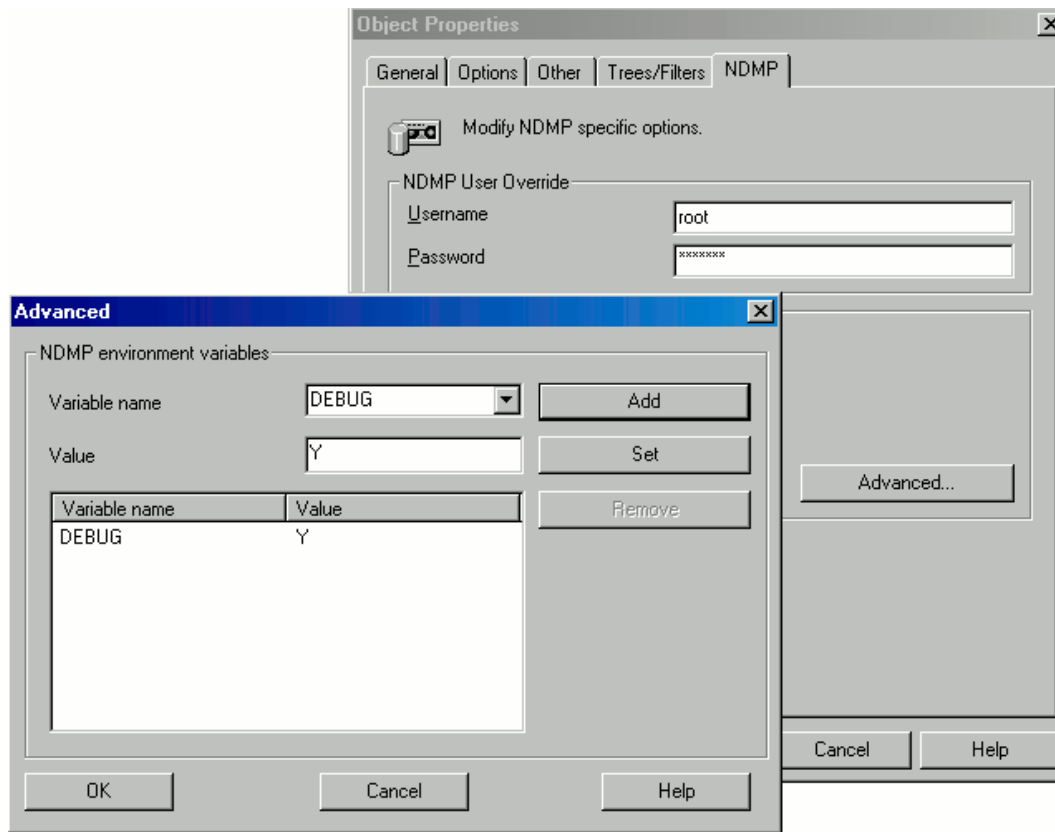
8. Review the summary of the backup specification.

To specify the NDMP 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 selected NAS device system.

To set the NDMP environment variables, click **Advanced**. See [Specifying advanced NetApp options below](#). For more information, see [NDMP environment variables on page 61](#).

Figure 33: Specifying advanced NetApp options



For an EMC Celerra NDMP client, in **NDMP backup type**, select either **dump** or **NVB**.

For a NetApp NDMP client, in **NDMP backup type**, select either **dump** or **SMTape**.

Click **Next**.

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

Tip: Preview backup session for your backup specification before using it. For details, see the *HP Data Protector Help* index: "previewing a backup".

Backing up large data set

There is no limit to the number of files that you can backup in a single NDMP backup specification. However, performance can be significantly impacted if minimum configurations are not met. For a large backup data set the following recommendations should be followed:

- Set the OB2NDMPMEMONLY variable to 0.
- Minimum of 4 GB of memory for the NDMP media server.
- 64 bit compute platform for the NDMP media server.
- Dual core or larger CPU for the NDMP media server.

Note: This is applicable only for NetApp ONTAP systems.

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 use and click **Start Backup**.
3. Select a Backup type and Network load. Click **OK**.

Restore

Restore filesystems using the Data Protector GUI or CLI.

Prerequisites

- To enable restoring individual files or directories from a NetApp NAS device that uses the ONTAP 8.1.x operating system, set either the ENHANCED_DAR_ENABLED NDMP environment variable in the restore wizard or the ENHANCED_DAR_ENABLED omnirc option.

For details, see [NDMP environment variables on page 61](#) and [The NDMP specific omnirc options on page 64](#).

Limitations

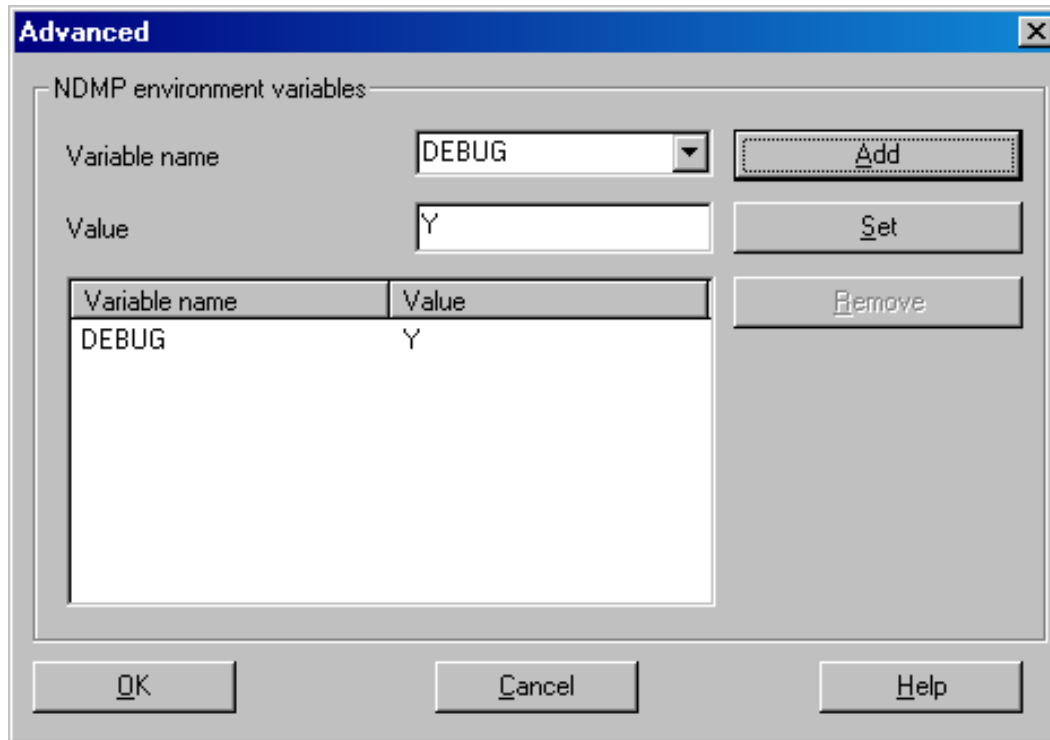
- Once you have selected a directory, you cannot exclude any subdirectories or files from restore. Specifically, the following options are not supported:
 - Data Protector GUI options: `Restore only` and `Skip`.
 - Data Protector `omnir` command: `-only`, `-skip` and `-exclude`.
- The data that was backed up from an NDMP Server of a particular type (for example, NDMP-NetApp) cannot be restored to an NDMP Server of another type (for example, NDMP-Celerra).
- When restoring to another NDMP Server, the device to restore from must be connected directly to the target NDMP Server, and the device must be selected or specified as the restore device in the Data Protector GUI or CLI.
- Restore preview is not supported.
- Restoring data using the Data Protector **Restore by Query** functionality is not supported.

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 a restore target client for every selected object. By default, data is restored to the original location, from where the data was originally backed up. To restore to a new location, select **Restore to new location** and type the new path.
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** ([NDMP advanced restore options on the next page](#)). For more information, see [NDMP environment variables on page 61](#).

Figure 34: NDMP advanced restore options



6. In the Devices page, select devices you want to use for the restore.

For more information, see the *HP Data Protector Help* index: "restore, selecting devices for".

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**.

Note: Select **Display statistical information** to view the restore profile messages in the session output.

12. Click **Finish** to start the restore.

The statistics of the restore session, along with the message `Session completed successfully` is displayed at the end of the session output.

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.

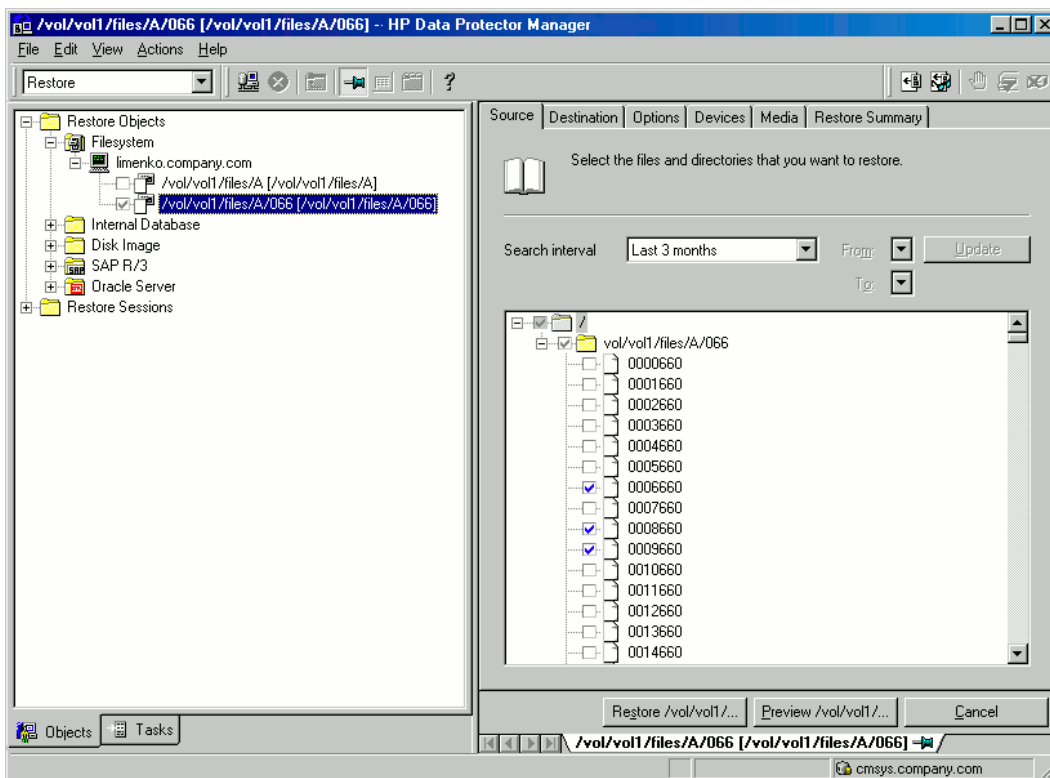
During restore, Data Protector first computes which segment contains the requested file or directory, then locates the segment, and finally starts reading through it to locate the beginning of the file or directory.

Prerequisites

File history tracking must be turned on during the backup. For information of how to enable file history tracking, see [NDMP environment variables on the next page](#).

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 [Selecting NDMP Server Data for direct access restore below](#).

Figure 35: Selecting NDMP Server Data for direct access restore



Limitations

- Directory direct access restore (DDAR) is not supported when using Hitachi BlueArc or Hitachi NAS devices.

- If you select both a directory and individual files from another directory and start the restore, only the selected files are restored. To restore both, use standard restore (set the NDMP environment variable DIRECT to N).
- **NetApp:**
 - Direct access restore (DAR) of files is supported on ONTAP version 6.1.x and later.
 - Directory direct access restore (DDAR) is supported on ONTAP version 6.4.x and later.
 - Direct access restore (DAR) of files cannot be used with backup images created with the Snap mirror to tape backup (SMTape backup) option selected.
 - Directory direct access restore (DDAR) cannot be used with backup images created with the Snap mirror to tape backup (SMTape backup) option selected.
- **Celerra:**
 - Directory direct access restore (DDAR) cannot be used with backup images created with the NDMP volume backup (NVB) option selected.

Recommendations for single-file restore from a large backup set

A single-file restore typically takes a long time if the backup data set is very large. Also, the time taken for a single-file restore exponentially increases with the increase in backup data set size. For example, the approximate time taken for a single-file restore from a backup set of 10 million is about 1.5 hours, whereas it increases to about 7 hours from a backup set of 30 million files.

Note: This is applicable only for NetApp ONTAP systems.

Restoring using another device

You can restore using a device other than that used for a backup. For more information, see the *HP Data Protector Help* index: "restore, selecting devices for".

NDMP environment variables

Set the NDMP environment variables for the selected NAS devices using the Data Protector GUI. See [Specifying advanced NetApp options on page 56](#) and [NDMP advanced restore options on page 59](#).

The following tables show the supported NDMP environment variables:

Table 10: NDMP variables for NetApp NAS device

Variable	Value	Function
HIST	y/n Default: y	Turns on/off file history tracking.
DIRECT	y/n Default: y	Enables direct access restore.
LEVEL	0, 1, 2, ... 9 Default: 0 (full)	Specifies backup level.
SMTAPE_SNAPSHOT_NAME ¹	<i>Snapshot_copy_name</i> Default: Invalid	Specifies the snapshot copy name. The specified snapshot and all older snapshot copies are backed up to a tape.
SMTAPE_DELETE_SNAPSHOT	y/n Default: n	Deletes the auto snapshot copy created after the backup.
SMTAPE_BREAK_MIRROR	y/n Default: n	Disconnects the SnapMirror after the restore. Note: After a successful restore, the restored volume is in the restricted state and does not become writable unless the SMTAPE_BREAK_MIRROR variable is set to y.
ENHANCED_DAR_ENABLED	F/T Default: T	When set to F, this variable enables restoring individual files or directories from a NetApp NAS device that uses the ONTAP 8.1.x operating system.

¹Supported only on the Data ONTAP version 8.0.7 or later.

Table 11: NDMP variables for Celerra NAS device

Variable	Value	Function
HIST	y/n Default: y	Turns on/off file history tracking.
DIRECT	y/n Default: y	Enables direct access restore.
LEVEL	0, 1, 2, ... 9 Default: 0 (full)	Specifies backup level.
BASE_DATE	<i>32bit level 32bit date</i>	Incremental backup based on a specific date.
OPTIONS	LK	Follow symbolic links.
	AT	Preserve access time.
	NT	Save NT attributes.
	MI/MD/MM	Restore collision policy for localization.

Table 12: NDMP variables for Hitachi BlueArc and Hitachi NAS device ¹

Variable	Value	Function
HIST	y/n Default: y	Turns on/off file history tracking.
DIRECT	y/n Default: y	Enables direct access restore.
LEVEL	0, 1, 2, ... 9 Default: 0 (full)	Specifies backup level.
TYPE	dump/tar Default: dump	Specifies backup type.
UPDATE	y/n Default: y	Keeps a record of the backup time. Later incremental backups can use this backup as a base.

¹ For information on other variables, see the vendor-specific documentation.

Variable	Value	Function
FILESYSTEM	<i>directory_name</i> Default: none	Specifies the directory to be backed up.
EXCLUDE	A separated list of files to be excluded from backup. Default: none	Specifies files or directories to be excluded from the backup.

Note: You can also set some NDMP environment variables using the `omnirc` file. For more information, see [The NDMP specific omnirc options below](#).

The NDMP specific omnirc options

For details of how to set the `omnirc` options, see the *HP Data Protector Help* index: "omnirc options".

Note: You can also set some options using the Data Protector GUI. See [Specifying the NDMP Server mountpoints for backup \(UNIX systems\) on page 55](#), [Specifying advanced NetApp options on page 56](#), and [NDMP environment variables on page 61](#).

The GUI setting overrides the setting in the `omnirc` file.

The NDMP specific `omnirc` options are:

- **OB2NDMPFH** (Y/N)

Default: 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: 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: 1

This option 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 option to 0 whenever the number of files in the backup specification exceeds 5 million.

As an example, to back up files that number in millions, 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 8 GB of system memory and 2.8 GB of disk space for a dual-processor, x64 Bit system.

• **OB2NDMPCATQUESIZE**

Default: 5

This option 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 option to higher values if the number of files in the backup specification is less than 20 million and enough system memory is available.

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

• **OB2NDMPFHFILEOPT**

Defaults:

Windows systems: *Data_Protector_program_data\tmp, 32, 1024*

UNIX systems: */var/opt/omni/tmp, 32, 1024*

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

- a. Pathname of the directory where the file history swap files are stored.
- b. Maximum number of file history swap files, created by Data Protector on the NDMP client's disk.
- c. 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 systems: *C:\tmp, 32, 1024; D:\tmp\tmp_1, 10, 1024*

UNIX systems: */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:

$$EstConsumption = (NumOfFiles + NumOfDirs) \times (136 + AverageFileNameSize)$$

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

See the calculations in [Approximate disk consumption by file history swap files below](#) 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 13: Approximate disk consumption by file history swap files

Number of backed up files and directories	Approximate disk consumption by file history swap files
5 million	0.7 GB
10 million	1.4 GB
20 million	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 NDMP-Celerra backup sessions.

For more information, see the *HP Data Protector 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 *HP Data Protector Help* index: "patches".
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

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 NDMP 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 `OB2SCTLMOVETIMEOUTomnirc` option on the NDMP client to a higher value (for example, 360000 or higher).

For details of how to set the `omnirc` options, see the *HP Data Protector 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 51](#).

Appendix A: Data Protector NetApp SnapManager solution

Introduction

This appendix describes the Data Protector NetApp SnapManager solution which, when used together with standard Data Protector functionality, enables backup and restore of NetApp SnapManager snapshots to and from Data Protector backup media.

The solution is available only on Windows systems.

Concepts

The NetApp SnapManager for Microsoft Exchange (SME) and NetApp SnapManager for Microsoft SQL Server (SMSQL) are solutions that create snapshots of the Microsoft Exchange Server and Microsoft SQL Server data on the NetApp storage system.

Data Protector supports SME and SMSQL through the `omnisnapmgr.pl` script, enabling you to archive existing NetApp SnapManager snapshots to Data Protector backup media. The `omnisnapmgr.pl` script uses the NetApp SnapDrive command line interface to perform queries, mount, and dismount volumes.

To back up SME and SMSQL snapshots to Data Protector backup media, you must create a standard Data Protector filesystem backup specification and specify the `omnisnapmgr.pl` script as a pre- and post-exec script to this backup specification.

At the start of the backup session (in the pre-exec phase), `omnisnapmgr` mounts the latest SME or SMSQL snapshots that were not archived yet, to the Windows client on which it runs. The Data Protector Disk Agent will then perform the backup of the files from the mounted volumes to the Data Protector backup device(s). At the end of the backup session (in the post-exec phase), `omnisnapmgr.pl` will dismount volumes that were mounted at the start of the backup session.

Configuration

Prerequisites

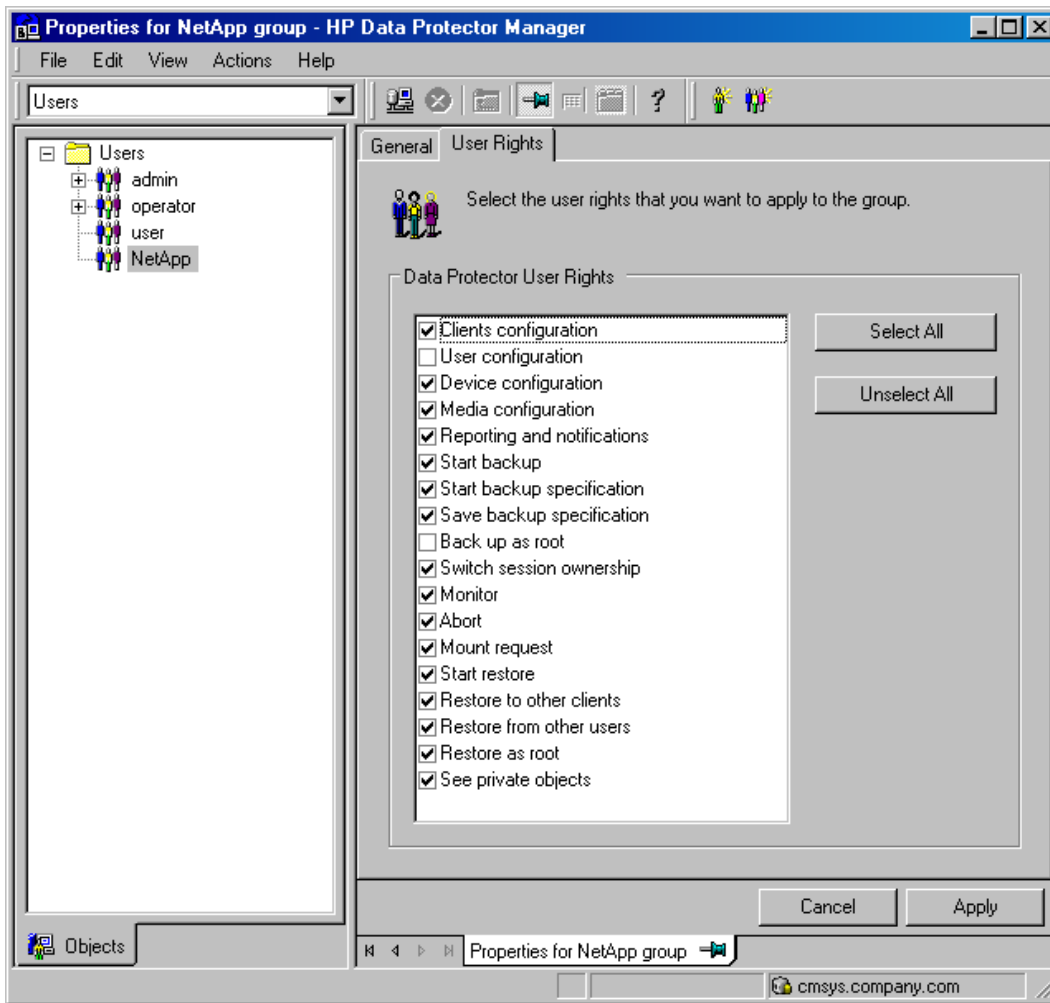
- The system on which you will install the Data Protector SnapManager solution (the backup system) must have the Data Protector Disk Agent installed and must be a member of a Data Protector cell.
- At least one Data Protector backup device must be configured in the Data Protector cell.
- NetApp SnapDrive must be installed and configured on the backup system.

Configuration procedure

If the Data Protector Inet on the backup system does *not* run under the SnapDrive account, you can:

- Change the Data Protector's Inet account to that of SnapDrive. See the *HP Data Protector Help* index: "Inet, changing account".
- Configure the SnapDrive account for use with Data Protector:
 - Create a new Data Protector user group with sufficient rights for the backup and remove the right Backup as root:

Figure 36: Setting the user rights for the new user group



- Add the SnapDrive account to this group or, if the SnapDrive account is already in the Data Protector user list, move it to the newly created group.

Backup

Limitations

- You must back up SME and SMSQL snapshots in separate backup specifications.

Creating a backup specification

To create a NetApp SnapManager backup specification:

1. Create a filesystem backup specification for the Windows client to which Data Protector will mount the SnapManager snapshots. See the *HP Data Protector Help* index: “creating, backup specifications”.

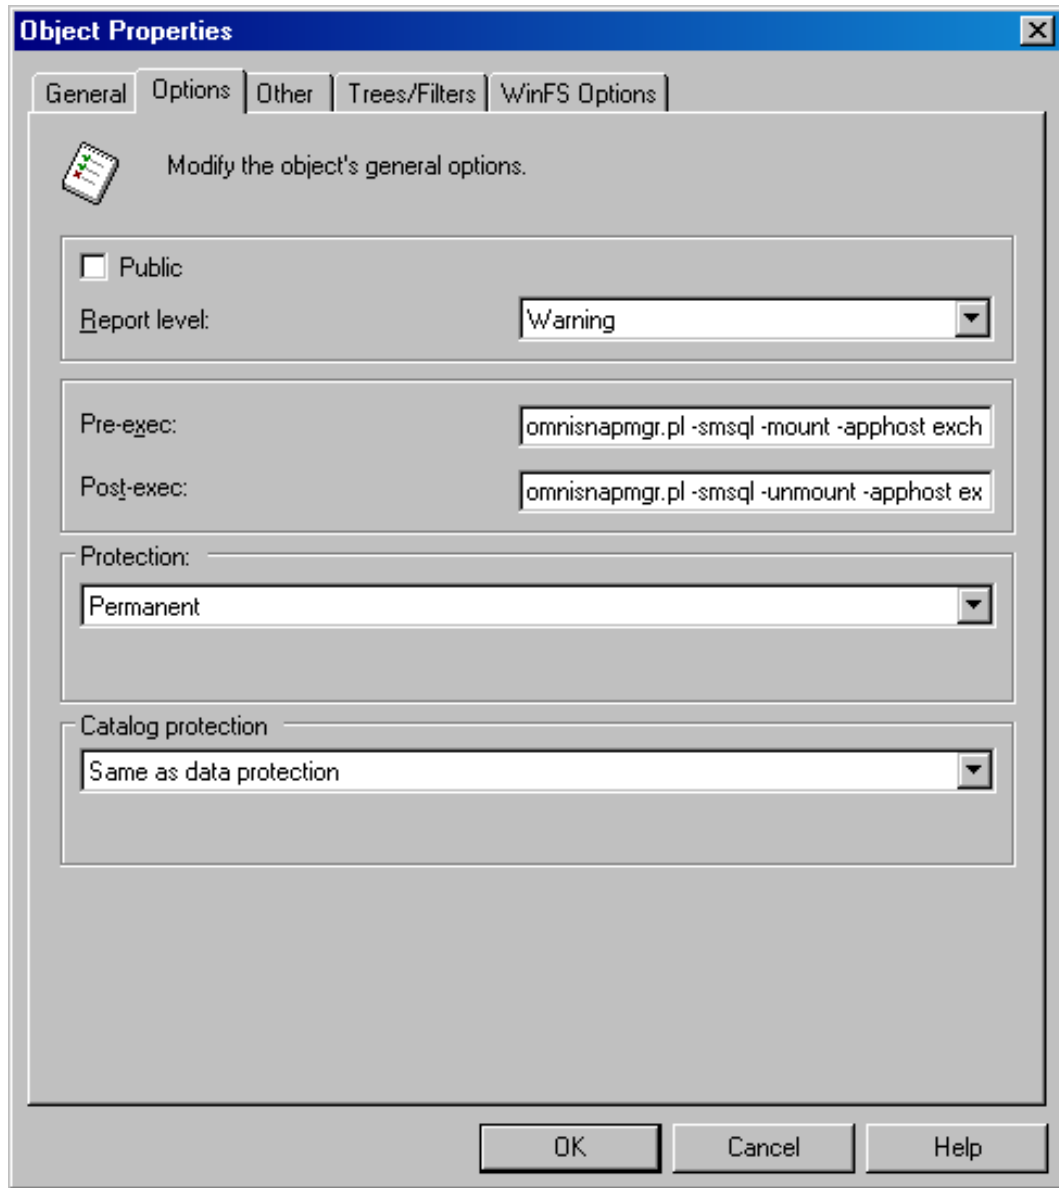
In the Source Property page, select the folder to which the Data Protector `omnisnapmgr` script will mount the volumes. This folder and all volumes mounted under it will be backed up.

Note: Data Protector excludes some temporary folders from being backed up, even if they are selected in the backup specification. You must select a folder that is not an operating system temporary folder or a Data Protector temporary folder or its sub folder, such as `C:\Windows\Temp` or `Data_Protector_home\tmp`. For a list of excluded folders, see the *HP Data Protector Help* index: “Windows, systems backup”.

2. Specify the `omnisnapmgr.pl` pre- and post-exec scripts:
 - a. Under **Filesystem** options, click **Advanced**.
 - b. In the Options pane, enter the **Pre-exec** and **Post-exec** scripts.

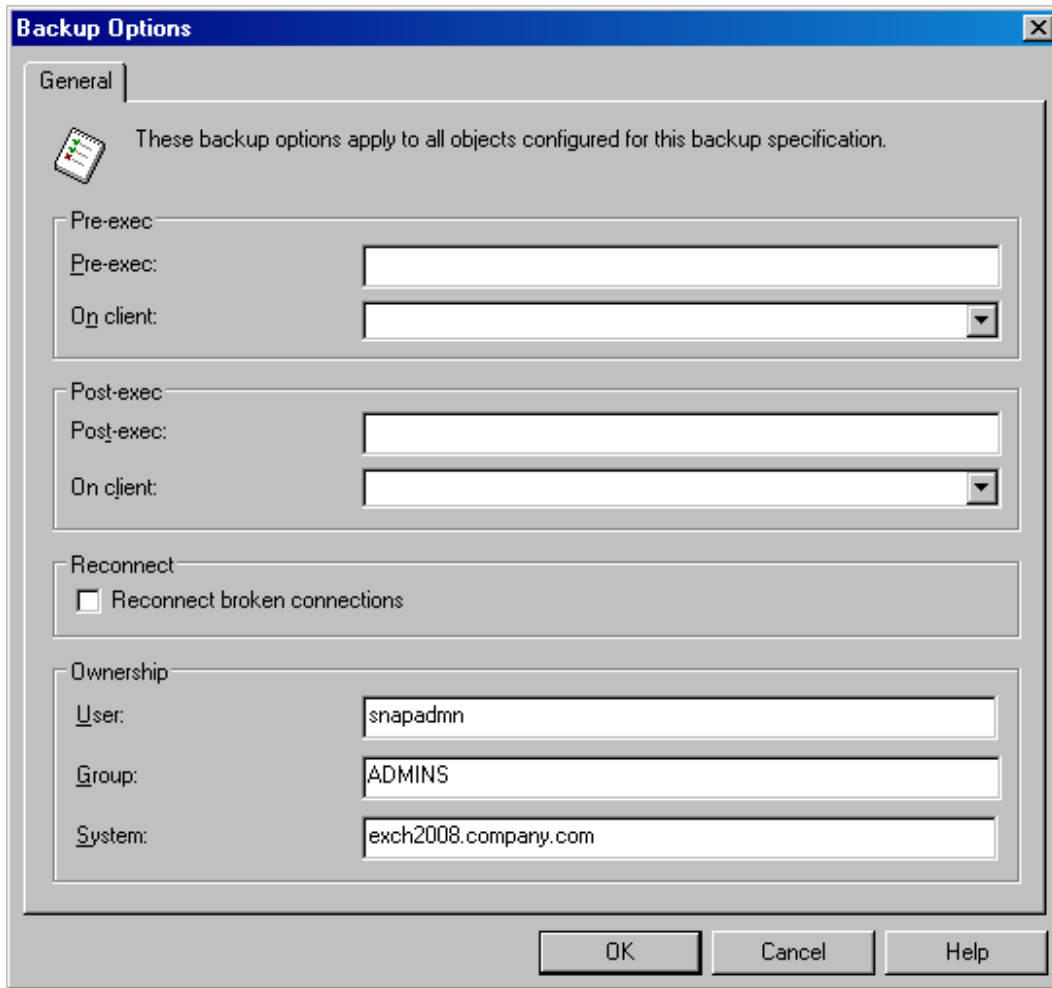
For the `omnisnapmgr.pl` syntax, available options, and examples, see the [omnisnapmgr.pl reference page on page 76](#).

Figure 37: Specifying the pre- and post-exec commands



3. If the Data Protector Inet account is not running under the SnapDrive account, specify the SnapDrive account as the backup owner:
 - a. Under **Filesystem** options, click **Advanced**.
 - b. In the Backup options window, under Ownership, enter the user name, group, and the client system name.

Figure 38: Specifying the SnapDrive account



4. Save the backup specification and run or schedule the backup session.

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

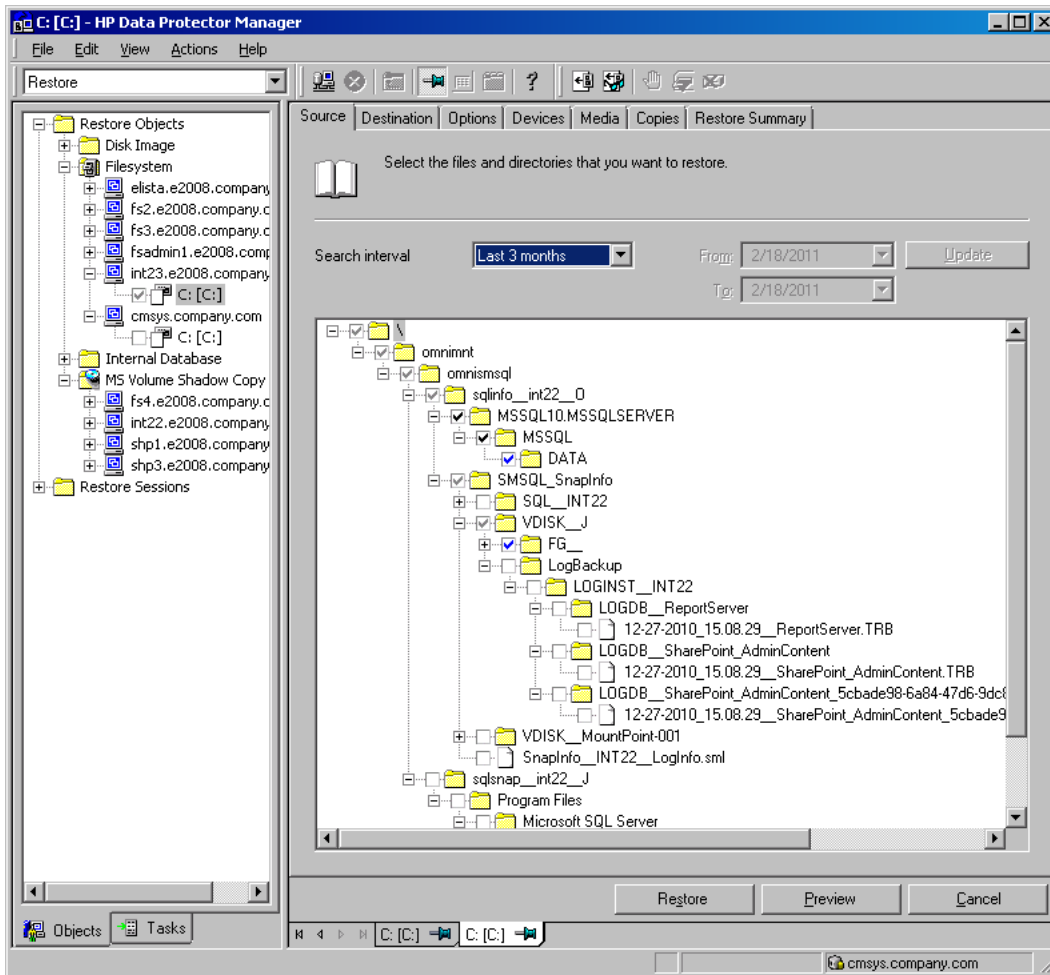
Restore

To restore SME or SMSQL data, use Data Protector and SnapManager:

1. Restore the database files and log files from Data Protector backup media to a temporary directory using the standard Data Protector restore procedure. See the *HP Data Protector Help* index: "restoring".

For an example of the backed up SMSQL objects, see [Selecting the backed up the NetApp SnapManager objects for restore](#) .

Figure 39: Selecting the backed up the NetApp SnapManager objects for restore



Note: In case of a disaster, when your application installation is lost, you need to restore or reinstall the application to the original location first and then use the SnapManager Recovery Wizard to recover the application data.

2. Use the SnapManager Recovery Wizard to recover the Microsoft Exchange Server or Microsoft SQL Server data. Select **Restore from unmanaged media** and follow the instructions.

For information, see the SnapManager documentation.

Figure 40: Selecting the restore mode in SnapManager Restore Wizard

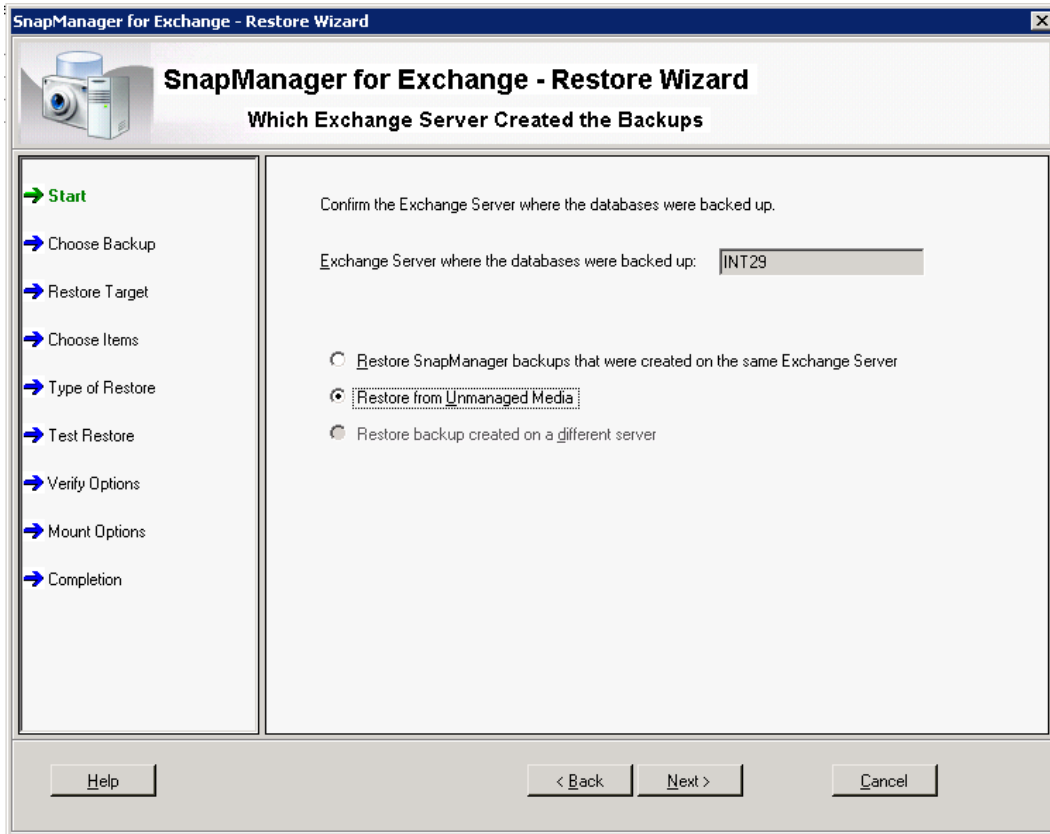
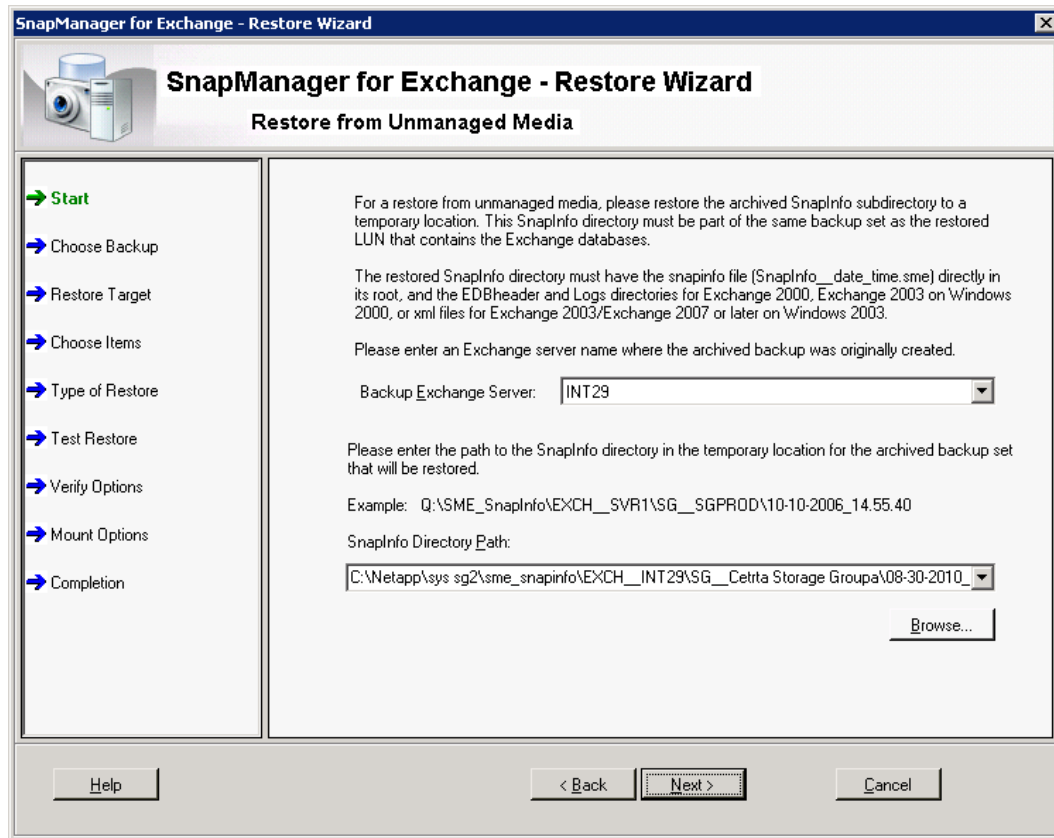


Figure 41: Specifying SnapManager Restore Wizard options



omnisnapmgr.pl reference page

SYNOPSIS

```
omnisnapmgr.pl-version | -help
```

```
omnisnapmgr.pl [-sme | -smsql] {-mount | -unmount} -apphostClientName [-force] [-partial] [-preview]
```

```
omnisnapmgr.pl-query
```

DESCRIPTION

The `omnisnapmgr.pl` script is used to mount or dismount the SME or SMSQL snapshots (when started as a pre- or post-exec script) or query NetApp snapshots (when run from a Command Prompt window).

OPTIONS

-version	Displays the Data Protector version.
----------	--------------------------------------

-help	Displays the usage synopsis.
-sme	Specifies that SME snapshots will be backed up. If not specified, <code>omnisnapmgr</code> assumes by default that SME is backed up.
-smsql	Specifies that SMSQL snapshot will be backed up.
-mount	Mounts all volumes from the last SnapManager backup of the application system.
-unmount	Dismounts all volumes mounted by the <code>-mount</code> option
- apphost <i>ClientName</i>	Specifies the application server (Exchange or SQL Server) system. If not specified, the local system is used.
-force	Performs a backup of the snapshots even if they were already backed up.
-partial	Performs a backup the SME or SMSQL snapshot even if some of the volumes cannot be mounted.
-preview	Displays the SnapDrive mount commands without executing them.
-query	Lists the snapshots and volumes contained in the last SME/SMSQL backup session.

NOTES

The `omnisnapmgr.pl` script is available on Windows systems only.

EXAMPLES

1. To back up to a SnapManager Microsoft Exchange Server snapshot to Data Protector backup media, where Microsoft Exchange Server is running on the client "exch1.company.com", specify the following pre- and post-exec commands:

Pre-exec:

```
perl omnisnapmgr.pl -sme -mount -apphost exch1.company.com
```

Post-exec:

```
perl omnisnapmgr.pl -sme -unmount -apphost exch1.company.com
```

2. To back up a SnapManager Microsoft SQL snapshot Data Protector backup media, where Microsoft SQL Server is running on the client "sql2.company.com", even if the snapshots were backed up and to ensure that the backup is performed even if some volumes cannot be mounted, specify the following pre- and post-exec commands:

Pre-exec:

```
perl omnisnapmgr.pl -smsql -mount -apphost sql2.company.com
```

Post-exec:

```
perl omnisnapmgr.pl -ssql -unmount -apphost sql2.company.com \  
-force -partial
```

Glossary

A

access rights

See user rights.

ACSL (StorageTek specific term)

The Automated Cartridge System Library Server (ACSL) software that manages the Automated Cartridge System (ACS).

Active Directory (Windows specific term)

The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access, and manage resources regardless of the physical system they reside on.

AES 256-bit encryption

The Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

AML (ADIC/GRAU specific term)

Automated Mixed-Media library.

AMU (ADIC/GRAU specific term)

Archive Management Unit.

application agent

A component needed on a client system to back up or restore online database integrations. See also Disk Agent.

application system (ZDB specific term)

A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume.

archive logging (Lotus Domino Server specific term)

Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

archived log files (Data Protector specific term)

Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online or offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.

archived redo log (Oracle specific term)

Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode. NOARCHIVELOG - The filled online redo log files are not archived. See also online redo log.

ASR set

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of

the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory <Data_Protector_program_data>\Config\Server\dr\asr (Windows systems) or /etc/opt/omni/server/dr/asr (UNIX systems), as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs

Data files to which auditing information is stored.

audit report

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

auditing information

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

autochanger

See library.

autoloader

See library.

Automatic Storage Management (ASM) (Oracle specific term)

A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

auxiliary disk

A bootable disk that has a minimal operating system with networking and

Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

B

BACKINT (SAP R/3 specific term)

A Data Protector interface program that lets the SAP R/3 backup programs communicate with the Data Protector software via calls to an open interface. For backup and restore, SAP R/3 programs issue commands through the Data Protector backint interface.

backup API (Oracle specific term)

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 reading and writing of data to media and also to allow creation, searching, and removing of backup files.

backup chain

See restore chain.

backup device

A device configured for use with Data Protector that 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 systems and mount point on UNIX systems), for integration objects - backup stream identification, indicating the backed up database/application items), description (for filesystem objects - uniquely defines objects with identical client name and mount point, for integration objects - displays the integration type), and type (for filesystem objects - filesystem type, 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. 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. All clients configured in one backup specification are backed up at the same time in one backup session using the same backup type (full or incremental).

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 view

Data Protector provides different views of your 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 Continuances are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

BC Process (EMC Symmetrix specific term)

A protected storage environment solution, which uses specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. See also BCV.

BCV (EMC Symmetrix specific term)

Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as 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 Data Protector Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

A volume/disk/partition containing 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 BRARCHIVE and BRBACKUP.

BSM

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

C

CAP (StorageTek specific term)

The Cartridge Access Port built into the door panel of a library. The purpose is to enter or eject media.

Catalog Database (CDB)

A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.

catalog protection

Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.

CDB

See Catalog Database (CDB).

CDF file (UNIX systems specific term)

A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell

A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager

The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing

Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs. See also MoM.

Centralized Media Management Database (CMMDB)

See CMMDB.

Certificate Server

A Windows certificate server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal (Windows specific term)

A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider

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' or type 'sbt_tape'. If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging (Microsoft Exchange Server and Lotus Domino Server specific term)

Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup

A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification. If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up. If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the

backup specification was created are not backed up.

client or client system

Any system configured with any Data Protector functionality and configured in a cell.

cluster continuous replication (Microsoft Exchange Server specific term)

Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node. A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group. See also Exchange Replication Service and local continuous replication.

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on HP Serviceguard), application services, IP names and addresses ...).

CMD script for Informix Server (Informix Server specific term)

A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script

is a set of system commands that export environment variables for Informix Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended. See also MoM.

COM+ Class Registration Database (Windows specific term)

The COM+ Class Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guaranties consistency among these attributes and provides common operation on top of these attributes.

command device (HP P9000 XP Disk Array Family specific term)

A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

command-line interface (CLI)

A set commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency

See Disk Agent concurrency.

container (HP P6000 EVA Disk Array Family specific term)

Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file (Oracle and SAP R/3 specific term)

A data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set (HP P6000 EVA Disk Array Family specific term)

A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA. See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

CRS

The Data Protector Cell Request Server process (service), which runs on the Cell Manager, starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. The CRS runs under the account root on UNIX systems. On Windows systems it runs under the account of the user, specified at installation time.

CSM

The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D**data file (Oracle and SAP R/3 specific term)**

A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection

Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions. See also catalog protection.

Data Protector user account

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data replication (DR) group (HP P6000 EVA Disk Array Family specific term)

A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log. See also copy set.

data stream

Sequence of data transferred over the communication channel.

Data_Protector_home

A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is %ProgramFiles%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_program_data.

Data_Protector_program_data

A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012. Its default path is %ProgramData%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_home.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

database parallelism

More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject (Informix Server specific term)

An Informix Server physical database object. It can be a blob space, db space, or logical log file.

DC directory

A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. See also Detail Catalog Binary Files (DCBF) and Internal Database (IDB).

DCBF

See Detail Catalog Binary Files (DCBF).

delta backup

A delta backup is a backup containing all the changes made to the database from the last backup of any type. See also backup types.

Detail Catalog Binary Files (DCBF)

A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. See also DC directory and Internal Database (IDB).

device

See backup device.

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. -Get configuration, status, and performance statistics by device group. -Issue control operations that apply to all devices in the device group.

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 writing to the tape, and so on. If data is written to the tape slower than it is delivered to the device then the device is streaming. Streaming significantly improves the use of space and the performance of the device.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup

An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. See also incremental backup.

differential backup (Microsoft SQL Server specific term)

A database backup that records only the data changes made to the database after the last full database backup. See also backup types.

differential database backup

A differential database backup records only those data changes made to the database after the last full database backup.

directory junction (Windows specific term)

Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery

A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

disaster recovery operating system

See DR OS.

Disk Agent

A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

Disk Agent concurrency

The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk group (Veritas Volume Manager specific term)

The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (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.

DR image

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS

An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium. It can also read the data from the medium and send it to the computer system.

drive index

A number that identifies the mechanical position of a drive inside a library device.

This number is used by the robotic control to access a drive.

drive-based encryption

The Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the metadata that is written to the medium.

E

EMC Symmetrix Agent

A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file (Informix Server specific term)

The Informix Server configuration file `ixbar.<server_id>` that resides in the directory `<INFORMIXDIR>/etc` (on Windows systems) or `<INFORMIXDIR>\etc` (on UNIX systems). `<INFORMIXDIR>` is the Informix Server home directory and `<server_id>` is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

encrypted control communication

Data Protector secure communication between the clients in the Data Protector cell is based on a Secure Socket Layer (SSL) that uses export grade SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key

A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.

encryption KeyID-StoreID

Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. KeyID identifies the key within the keystore. StoreID identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may be several StoreIDs used on the same Cell Manager.

enhanced incremental backup

Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

enterprise backup environment

Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept. See also MoM.

Event Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The 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 (Windows specific term)

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.

exchanger

See library.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and filenames, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. See also importing media.

**Extensible Storage Engine (ESE)
(Microsoft Exchange Server specific
term)**

A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F

failover

Transferring of the most important cluster data, called group (on Windows systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

**failover (HP P6000 EVA Disk Array
Family specific term)**

An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations. See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

FC bridge

See Fibre Channel bridge.

Fibre Channel

An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge

A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries

to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot

A file containing the data from a backup to a file library device.

file jukebox device

A device residing on disk consisting of multiple slots used to store file media.

file library device

A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)

A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file tree walk

The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror (HP P9000 XP Disk Array Family specific term)

A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used. See also primary volume and mirror unit (MU) number.

flash recovery area (Oracle specific term)

A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files). See also recovery files.

formatting

A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Use the Force operation option to reformat Data Protector media with non-protected data. Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup

A backup in which all selected objects are backed up, whether or not they have been recently modified. See also backup types.

full database backup

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup

A full mailbox backup is a backup of the entire mailbox content.

full ZDB

A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup. See also incremental ZDB.

G

global options

A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager.

group (Microsoft Cluster Server specific term)

A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI

A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.

H**hard recovery (Microsoft Exchange Server specific term)**

A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat

A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)

A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file

A file that contains information about holidays. You can set different holidays by editing the Holidays file that resides on the Cell Manager at the following location: <Data_Protector_program_data>\Config\Server\holidays (Windows systems) and /etc/opt/omni/server/Holidays (UNIX systems).

hosting system

A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP Business Copy (BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware. See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

HP Business Copy (BC) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P4000 SAN Solutions configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system. See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit (MCU), application system, and backup system.

HP Command View (CV) EVA (HP P6000 EVA Disk Array Family specific term)

The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, or mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed

by a Web browser. See also HP P6000 / HP 3PAR SMI-S Agent and HP SMI-S P6000 EVA Array provider.

HP Continuous Access (CA) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs). See also HP BC P9000 XP (HP P9000 XP Disk Array Family specific term), Main Control Unit (MCU), and LDEV.

HP Continuous Access + Business Copy (CA+BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array. See also HP BC P6000 EVA, replica, and source volume.

HP P6000 / HP 3PAR SMI-S Agent

A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 / 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface. See

also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP P9000 XP Agent

A Data Protector software component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It communicates with the P9000 XP Array storage system via the RAID Manager Library.

HP SMI-S P6000 EVA Array provider

An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses. See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

I

ICDA (EMC Symmetrix specific term)

EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

See Internal Database (IDB).

IDB recovery file

A file that maintains information about completed IDB backup sessions and the

backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

importing media

A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media. See also exporting media.

incremental (re-)establish (EMC Symmetrix specific term)

A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length. See also backup types.

incremental backup (Microsoft Exchange Server specific term)

A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up. See also backup types.

incremental mailbox backup

An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental restore (EMC Symmetrix specific term)

A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB

A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape. See also full ZDB.

Incremental1 Mailbox Backup

An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet

A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication

between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. See also Key Management Service and Site Replication Service.

Informix Server (Informix Server specific term)

Refers to Informix Dynamic Server.

initializing

See formatting.

Installation Server

A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (ZDB specific term)

A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore

from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery. See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

integration object

A backup object of a Data Protector integration, such as Oracle or SAP MaxDB.

Internal Database (IDB)

An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It is implemented with an embedded database and a collection of proprietary data files which reside on the Cell Manager. See also DC directory and Detail Catalog Binary Files (DCBF).

Internet Information Server (IIS) (Windows specific term)

Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

ISQL (Sybase specific term)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

jukebox

See library.

jukebox device

A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the 'file jukebox device'.

K

Key Management Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service.

keychain

A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

keystore

All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS

Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

L

LBO (Symmetric specific term)

A Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as

one entity and can only be restored as a whole.

LDEV (HP P9000 XP Disk Array Family specific term)

A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library

Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or unattended operation

A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA (Oracle specific term)

An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing

By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during

backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication (Microsoft Exchange Server specific term)

Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and

can be backed up (using VSS) as a normal storage group. See also cluster continuous replication and Exchange Replication Service.

lock name

You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script (Informix Server UNIX systems specific term)

A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to `<INFORMIXDIR>/etc/no_log.sh`.

logging level

An option that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID (Microsoft SQL Server specific term)

The name a user needs to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database (Oracle and SAP R/3 specific term)

The format of the login information is <user_name>/<password>@<service>, where: <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights. <password> must be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration. <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database (Oracle specific term)

The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database.

In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API (Lotus Domino Server specific term)

An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

M

Magic Packet

See Wake ONLAN.

mailbox (Microsoft Exchange Server specific term)

The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store (Microsoft Exchange Server specific term)

A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP Array or HP CA+BC P9000 XP Array configuration and acts as a master device. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

maintenance mode

An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the installation.

make_net_recovery

make_net_recovery is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX make_boot_tape command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX bootsys command or interactively specified on the boot console.

make_tape_recovery

make_tape_recovery is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

Manager-of-Managers

See MoM.

MAPI (Microsoft Exchange specific term)

The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU

See Main Control Unit (MCU).

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy

Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition

The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of

read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors

The user-assigned age threshold and overwrite threshold used to assert the condition of a medium.

media label

A user-defined identifier used to describe a medium.

media location

A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session

A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool

A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set

The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type

The physical type of media, such as DDS or DLT.

media usage policy

The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

medium ID

A unique identifier assigned to a medium by Data Protector.

merging

This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. See also overwrite.

Microsoft Exchange Server

A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) (Windows specific term)

An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server

A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy Service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow

copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets. See also shadow copy, shadow copy provider, replica, and writer.

mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

See target volume.

mirror rotation (HP P9000 XP Disk Array Family specific term)

See replica set rotation.

mirror unit (MU) number (HP P9000 XP Disk Array Family specific term)

A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family. See also first-level mirror.

mirrorclone (HP P6000 EVA Disk Array Family specific term)

A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD

The Media Management Daemon process (service) (MMD) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots

configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells. See also CMMDB and Catalog Database (CDB).

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount point

The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX systems, the mount points are displayed using the bdf or df command.

mount request

A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

MSM

The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping (HP P6000 EVA Disk Array Family specific term)

Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot. See also snapshot.



OBDR capable device

A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

obdrindex.dat

See IDB recovery file.

object

See backup object.

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session

A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

object copy

A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session

A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying

The process of copying selected object versions to a specific media set. You can

select object versions from one or several backup sessions to be copied.

object ID (Windows specific term)

The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror

A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring

The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

object verification

The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session

A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

offline backup

A backup during which an application database cannot be used by the

application. For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished. For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process. See also zero downtime backup (ZDB) and online backup.

offline recovery

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Cell Manager can only be recovered offline.

offline redo log

See archived redo log.

ON-Bar (Informix Server specific term)

A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: the onbar command, Data Protector as the backup solution, the XBSA interface, and ON-Bar catalog tables, which are used to back up dobjects and track instances of dobjects 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 located in

the directory <INFORMIXDIR>\etc (on Windows systems) or <INFORMIXDIR>/etc/ (on UNIX systems).

online backup

A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to tape is finished. For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process. In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. See also zero downtime backup (ZDB) and offline backup.

online recovery

A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.

online redo log (Oracle specific term)

Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused. See also archived redo log.

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.

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. When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

P1S file

P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into the

directory <Data_Protector_program_data>\Config\Server\dr\p1s (Windows systems) or /etc/opt/omni/dr/p1s (UNIX systems) with the filename recovery.p1s.

package (HP ServiceGuard and Veritas Cluster Specific Term)

A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

pair status (HP P9000 XP Disk Array Family specific term)

The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP P9000 XP Agent:

PAIR - The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.

SUSPENDED - The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.

COPY - The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

parallel restore

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical

volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism

The concept of reading multiple data streams from an online database.

phase 0 of disaster recovery

Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery

Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery

Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.

phase 3 of disaster recovery

Restoration of user and application data.

physical device

A physical unit that contains either a drive or a more complex unit such as a library.

post-exec

A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also pre-exec.

pre- and post-exec commands

Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems.

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 systems and as shell scripts on UNIX systems. See also post-exec.

primary volume (P-VOL) (HP P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU). See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and catalog protection.

public folder store (Microsoft Exchange Server specific term)

The part of the Information Store that maintains information in public folders. A

public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be: public, that is visible (and accessible for restore) to all Data Protector users private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID

Redundant Array of Independent Disks.

RAID Manager Library (HP P9000 XP Disk Array Family specific term)

A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands. See also HP P9000 XP Agent.

RAID Manager P9000 XP (HP P9000 XP Disk Array Family specific term)

A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.

rawdisk backup

See disk image backup.

RCU

See Remote Control Unit (RCU).

RDBMS

Relational Database Management System.

RDF1/RDF2 (EMC Symmetrix specific term)

A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

Recovery Catalog (Oracle specific term)

A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about the physical schema of the Oracle target database, data file and archived log backup sets, data file copies, archived redo logs, and stored scripts.

Recovery Catalog Database (Oracle specific term)

An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files (Oracle specific term)

Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area.

Recovery Manager (RMAN) (Oracle specific term)

An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the

recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

RecoveryInfo

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

recycle or unprotect

A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log (Oracle specific term)

Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit (RCU) (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.

Removable Storage Management Database (Windows specific term)

A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications

to access and share the same media resources.

reparse point (Windows specific term)

A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica (ZDB specific term)

An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on a UNIX system, the whole volume/disk group containing a backup object is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition is replicated. See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set (ZDB specific term)

A group of replicas, all created using the same backup specification. See also replica and replica set rotation.

replica set rotation (ZDB specific term)

The use of a replica set for regular backup production: Each time the same backup

specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set. See also replica and replica set.

restore chain

Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.

restore session

A process that copies data from backup media to a client.

resync mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL. See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

RMAN (Oracle specific term)

See Recovery Manager.

RSM

The Data Protector Restore Session Manager controls restore and object

verification sessions. This process always runs on the Cell Manager system.

RSM (Windows specific term)

Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

S

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler

A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

secondary volume (S-VOL) (HP P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration. See also primary volume (P-VOL) and Main Control Unit (MCU).

session

See backup session, media management session, and restore session.

session ID

An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.

session key

This environment variable for the pre- and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat, and omniabort CLI commands.

shadow copy (Microsoft VSS specific term)

A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant. See also Microsoft Volume Shadow Copy Service and replica.

shadow copy provider (Microsoft VSS specific term)

An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays). See also shadow copy.

shadow copy set (Microsoft VSS specific term)

A collection of shadow copies created at the same point in time. See also shadow copy and replica set.

shared disks

A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

Site Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service. See also Information Store and Key Management Service.

slot

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB

See split mirror backup.

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA)

See HP P6000 / HP 3PAR SMI-S Agent.

snapshot (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A type of target volumes created using a specific replication technology.

Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume. See also replica and snapshot creation.

snapshot backup

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation. See also snapshot.

source (R1) device (EMC Symmetrix specific term)

An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type. See also target (R2) device.

source volume (ZDB specific term)

A storage volume containing data to be replicated.

sparse file

A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -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 P9000 XP Disk Array Family specific term)

A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes. See also replica and split mirror creation.

split mirror backup (EMC Symmetrix specific term)

See ZDB to tape.

split mirror backup (HP P9000 XP Disk Array Family specific term)

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes. See also split mirror.

split mirror restore (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method. See also ZDB to tape, ZDB to disk+tape, and replica.

sqlhosts file or registry (Informix Server specific term)

An Informix Server connectivity information file (on UNIX systems) or registry (on Windows systems) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file

A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster. See also target system.

SRDF (EMC Symmetrix specific term)

The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent (SSEA)

See HP P9000 XP Agent.

sst.conf file

The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file

The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device

A file device is a file in a specified directory to which you back up data.

Storage Group (Microsoft Exchange Server specific term)

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

storage volume (ZDB specific term)

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.

Typically, these can be created or exist within a storage system such as a disk array.

StorageTek ACS library (StorageTek specific term)

Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

switchover

See failover.

Sybase Backup Server API (Sybase specific term)

An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server (Sybase specific term)

The server in the Sybase "client-server" architecture. The Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

SYMA

See EMC Symmetrix Agent.

synthetic backup

A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup

The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

System Backup to Tape (SBT) (Oracle specific term)

An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases (Sybase specific term)

The four system databases on a newly installed Sybase SQL Server are the: - master database (master) -temporary database (tempdb) -system procedure database (sybssystemprocs) -model database (model).

System Recovery Data file

See SRD file.

System State (Windows specific term)

The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SysVol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition

A volume/disk/partition containing operating system files. Microsoft

terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (Windows specific term)

A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

T

tablespace

A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup (ZDB specific term)

See ZDB to disk.

target (R2) device (EMC Symmetrix specific term)

An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. See also source (R1) device.

target database (Oracle specific term)

In RMAN, the target database is the database that you are backing up or restoring.

target system (disaster recovery specific term)

A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.

target volume (ZDB specific term)

A storage volume to which data is replicated.

Terminal Services (Windows specific term)

Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread (Microsoft SQL Server specific term)

An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder (EMC Symmetrix specific term)

A business continuation process that creates an instant copy of single or multiple EMC 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

See Tape Library Unit.

TNSNAMES.ORA (Oracle and SAP R/3 specific term)

A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction

A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes

transaction backup

Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup (Sybase and SQL specific term)

A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup

Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files

Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction log table (Sybase specific term)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (Microsoft VSS specific term)

A shadow copy that is created on the application system and can be presented to the backup system which performs the backup. See also Microsoft Volume Shadow Copy Service (VSS).

U

unattended operation

See lights-out operation.

user account (Data Protector user account)

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

User Account Control (UAC)

A security component in Windows Vista and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas

NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group

Each Data Protector user is member of a User Group. Each User Group has a set

of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile (Windows specific term)

Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights

User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

user_restrictions file

A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than admin and operator.

V

vaulting media

The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

verify

A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS) (HP P6000 EVA Disk Array Family specific term)

The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. See also HP Command View (CV) EVA.

Virtual Device Interface (Microsoft SQL Server specific term)

This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk (HP P6000 EVA Disk Array Family specific term)

A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. See also source volume and target volume.

virtual full backup

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

Virtual Library System (VLS)

A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

virtual server

A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape

An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. See also Virtual Library System (VLS) and virtual tape library (VTL).

virtual tape library (VTL)

Data storage virtualization software that is able to emulate tape devices and libraries thus providing the functionality of traditional tape-based storage. See also Virtual Library System (VLS).

volser

A VOLume SERIAL number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group

A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (Windows specific term)

An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to

the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy Service

See Microsoft Volume Shadow Copy Service (VSS).

VSS

See Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching 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.

W

Wake ONLAN

Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting

The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character

A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows configuration backup

Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry

A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server

A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer (Microsoft VSS specific term)

A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

X

XBSA Interface (Informix Server specific term)

ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z

ZDB

See zero downtime backup.

ZDB 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 with specific disk array families, split mirror restore. See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape (ZDB specific term)

A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used. See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation. See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index

A

architecture

NDMP integration 38

Sybase integration 8

B

backing up

NetApp SnapManager data 70

backing up NDMP 52

backup specification, creating 53

backup specification, modifying 57

backup types 37

starting backups 57

backing up Sybase 14

backup options 20

backup specifications, creating 15

backup specifications, modifying 20

backup types 7

database objects backup 14

full backups 7, 14

previewing backups 21

scheduling backups 20

scheduling backups, example 20

starting backups 22

transaction logs backups 7, 14

backup options

Sybase integration 20

backup sessions, scheduling

Sybase integration 20

backup specifications, creating

NDMP integration 53

Sybase integration 15

backup specifications, modifying

NDMP integration 57

Sybase integration 20

backup types

NDMP integration 37

Sybase integration 7

block size

NDMP integration 51

C

Celerra NAS devices

NDMP integration 39, 49, 52, 61

checking configuration

Sybase integration 13

command-line interface reference

NetApp SnapManager solution 76

concepts

NDMP integration 38

NetApp SnapManager solution 69

Sybase integration 7

configuration

NetApp SnapManager solution 69

configuring NDMP 40

configuring NDMP devices 42

creating media pools 42

importing NDMP Servers 40

configuring Sybase 9

checking configuration 13

creating backup specifications

NDMP integration 53

Sybase integration 15

E

environment variables

NDMP integration 61

examples, Sybase integration

restore 32

scheduling backups 20

F

file history swap files

NDMP integration 39

full backups

Sybase integration 7, 14

H

Hitachi BlueArc NAS devices

NDMP integration 39, 50, 60

Hitachi NAS devices

NDMP integration 39, 50, 60

HP-X9000 NAS device

NDMP integration 39

I

interactive backups

NDMP integration 57

Sybase integration 22

introduction

NDMP integration 37

Sybase integration 7

M

media management

NDMP integration 66

modifying backup specifications

NDMP integration 57

Sybase integration 20

monitoring sessions

Sybase integration 34

N

NDMP backup 52

backup specification, creating 53

backup specification, modifying 57

backup types 37

starting backups 57

NDMP configuration 40

configuring NDMP devices 42

creating media pools 42

importing NDMP Servers 40

NDMP integration

architecture 38

backup 52

concepts 38

configuration 40

environment variables 61

file history swap files 39

introduction 37

media management 66

omnirc options 64

restore 57

troubleshooting 67

NDMP restore 57

direct access restore 60

using another device 61

using GUI 58

- NDMP troubleshooting 67
- NetApp NAS devices
 - NDMP integration 39, 48, 51, 61
- NetApp SnapManager solution
 - backing up data 70
 - command-line interface reference 76
 - concepts 69
 - configuration 69
 - restoring data 73
- O**
- omnirc options
 - NDMP integration 64
- P**
- previewing backups
 - Sybase integration 21
- R**
- restoring
 - NetApp SnapManager data 73
- restoring NDMP 57
 - direct access restore 60
 - using another device 61
 - using GUI 58
- restoring Sybase 24
 - examples 32
 - finding information for restore 24
 - using another device 33
 - using the Sybase isql command 30
- running backups 22, 57
- S**
- scheduling backups
 - Sybase integration 20
- starting backups
 - NDMP integration 57
 - Sybase integration 22
- Sybase backup 14
 - backup options 20
 - backup specifications, creating 15
 - backup specifications, modifying 20
 - backup types 7
 - database objects backup 14
 - full backups 7, 14
 - previewing backups 21
 - scheduling backups 20
 - scheduling backups, example 20
 - starting backups 22
 - transaction logs backups 7, 14
- Sybase configuration 9
 - checking configuration 13
- Sybase integration
 - architecture 8
 - backup 14
 - concepts 7
 - configuration 9
 - introduction 7
 - monitoring sessions 34
 - restore 24
 - troubleshooting 34
- Sybase restore 24
 - examples 32
 - finding information for restore 24
 - using another device 33
 - using the Sybase isql command 30

Sybase troubleshooting 34

T

transaction logs backups

 Sybase integration 7, 14

troubleshooting NDMP 67

troubleshooting Sybase 34

We appreciate your feedback!

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Integration guide for Sybase and Network Data Management Protocol Server (Data Protector 8.10)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

