

HP OpenView Storage Data Protector Integration Guide

for

Sybase Network Node Manager Network Data Management Protocol

Manual Edition: October 2004



i n v e n t

Manufacturing Part Number: B6960-90111

Release A.05.50

© Copyright Hewlett-Packard Development Company, L.P.2004.

Legal Notices

©Copyright 2004 Hewlett-Packard Development Company, L.P.

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

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

UNIX® is a registered trademark of The Open Group.

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

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

Java™ is a U.S. trademark of Sun Microsystems, Inc.

ARM® is a registered trademark of ARM Limited.

1. Integrating Sybase and Data Protector

In This Chapter	2
Overview	3
Prerequisites and Limitations	5
Integration Concepts	7
Data Protector Sybase Configuration File	9
Setting, Retrieving, and Listing Data Protector Sybase Configuration File Parameters Using the CLI	11
Configuring the Integration	14
Before You Begin Configuring	14
Configuring a Sybase User in Data Protector	17
Configuring a Sybase Server	20
Configuring a Sybase Backup	30
Testing the Integration	42
Using the Data Protector GUI	42
Using the Data Protector CLI	42
Backing Up a Sybase Database	45
Scheduling an Existing Backup Specification	47
Running an Interactive Backup	50
Backing Up Using Sybase Commands	52
Restoring a Sybase Database	54
Finding Information Needed for Restore	54
Restore	57
The Data Protector <code>syb_tool</code> Command	65
Restoring Using Another Device	72
Disaster Recovery	73
Monitoring a Sybase Backup and Restore Session	75
Monitoring Current Sessions	75
Viewing Previous Sessions	76
Sybase Character Sets	77
Troubleshooting	78
Before You Begin	78
Troubleshooting on Windows Systems	78
Troubleshooting on UNIX Systems	84

2. Integrating Network Node Manager and Data Protector

In This Chapter	98
Overview	99

Contents

Prerequisites and Limitations	101
Integration Concept	102
Configuring an NNM Backup	104
Tasks for the NNM Administrator	104
Creating a New Template	105
Creating a Backup Specification	105
NNM Backup Options	107
Testing the Integration	107
Backing Up an NNM Database	109
Scheduling a Backup	110
Starting an Interactive Backup	111
Restoring NNM	113
Disaster Recovery	113
Monitoring an NNM Backup and Restore	115
Acceptable Warnings on the Windows Systems	115
Troubleshooting	117
Error and Warning Messages	117
Backup and Restore Problems	120

3. Integrating the NDMP Server and Data Protector

In This Chapter	122
Overview	123
Prerequisites and Limitations	125
Integration Concept	127
Network Data Management Protocol (NDMP)	130
Configuring the Integration	135
Library Configurations	135
Configuration Procedure	138
Importing the NDMP Server Host	138
Creating a Media Pool	140
Configuring an NDMP Backup Device	141
Network Appliance Configuration	146
EMC Celerra Configuration	148
Backing Up the NDMP Server Data	149
Restoring the NDMP Server Data	154
Direct Access Restore	155
Restore Using Another Device	157
NDMP Environment Variables	158

Contents

The NDMP Related omnirc File Variables	160
Media Management	164
Troubleshooting	165
Error Messages	165
Catalog Data Does Not Fit	165
Importing NDMP Media	165
Use of Media on Different Types of NDMP Servers	166
Use of NDMP Dedicated Media Pools with Standard Non-NDMP Devices	166
A Tape Remains in the Drive After the Scan Operation	166

Glossary

Index

Contents

Printing History

The manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

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

Table 1

Edition History

Part Number	Manual Edition	Product
B6960-90111	October 2004	Data Protector Release A.05.50

Conventions


The following typographical conventions are used in this manual.

Table 2

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

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information about the Data Protector graphical user interface.

Figure 1 Data Protector Graphical User Interface



Contact Information

General Information

General information about Data Protector can be found at

<http://www.hp.com/go/dataprotector>

Technical Support

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

<http://support.openview.hp.com/support.jsp>

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

Information about the latest Data Protector patches can be found at

http://support.openview.hp.com/patches/patch_index.jsp

For information on the Data Protector required patches, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

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

Documentation Feedback

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

http://ovweb.external.hp.com/lpe/doc_serv/

Training Information

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

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

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

Data Protector Documentation

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

Manuals

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

HP OpenView Storage Data Protector Concepts Guide

This manual describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented *HP OpenView Storage Data Protector Administrator's Guide*.

HP OpenView Storage Data Protector Administrator's Guide

This manual describes typical configuration and administration tasks performed by a backup administrator, such as device configuration, media management, configuring a backup, and restoring data.

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Integration Guide

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

- *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*

This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft Exchange Server 5.x, Microsoft SQL Server 7/2000, and Volume Shadow Copy Service.

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

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

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server 2000/2003, and Microsoft SQL Server 2000 databases. The manual also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

HP OpenView Storage Data Protector MPE/iX System User Guide

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

HP OpenView Storage Data Protector Media Operations User's Guide

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

HP OpenView Storage Data Protector Software Release Notes

This manual gives a description of new features of HP OpenView Storage Data Protector A.05.50. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.openview.hp.com/products/datapro/spec_0001.html.

Online Help

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

In This Book

The *HP OpenView Storage Data Protector Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol* describes how to configure and use Data Protector with Sybase, Network Node Manager, and Network Data Management Protocol.

Audience

This manual is intended for backup administrators who are responsible for the planning, setup, and maintenance of network backups. It assumes that you are familiar with:

- Basic Data Protector functionality
- Database administration

Conceptual information can be found in the *HP OpenView Storage Data Protector Concepts Guide*, which is recommended to fully understand the fundamentals and the model of Data Protector.

Organization

The manual is organized as follows:

- Chapter 1** “Integrating Sybase and Data Protector” on page 1.
- Chapter 2** “Integrating Network Node Manager and Data Protector” on page 97.
- Chapter 3** “Integrating the NDMP Server and Data Protector” on page 121.
- Glossary** Definition of terms used in this manual.

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*:

- Microsoft SQL Server 7.0/2000
- Microsoft Exchange Server 5.x
- Microsoft Exchange Server 2000/2003
- Microsoft Volume Shadow Copy Service

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for Oracle and SAP*:

- Oracle
- SAP R/3
- SAP DB

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*:

- Informix
- IBM DB2 UDB
- Lotus Notes/Domino

The integrations of Data Protector ZDB integrations with the following applications or operating system services are described in the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*:

- Oracle
- SAP R/3
- Microsoft SQL Server 7.0/2000
- Microsoft Volume Shadow Copy Service
- Microsoft Exchange Server 2000/2003

1 Integrating Sybase and Data Protector

In This Chapter

This chapter explains how to configure and use the Data Protector Sybase integration. It explains the concepts and methods that you need to understand to back up and restore Sybase data.

The chapter is organized into the following sections:

“Overview” on page 3

“Prerequisites and Limitations” on page 5

“Integration Concepts” on page 7

“Data Protector Sybase Configuration File” on page 9

“Configuring the Integration” on page 14

“Testing the Integration” on page 42

“Backing Up a Sybase Database” on page 45

“Restoring a Sybase Database” on page 54

“Monitoring a Sybase Backup and Restore Session” on page 75

“Sybase Character Sets” on page 77

“Troubleshooting” on page 78

Overview

Data Protector integrates with Sybase SQL Server to offer the online backup of your Sybase databases.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for up-to-date information about platforms supported by the Data Protector Sybase integration.

The online backup concept is now widely accepted because it addresses the business requirement of high application availability. During the backup, the database is online and actively used. The backup is performed quickly and efficiently, with least impact on database performance.

Backup Types

You can perform the following types of backups of your Sybase databases from the Data Protector User Interface or from Sybase via the **isql** command:

- Interactive backup using any of the following backup modes:
 - **Full**, at which part of or the entire database, including both the data and the **transaction log**, is backed up
 - **Trans**, at which the transaction log is backed up, providing a record of any changes made since the last full or trans backup
- Scheduled backup of selected Sybase databases. Data Protector allows you to define the date and time for your unattended backup to start. You can also use predefined backup schedules to simplify your configuration.

A backup is always executed on the **Sybase Server** via the **isql** utility. The **isql** utility communicates backup and restore requests to Sybase Backup Server.

Restore Types

You can perform the following types of restores of your Sybase databases using Sybase **isql** commands:

- Restore all or part of the database
- Restore the database to a special point in time

Why Use the Data Protector User Interface?

Backing up using the integration offers various advantages over backing up using Sybase Backup Server alone:

- Central Management for all backup operations

You can manage backup operations from a central point. This is especially important in large business environments.

- Media Management

Data Protector has an advanced media management system, which allows you to keep track of all media and the status of each medium, set protection for stored data, fully automate operation, as well as organize and manage devices and media.

- Backup Management

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

- Scheduling

Data Protector has a built-in scheduler that allows you to automate backups to run periodically. With the Data Protector Scheduler, the backups you set will run unattended at the times you specify.

- Device Support

Data Protector supports a wide range of devices; from standalone drives to complex multiple drive libraries. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported devices and other information.

- Reporting

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

- Monitoring

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

All backup sessions are logged in the built-in IDB, providing you with a history of activities that can be queried at a later time.

Prerequisites and Limitations

Prerequisites

- You need a license to use the Data Protector Sybase integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for licensing instructions.
- Before you begin, ensure that you have correctly installed and configured Sybase Server and Data Protector. For additional information, refer to the:
 - *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures.
 - *HP OpenView Storage Data Protector Administrator's Guide* for instructions on how to configure and run backups.
 - *Sybase SQL Server System Administration Guide* and *Sybase SQL Server Installation and Configuration Guide* for more information on Sybase.

Audience

- The primary audience of this chapter is the administrator who must backup and restore Sybase data using the Data Protector Sybase integration. This chapter assumes that you are familiar with Sybase SQL Server, Sybase Backup Server, Windows or UNIX operating system, and basic Data Protector functionality. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for Data Protector details.

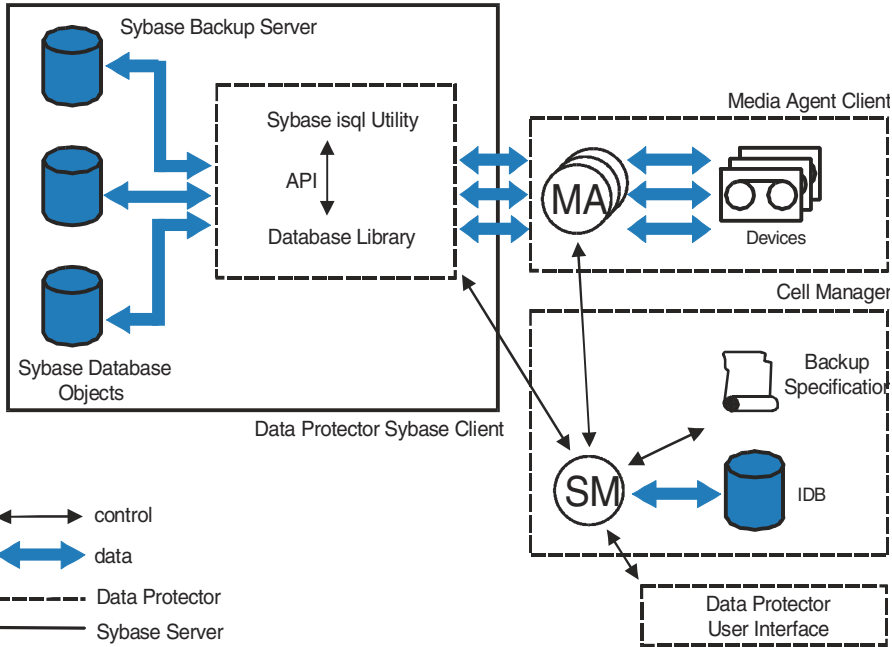
Limitations

Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a list of supported platforms and general Data Protector limitations and requirements. This section describes limitations specific to this integration.

- Do not use double quotes for object-specific pre-exec and post-exec commands. These commands are optionally entered as integration-specific options during the creation of backup specifications.
- The concurrency value greater than 1 is supported only with Sybase 12.x.
- On Windows, the maximum number of parallel streams is 13 for local backups and four for network backups.

Integration Concepts

Figure 1-1 Sybase Backup Concept



Data Protector integrates with Sybase Backup Server through the **Data Protector Database Library** based on a common library called Data Protector BAR (**B**ackup **A**nd **R**estore). The Data Protector Database Library channels communication between the Data Protector Session Manager (SM), and, via the Sybase Backup Server Application Programming Interface (API), to the Sybase isql utility. See Figure 1-1 for the architecture of the Data Protector Sybase integration.

Table 1-1 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
----	--

Table 1-1

Legend

API	Sybase Backup Server Application Programming Interface
Database Library	The Data Protector set of routines that enables the data transfer between the Sybase Backup Server and Data Protector.
MA	Data Protector General Media Agent

The isql Utility

The isql utility is a standalone program that sends commands to Sybase Backup Server, formatting the results and printing them on the standard output. When a request to execute a backup or restore is received, isql initiates a session with both Sybase Backup Server and Data Protector.

Backup Specification

Backup and restore commands are issued via the Data Protector User Interface or using the Sybase isql command-line interface. The list of objects to be backed up, together with backup options and the set of devices to be used are kept in the Data Protector backup specification.

Sybase Backup Server API

For a backup, Data Protector receives databases and transaction logs from Sybase Backup Server through the **Sybase Backup Server API** and writes them to devices on Data Protector clients using the General Media Agents. The Data Protector `ob2sybase` command starts the `sybase_<SYBASESERVERNAME>.sh` script (UNIX systems) or `syback.exe` program (Windows systems), which then starts the isql commands used for backup. The Data Protector `ob2sybase` command keeps the number of parallel backup streams to an optimum level during backup.

For a restore, Data Protector retrieves the requested databases and transaction logs from media and sends them through the Sybase Backup Server API to Sybase Backup Server, which writes them to disks.

While Sybase Backup Server is responsible for read/write operations to disk, Data Protector manages devices and media used for backup and restore sessions, and provides other powerful media management features before, during, and after backup sessions.

What's Next?

Equipped with the working concepts of the Data Protector Sybase integration, go on to upgrade or install the integration.

Data Protector Sybase Configuration File

Data Protector stores the Sybase integration parameters for every configured Sybase *instance* in the following file on the Cell Manager:

- `/etc/opt/omni/server/integ/config/Sybase/<client_name>%<instance_name>` (HP-UX and Solaris systems)
- `<Data_Protector_home>\Config\Server\Integ\Config\Sybase\<client_name>%<instance_name>` (Windows systems).

The parameters stored in the configuration file are those entered during the configuration of this integration, as described in “Configuring a Sybase Server” on page 20. These parameters are:

- the Sybase home directory,
- the full path for the Sybase `isql` command,
- the Sybase backup operator username and password,
- the name of the Sybase `<SYBASE_ASE>` directory (Sybase 12.x only),
- the name of the Sybase `<SYBASE_OCS>` directory (Sybase 12.x only),
- environmental variables (for example, various language environments support; see “Sybase Character Sets” on page 77)

The Data Protector Sybase configuration file is generated and parameters are written to it during the following events:

- during configuration of the integration (using the `util_sybase.exe` command or the Data Protector GUI)
- during creation of a backup specification if the configuration parameters are changed
- when the configuration parameters are changed (using the `util_sybase.exe` command, the `util_cmd` command, or the Data Protector GUI)

See “Configuring a Sybase Server” on page 20 for more information on configuring the integration using the `util_sybase.exe` command or the Data Protector GUI. See “Setting, Retrieving, and Listing Data Protector Sybase Configuration File Parameters Using the CLI” on page 11 for more information using the `util_cmd` command.

Configuration File Syntax The syntax of the file is as follows:

IMPORTANT

To avoid problems with your backups, ensure that the syntax of your configuration file matches the examples by using the `util_sybase.exe` command, the `util_cmd` command, or the Data Protector GUI. Do not edit the file manually.

```
Home_Dir="<SYBASE_HOME>" //homedir path
ISQL_path="<ISQL_PATH>" //path of isql executable
SA_user="<SYBASE_USER>"
SA_password="<SYBASE_PASSWORD>"
SYBASE_ASE="<SYBASE_ASE>" //path of ASE directory inside of
Home_Dir
SYBASE_OCS="<SYBASE_OCS>" //path of OCS directory inside of
Home_Dir
Isql_opts="<options>" //options, passed to isql executable
Environment="{ }
```

Example of Configuration File

This is an example of the Data Protector Sybase configuration file:

```
Home_Dir="/applications/sybase/"
ISQL_path="/applications/sybase/bin/isql"
SA_user="user1"
SA_password="3245s3sf4gsd5"
SYBASE_ASE="ASE-12_0"
SYBASE_OCS="OCS-12_0"
Isql_opts="<options>"
Environment="{ }
```

Setting, Retrieving, and Listing Data Protector Sybase Configuration File Parameters Using the CLI

Data Protector Sybase configuration file parameters are normally written to the Data Protector Sybase configuration file after the completed configuration of the Sybase instance in Data Protector using the `util_sybase.exe` command or the Data Protector GUI.

The `util_cmd` Command

You can set, retrieve, or list the Data Protector Sybase configuration file parameters using the `util_cmd -putopt` (setting a parameter), `util_cmd -getopt` (retrieving a parameter), or `util_cmd -getconf` (listing all parameters) command on the Data Protector Sybase client. The command resides in the `/opt/omni/sbin` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_cmd` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

- On UNIX: `export OB2BARHOSTNAME=<virtual_hostname>`
- On Windows: `set OB2BARHOSTNAME=<virtual_hostname>`

The `util_cmd` Synopsis

The syntax of the `util_cmd` command is as follows:

```
util_cmd -getconf[ig] Sybase <Sybase_instance> [-local \  
<filename>]
```

```
util_cmd -getopt[ion] [Sybase <Sybase_instance>] \  
<option_name> [-sub[list] <sublist_name>] [-local \  
<filename>]
```

```
util_cmd -putopt[ion] [Sybase <Sybase_instance>] \  
<option_name> [<option_value>] [-sub[list] <sublist_name>] \  
[-local <filename>]
```

where:

`<option_name>` is the name of the parameter

`<option_value>` is the value for the parameter

`[-sub[list] <sublist_name>]` specifies the sublist in the configuration file which a parameter is written to or taken from.

`[-local <filename>]` specifies one of the following:

- When used with the `-getconf [ig]` option, it specifies a filename that the command output is written to. If the `-local` option is not specified, the output is written to the standard output.
- When used with the `-getopt [ion]`, it specifies a filename of the file from which the parameter and its value are to be retrieved from and then written to the standard output. If the `-local` option is not specified, the parameter and its value are retrieved from the Data Protector Sybase configuration file and then written to the standard output.
- When used with the `-putopt [ion]` option, it specifies a filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the Data Protector Sybase configuration file.

The `util_cmd` Example

To enable this support in the Data Protector Sybase integration, you need to specify the environmental variable `OB2_ISQL_OPTS` in the Data Protector Sybase configuration file as follows:

```
util_cmd -putopt [ion] Sybase <instance_name> OB2_ISQL_OPTS \  
="-J<char_set>" -sublist Environment
```

where `<char_set>` is the character set to be used. For example:

```
util_cmd -putopt [ion] Sybase <instance_name> OB2_ISQL_OPTS \  
="-Jsjis" -sublist Environment
```

The above command will enable the Sybase Japanese language environment support. See also “Sybase Character Sets” on page 77.

Return Values

The `util_cmd` command displays a short status message after each operation (written to the standard error):

- Configuration read/write operation successful.
This message is displayed when all the requested operations have been completed successfully.

- Configuration option/file not found.
This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.

- Configuration read/write operation failed.

This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable or the Data Protector Sybase configuration file is missing on the Cell Manager.

Configuring the Integration

After the installation, the integration is not yet ready for use. The following subsections provide instructions for configuring the integration so that it functions properly.

Configuration Overview

To configure the integration, follow these steps:

1. Configure a Sybase user

This is a user with appropriate rights in both Data Protector and Sybase environments as described in “Configuring a Sybase User in Data Protector” on page 17.

2. Configure a Sybase Server

This is a client running Sybase Backup Server. See “Configuring a Sybase Server” on page 20 for instructions about configuring this client.

3. Configure a Sybase backup

Configure the devices and media needed for your backup, and create a Data Protector backup specification (a file in which you specify the objects that you want to back up), the media and devices to which you want your data to be backed up, as well as powerful Data Protector backup options, which, for instance, allow you to schedule your backup to specific or periodic times). See “Configuring a Sybase Backup” on page 30 for instructions about configuring your backup.

Before You Begin Configuring

Check the following before you start configuring:

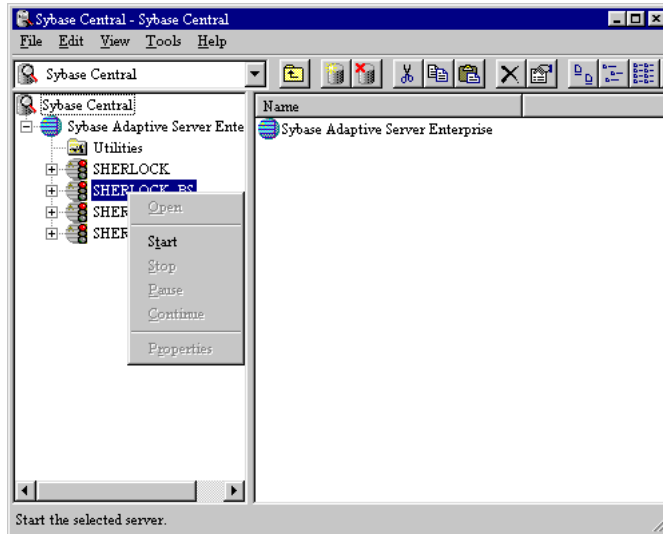
- ✓ The integration software has been installed on all Sybase Servers you want to back up. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.
- ✓ Sybase SQL Server and Sybase Backup Servers are correctly installed and running.

NOTE You must have Sybase Backup Servers running on the same machine as your Sybase SQL Server. The Sybase Backup Servers must be listed in the master..syservers table. This entry is created during installation or upgrade, and should not be deleted.

NOTE If your Sybase SQL Server is running a language other than English, see “Sybase Character Sets” on page 77.

Windows To start a Sybase server from Sybase Central, without connecting to it, right-click the server, and then click **Start** in the pop-up menu:

Figure 1-2 Starting Sybase Servers Using Sybase Central on Windows Systems



UNIX Proceed as follows:

1. Log on to your Sybase Backup Server as user `sybase`.

Integrating Sybase and Data Protector

Configuring the Integration

2. Type in the following command in the Sybase Backup Server home directory:

```
isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>
```

where <PASSWORD> is your password to Sybase SQL Server, <SYBASESERVERNAME> the name of Sybase SQL Server and <SA> is the Sybase user.

3. In the first line, type in the following:

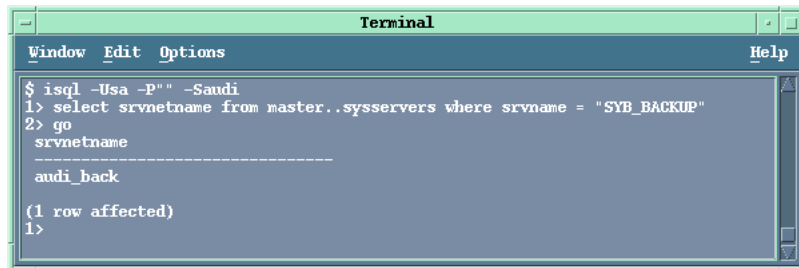
```
select srvnetname from master..sys.servers where srvname = "SYB_BACKUP"
```

and in the second line type go.

The name of Sybase Backup Server is returned.

In the Figure 1-3, the name of Sybase SQL Server is audi. The name of Sybase Backup Server is audi_back.

Figure 1-3 Checking if Sybase Backup Server Is Up on UNIX Systems



```
Terminal
Window Edit Options Help
$ isql -Usa -P'' -Saudi
1> select srvnetname from master..sys.servers where srvname = "SYB_BACKUP"
2> go
srvnetname
-----
audi_back

(1 row affected)
1>
```

- ✓ You can successfully run a filesystem backup of the Sybase Server. Configure and run a Data Protector filesystem backup of Sybase Server for test purposes. By doing this, you check whether the Sybase Server and the Data Protector Cell Manager can communicate properly. In case of errors, this type of backup is much easier to troubleshoot than the integration itself. The configuration procedure includes installing a Disk Agent on the Sybase Server, configuring appropriate devices and media (use any device), creating a filesystem backup specification, starting the backup, and then restoring the data. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.

- ✓ For each running Sybase instance you will have to supply the following information when configuring a backup specification:

Information Needed to Configure a Backup

- Sybase Backup Server home directory, for example, /applications/sybase (UNIX systems) or c:\sybase (Windows systems)
- Full pathname of the Sybase isql command, for example, /applications/sybase/bin/isql (UNIX systems), or c:\sybase\bin\isql.exe (Windows systems)
- Sybase SQL Server name
- Username and password of the Sybase user who has at least the backup role set in Sybase

In case of Sybase 12.x you will also need to provide the following:

- the name of the Sybase <SYBASE_ASE> directory and
- the name of the Sybase <SYBASE_OCS> directory.

For more information, refer to the *Sybase SQL Server System Administration Guide*.

Cluster-Aware Clients

- ✓ In a cluster environment, the environment variable OB2BARHOSTNAME must be defined as the virtual hostname before running the configuration from the command line (on the client). When running the configuration from the GUI, this is not required. The OB2BARHOSTNAME variable is set as follows:
 - On UNIX: `export OB2BARHOSTNAME=<virtual_hostname>`
 - On Windows: `set OB2BARHOSTNAME=<virtual_hostname>`

Configuring a Sybase User in Data Protector

On UNIX, to start a Sybase backup session, you need an operating system logon with sufficient privileges on the system where Sybase SQL Server is running.

Who Is the Sybase User?

To find a Sybase user with sufficient backup and restore privileges, run the following command on Sybase Server:

```
$ ls -l <SYBASE_HOME>/bin/isql (Sybase 11.9.3)
```

```
$ ls -l <SYBASE_HOME>/OCS-12_0/bin/isql (Sybase 12.x)
```

where `<SYBASE_HOME>` is the home directory of Sybase SQL Server.

Sybase SQL Server returns a user, in this case the user `sybase` in the group `sybase` with the following permissions. For example:

```
-rwsr-sr-x 1 sybase sybase 1569592 July 12 1999
/applications/sybase/bin/isql (Sybase 11.9.3)

-rwxr-xr-x 1 sybase sybase 1664672 Mar 20 2000
/applications/sybase.12/OCS-12_0/bin/isql (Sybase 12.x)
```

Owner of a Sybase Backup Specification

Using this logon, the user must be able to back up and restore Sybase objects. To start a backup of a Sybase object using Data Protector, the user must then become the owner of a Data Protector Sybase backup specification.

This user (for example, the user `sybase` in the group `sybase`) must be added to the Data Protector `admin` and `operator` groups.

Table 1-2 shows privileges of members of the Data Protector `operator` or `admin` groups. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for detailed information on user rights.

Table 1-2 Data Protector Admin and Operator User Groups and Their Access Rights

User Group	Access Rights
admin	Allowed to configure Data Protector and start backups, restores, and all other available operations. A member of this group has the rights of the <code>root</code> user on the UNIX or of the administrator on the Windows platform.
operator	Allowed to start backups and restores, and to respond to mount requests.

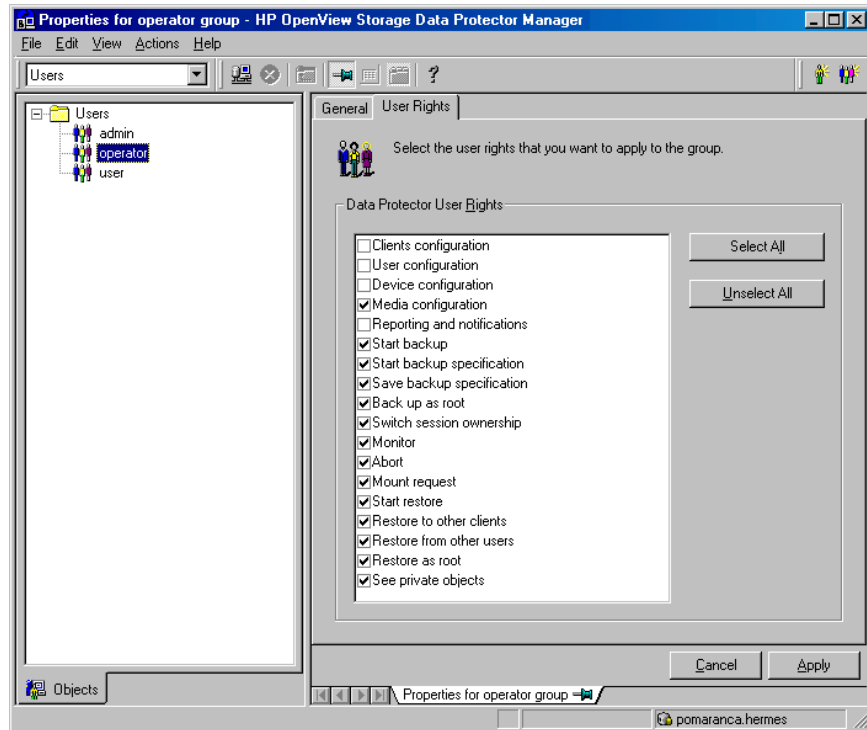
Data Protector User Rights

Data Protector user rights are user configurable. Ensure that the See private objects user right of the Data Protector operator group is selected. This right allows a user to browse private objects. Note that this does not give the user permission to restore data. To configure this user, proceed as follows:

Configuring Data Protector User Rights

1. In the Context List, select Users.
2. In the Results Area, right-click Operator and click Properties.

Figure 1-4 Data Protector Operator Group User Rights



3. If the See private objects user right is selected, click Apply.

What's Next?

In this section you configured the Sybase user, a user with appropriate rights in both the Data Protector and Sybase environments. You are now ready to configure your Sybase Server.

Configuring a Sybase Server

Each client running Sybase Backup Server must be configured for proper integration with Data Protector.

Cluster-Aware Clients

When configuring the Data Protector Sybase integration in a cluster environment, configure it only on one of the cluster nodes per one Sybase server, because the configuration files reside on the Cell Manager. Use the cluster virtual hostname when configuring the integration.

IMPORTANT

On UNIX, *after* you have configured the Sybase Server, using either the CLI or GUI as described further on, make sure that the Sybase user configured as described in the “Configuring a Sybase User in Data Protector” on page 17 has permissions to read the Data Protector Sybase configuration file. For more information on Data Protector Sybase configuration file, see “Data Protector Sybase Configuration File” on page 9.

Before you configure a Sybase Server, ensure that Sybase Backup Server is running. See “Before You Begin Configuring” on page 14 for instructions. On Windows, you can configure a Sybase Server using the Data Protector GUI, whereas on UNIX you can configure it either using the Data Protector GUI or the Data Protector CLI.

Using the Data Protector CLI (UNIX Systems Only)

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the configuration from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

- On UNIX: `export OB2BARHOSTNAME=<virtual_hostname>`
- On Windows: `set OB2BARHOSTNAME=<virtual_hostname>`

Configuring a Sybase Server

Log in as `root` and execute the following command script for each Sybase Server you want to configure:

```
/opt/omni/sbin/util_sybase.exe -CONFIG <SYBASE_SERVERNAME>  
<SYBASE_HOME> <ISQL_PATH> <SYBASE_USER> <SYBASE_PASSWORD>  
<SYBASE_ASE> <SYBASE_OCS>
```

NOTE

The `<SYBASE_ASE>` and `<SYBASE_OCS>` parameters are only required for Sybase 12.x.

Where:

- `<SYBASE_SERVERNAME>`
is the name of the Sybase SQL Server; if the Sybase SQL Server is configured in a cluster, this is the virtual hostname of the Sybase SQL Server,
- `<SYBASE_HOME>`
is the Sybase home directory, for example, `/applications/sybase/`,
- `<ISQL_PATH>`
is the full path for the Sybase `isql` command, for example, `/applications/sybase/bin/isql`
- `<SYBASE_USER>`
is the name of the Sybase user who has at least the backup role set in Sybase,
- `<SYBASE_PASSWORD>`
is the Sybase password for this user,
- `<SYBASE_ASE>`
is the name of the Sybase `<SYBASE_ASE>` directory (Sybase 12.x only) and
- `<SYBASE_OCS>`
is the name of the Sybase `<SYBASE_OCS>` directory (Sybase 12.x only).

In case of Sybase 12.x, the command and its output should look like:

```
util_sybase.exe -CONFIG koperton12 /applications/sybase.12/  
/applications/sybase.12/OCS-12_0/bin/isql sa "" ASE-12_0  
OCS-12_0  
*RETVAL*0
```

Integrating Sybase and Data Protector

Configuring the Integration

In case of Sybase 11.9.3, the command and its output should look like:

```
util_sybase.exe -CONFIG slaine /applications/sybase/  
/applications/sybase/isql sa ""  
  
*RETVAL*0
```

Upon successful configuration, a dialog box with the message `*RETVAL*0` is returned.

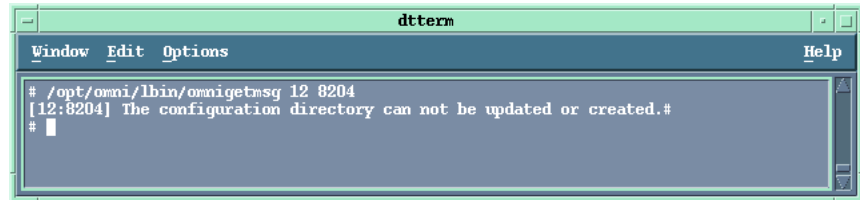
In case of an error, the error number is displayed in the form `*RETVAL*<error number>`.

To get the error description, start the command,
`/opt/omni/lbin/omnigetmsg 12 <error_number>`.

In the example above, an error message, number 8204 was received after failing to configure the Sybase Server. To get the error description, the `omnigetmsg` command was used:

Figure 1-5

Getting an Error Description on UNIX Systems



Using the Data Protector GUI

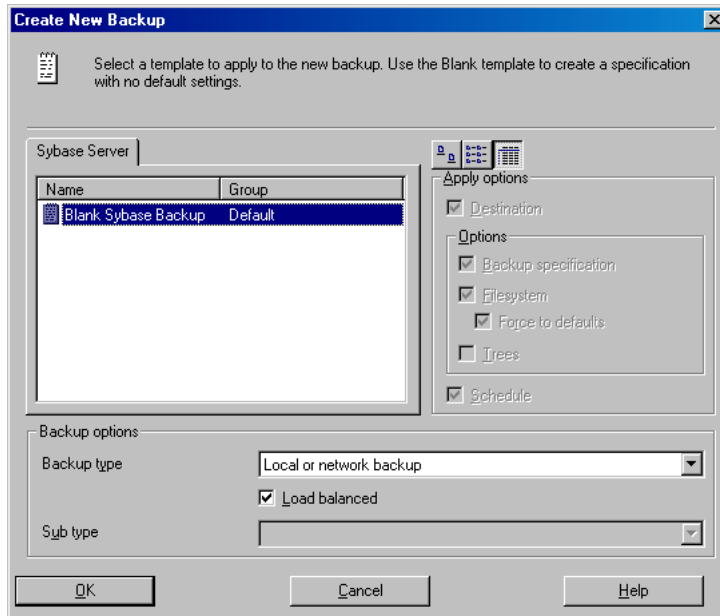
To configure a Sybase Server, perform the following steps in the HP OpenView Storage Data Protector Manager:

Configuring a Data Protector Client

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click Sybase and click Add Backup.

The Create New Backup dialog box is displayed.

Figure 1-6 **Creating a New Sybase Backup**



Select the **Load balanced** option, which enables Data Protector to automatically balance the usage of devices that you select for the backup.

3. Click **OK**.

On UNIX, in the **Results Area**, enter the following information:

The UNIX user name and group of the Sybase user, referred to in the section, “Configuring a Sybase User in Data Protector” on page 17. For example, user `sybase` in group `sybase`.

In a cluster environment, select the virtual hostname from the **Client** drop-down list.

Figure 1-7 **Configuring a Sybase Server on UNIX Systems**

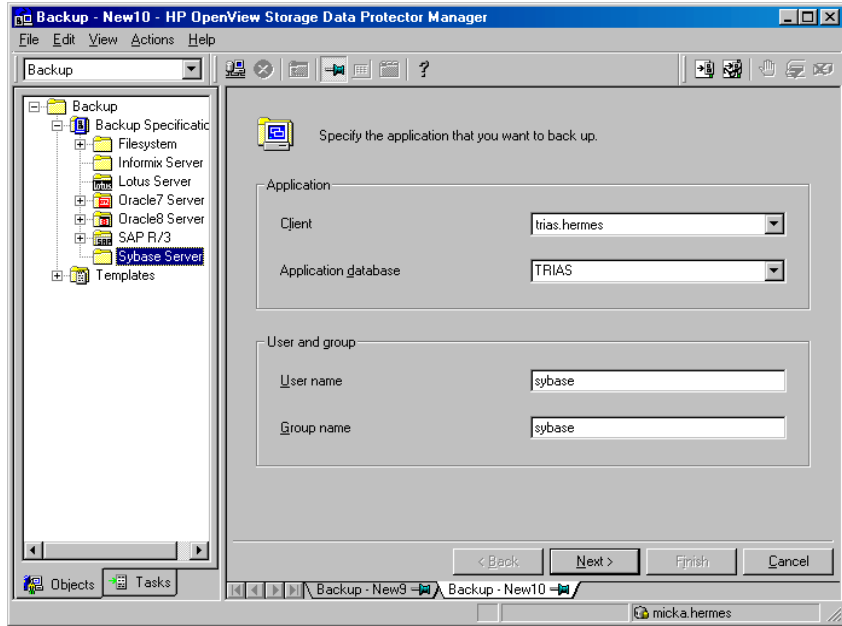
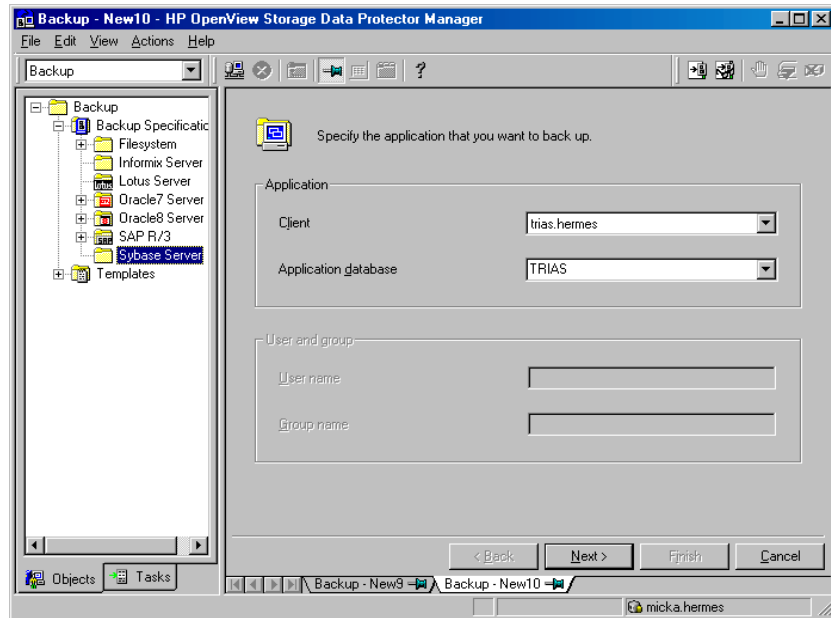


Figure 1-8 Configuring a Sybase Server on Windows Systems



Click Next.

An error message pertaining to the Sybase database is displayed. Click OK and the Configure Sybase dialog box is displayed.

4. Enter the Sybase SQL Server home directory, for example, /applications/sybase/ (UNIX systems) or c:\sybase (Windows systems), the full pathname of the Sybase isql command, for example, /applications/sybase/bin/isql (UNIX systems) c:\sybase\bin\isql.exe (Windows systems), the username and password of the Sybase user who has at least the backup role set in Sybase. In case of Sybase 12.x you also need to enter the Sybase <SYBASE_ASE> directory and the Sybase <SYBASE_OCS> directory. Note that in case of Sybase 12.x, the full pathname of the Sybase isql command is different, for example /applications/sybase/OCS-12_0/bin/isql (UNIX systems) c:\sybase\OCS-12_0\bin\isql.exe (Windows systems).

Figure 1-9 **Configuring a Sybase Server on UNIX Systems**

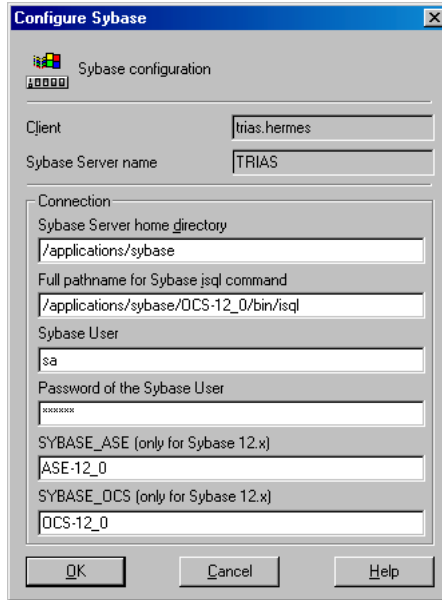
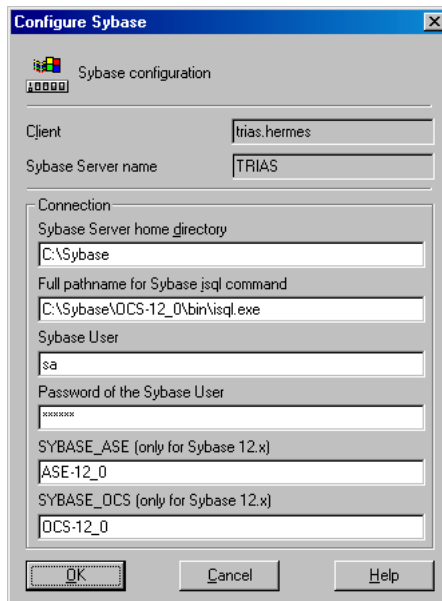


Figure 1-10 **Configuring a Sybase Server on Windows Systems**



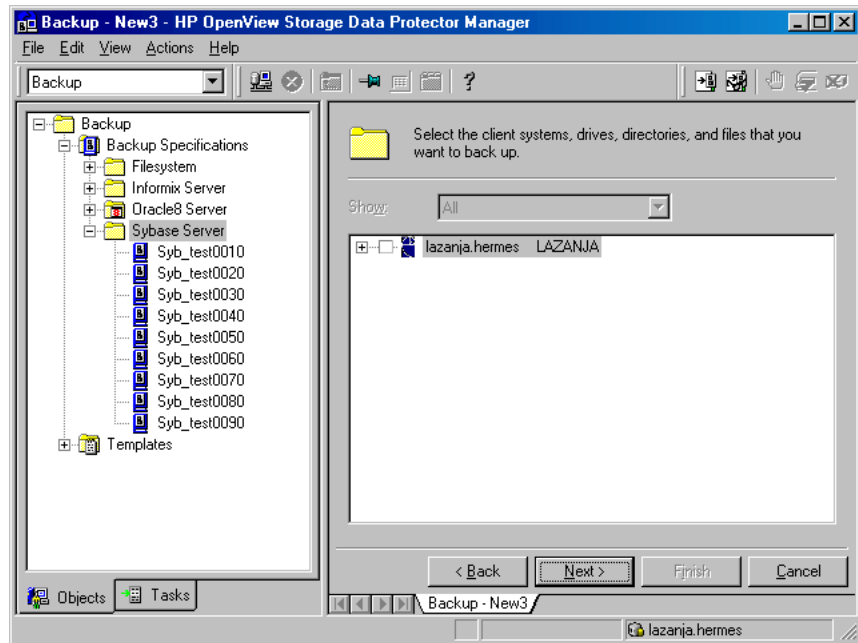
5. Click OK.

NOTE

If you receive a message that your Sybase Backup Server is not running, leave the configuration, start your Sybase Backup Server and proceed from there.

Upon successful configuration, the next step of the wizard is displayed in which you can start configuring your backup.

Figure 1-11 Successful Configuration



What Happens?

The following happens after saving the configuration:

Data Protector executes the `util_sybase.exe` command on the Sybase Server, which performs the following:

1. It saves the configuration parameters in the Data Protector Sybase configuration file. For more information on the Data Protector Sybase configuration file, see “Data Protector Sybase Configuration File” on page 9.
2. It creates the `sybase_<SYBASESERVERNAME>.sh` script (UNIX systems) or the `syback.exe` program (Windows systems).
3. It checks connections to Sybase Backup Server.

What’s Next?

You may want to check if the integration is properly configured before you start using it to make backups and restores. The next section shows you how.

Checking the Sybase Configuration Using the Data Protector CLI (UNIX Systems Only)

On UNIX, you can check the Sybase configuration either using the Data Protector CLI, or the Data Protector GUI. If your Sybase SQL Server is running language other than English, use the Data Protector GUI for checking the configuration. For more information on how to check the Sybase configuration using the Data Protector GUI, see “Checking the Sybase Configuration Using the Data Protector GUI” on page 29.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before checking the configuration from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

```
export OB2BARHOSTNAME=<virtual_hostname>
```

Checking the Configuration

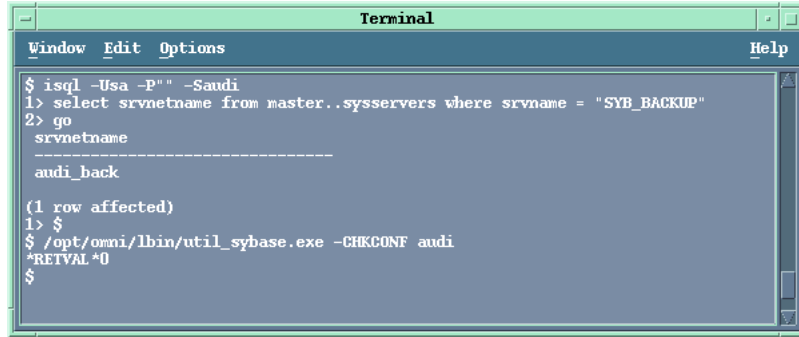
To check the Sybase configuration using the Data Protector CLI, log in as the Sybase user, and start the following command:

```
/opt/omni/lbin/util_sybase.exe -CHKCONF <SYBASESERVERNAME>
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server.

Figure 1-12 shows how to first verify that Sybase SQL Server, `audi`, is up and running and then check its configuration. The configuration is OK, since the `*RETV*0` message was returned.

Figure 1-12 **Checking the Sybase Configuration on UNIX Systems Using the CLI**



```
Terminal
Window Edit Options Help
$ isql -Usa -P" " -Saudi
1> select srvnetname from master..sys.servers where srvname = "SYB_BACKUP"
2> go
srvnetname
-----
audi_back

(1 row affected)
1> $
$ /opt/omni/lbin/util_sybase.exe -CHKCONF audi
*RETVAL *0
$
```

In case of an error, the error number is displayed in the form
`*RETVAL* <error number>`.

To get the error description, start the command,
`/opt/omni/lbin/omnigetmsg 12 <error_number>`.

What's Next?

Now that you have successfully configured your Sybase Server, go on and configure your backup.

Checking the Sybase Configuration Using the Data Protector GUI

On Windows, you can check the Sybase configuration using the Data Protector GUI only.

To check the configuration of your Sybase Server, proceed as follows in the HP OpenView Storage Data Protector Manager:

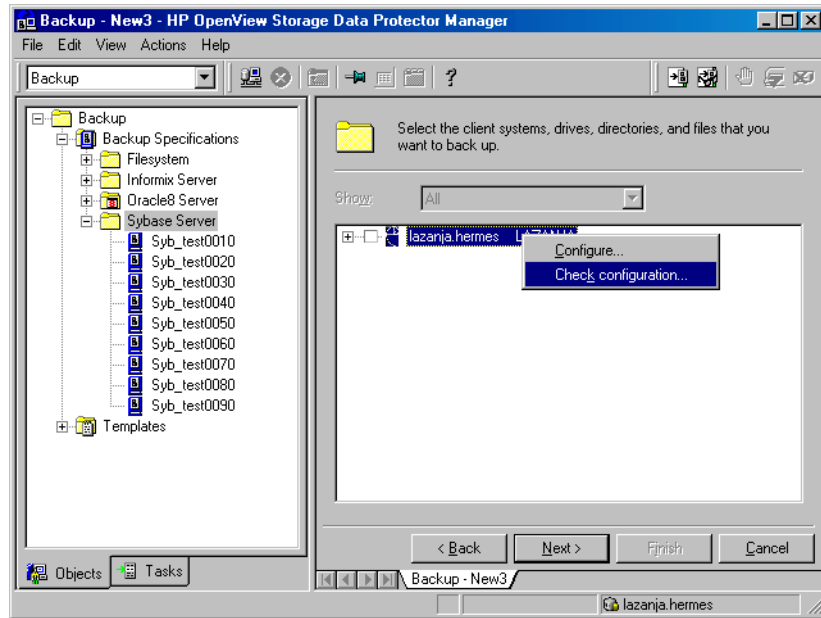
1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications, and then Sybase.
3. Go over the configuration procedure described in "Configuring a Sybase Server" on page 20 and the proceed with the procedure below.

Or, if you have already configured a backup specification, click it.

The Sybase Server is displayed. In the Context List, select Backup.

4. Right-click the client and then click Check Configuration.

Figure 1-13 Checking the Sybase Configuration on Windows Systems



5. If the integration is properly configured, a message is returned confirming this fact.

What's Next?

Now that you have successfully configured your Sybase Server, go on and configure your backup.

Configuring a Sybase Backup

To run backups and restores of your Sybase data, you need to configure Data Protector Sybase backup specifications. This section gives you instructions to this end.

To configure the backup of Sybase data, perform the following steps:

Configuration Steps

1. Configure devices, media, and media pools needed for the backup. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.

2. Create a Data Protector Sybase backup specification specifying the data that you want to back up, the media and devices to which you want your data to be backed up, as well as Data Protector backup options that define the behavior of your backup or restore session.

Creating a Data Protector Sybase Backup Specification

Sybase backup specifications are located in the following directory on the Cell Manager:

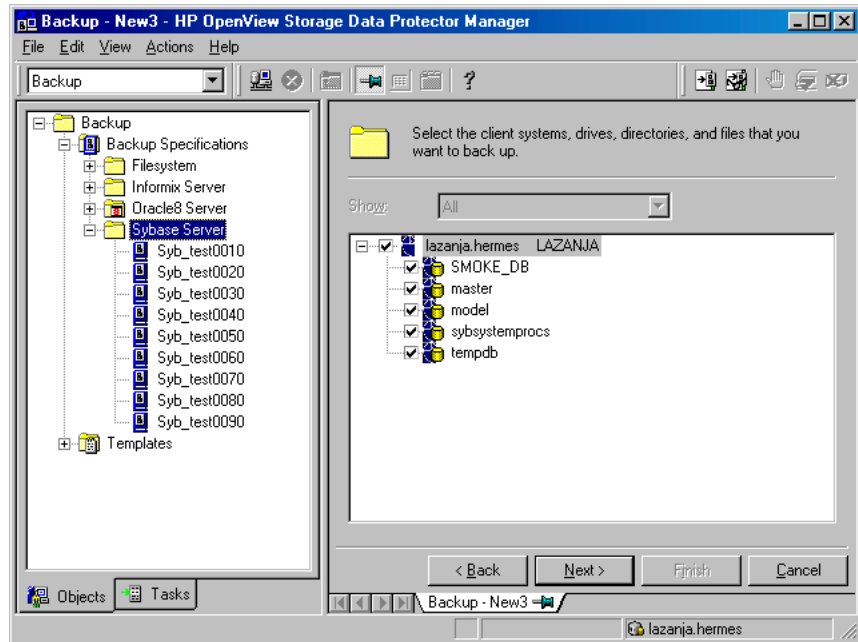
```
/etc/opt/omni/server/barlists/sybase (UNIX Cell Manager) or  
<Data_Protector_home>\config\server\barlists\sybase  
(Windows Cell Manager)
```

A Sybase backup specification is created using the Data Protector GUI. Ensure that you have appropriate privileges.

To create a Data Protector Sybase backup specification on a client with no backup specification configured, proceed from where you left off in “Configuring a Sybase Server” on page 20.

1. In the Results Area, select the databases you want to back up. The databases include user databases and **system databases**. In the example shown in Figure 1-14, all databases were selected for backup.

Figure 1-14 **Selecting Databases for Backup**



Click Next.

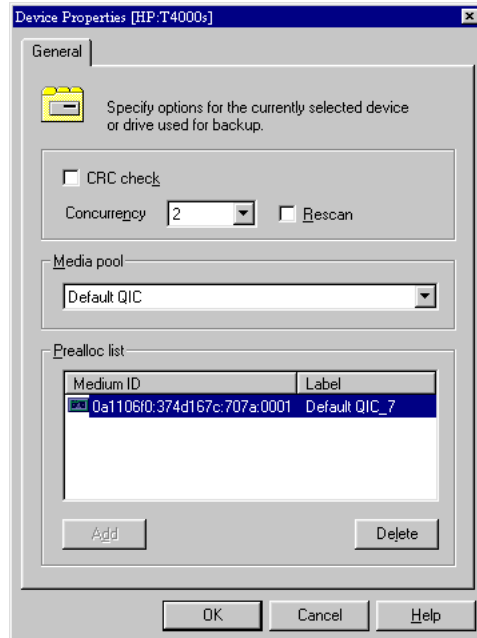
2. Select the device(s) you want to use for the backup.

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

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

3. Select the device you want to use and click Properties.
The Device Properties dialog box is displayed.

Figure 1-15 Specifying Device Properties



Specify the number of parallel backup streams in the *Concurrency* tab and the media pool you will use.

In the example shown in Figure 1-15, a concurrency of 2 and the Default QIC media pool were used.

IMPORTANT

Device concurrency values greater than one are possible only with the Sybase SQL Server 12.x version.

4. Click *Add*, to add specific media to the *Prealloc* list, a subset of media in the media pool used for backup, which also specifies the order in which media are used for backup.

Click *OK*.

5. Click *Next* to specify backup options.

**Object-Specific
Pre-Exec and
Post-Exec
Commands**

Specify the `Load Balancing` option. With this option set, Data Protector dynamically assigns backup objects to available devices. This enables devices to be used evenly and for backups to continue on available devices in case of failure of some device.

Under `Application Specific Options`, click `Advanced`, to specify pre-exec and post-exec commands that will be started on the Sybase Server *for each Sybase object*.

These commands are different from the pre-exec and post-exec commands in the `Backup Options` dialog box (which you reach by clicking `Advanced` under `Backup Specification Options`) in that they are valid only for the specific object you select and not for the whole client.

Under `General Information` in Figure 1-16, optionally specify the following:

- `Pre-exec`

a command that will be started on the Sybase Server before the backup. The command is started by the `ob2sybase.exe` command. On Windows, the command must reside in the `<Data_Protector_home>\bin` directory and only the filename must be provided in the backup specification. On UNIX, the full path for the command must be provided.

- `Post-exec`

a command that will be started on the Sybase Server after the backup. The command is started by the `ob2sybase.exe` command. On Windows, the command must reside in the `<Data_Protector_home>\bin` directory and only the filename must be provided in the backup specification. On UNIX, the full path for the command must be provided.

IMPORTANT

Do not use double quotes for object-specific pre-exec and post-exec commands.

Figure 1-16 Object-Specific Pre- and Post-Exec Commands on UNIX Systems

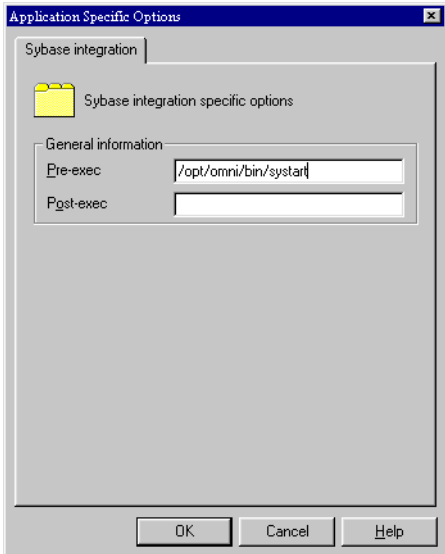
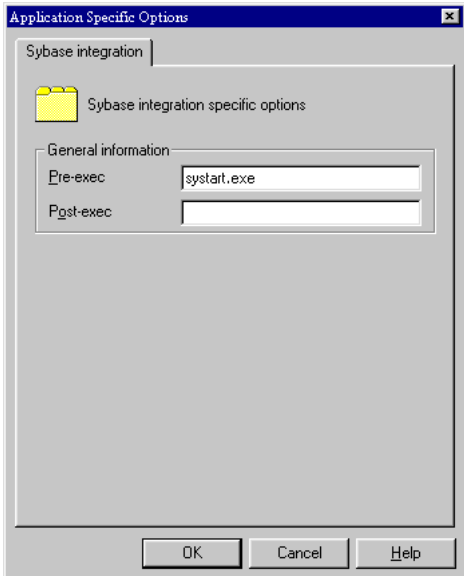


Figure 1-17 Object-Specific Pre- and Post-Exec Commands on Windows Systems



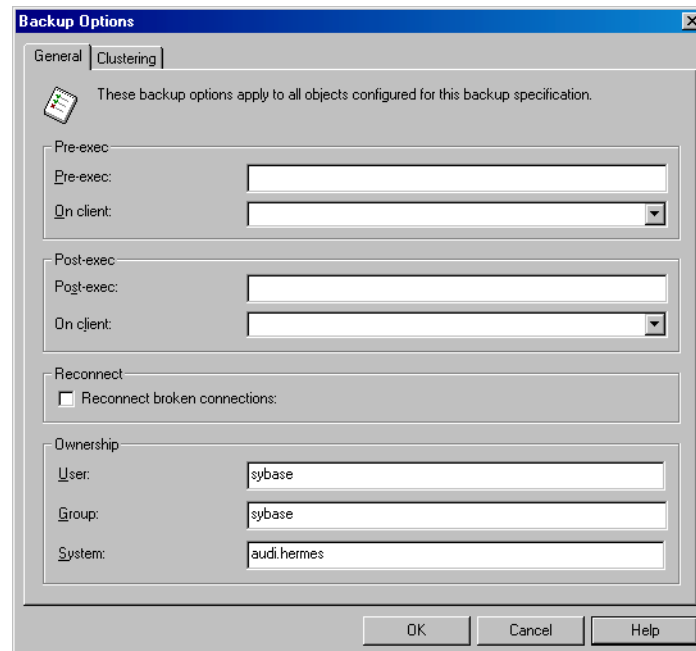
Changing the Sybase User on UNIX Systems

On UNIX, the Sybase user is the user who configured the backup. Changing ownership allows another user to start the configured backup and to later restore the backed up data. To change the Sybase user, proceed as follows:

1. In the Backup Specific Options group box, click Advanced.
The Backup Options dialog box is displayed.
2. Click General and edit the Ownership group box

Figure 1-18

Changing the Sybase User on UNIX Systems



Scheduling a New Backup Specification

Click OK and then Next, to schedule your backup specification. You can schedule your backup to start automatically and unattended on a specific date and time or at regular intervals for a period of up to a year in advance.

IMPORTANT

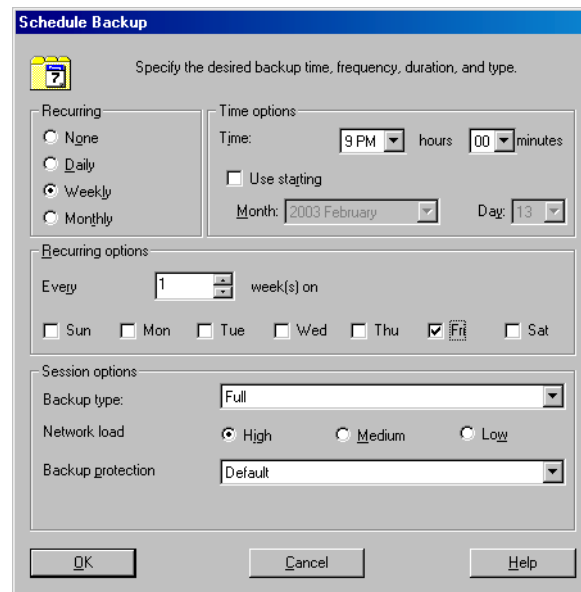
Sybase SQL Server allows the online backup of databases and transaction logs. Schedule frequent backups of your transaction log. The more often you back up your system, the less amount of work is lost should a system failure occur.

Scheduling Example

In this example, a full backup will be scheduled to start at 9.00 p.m. every Friday.

Click Add to open the Schedule Backup dialog box and specify the options as shown in Figure 1-19.

Figure 1-19 Scheduling a Weekly Full Backup



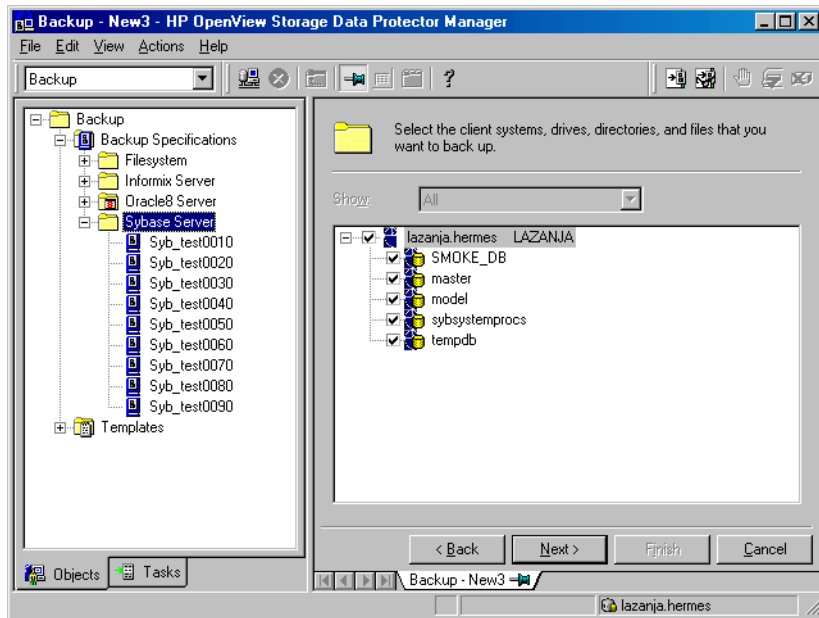
Click Next.

All the objects you selected for backup are displayed.

NOTE

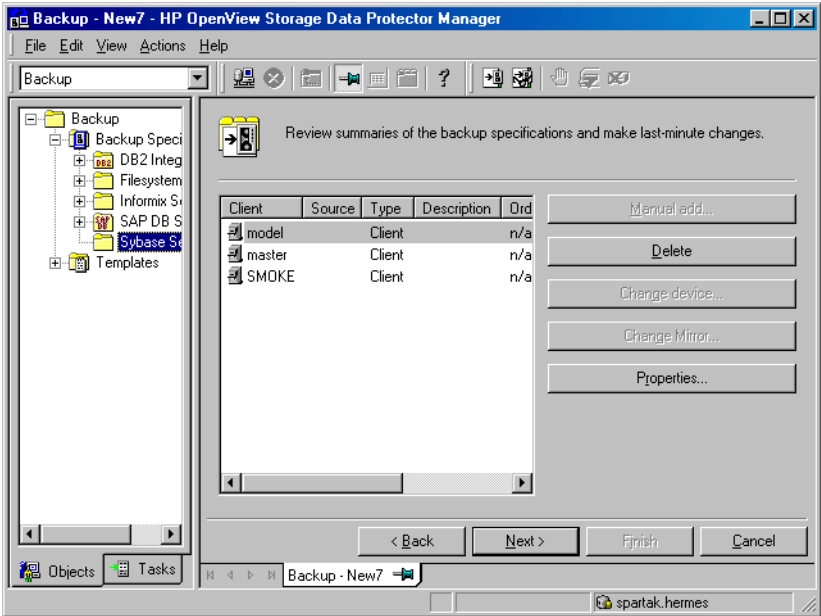
If you chose to back up the whole client by selecting the client as shown in Figure 1-20, then only the client is displayed in the Backup Specification Summary dialog box and not individual databases.

Figure 1-20 **Selecting the Whole Client For Backup**



In the example shown in Figure 1-21, individual objects were selected for backup. Hence, individual objects are displayed in the Backup Specification Summary dialog box.

Figure 1-21 Backup Specification Summary



You can also select the number of concurrent streams for each specific database by selecting the object and clicking Properties to open the Object Properties dialog box.

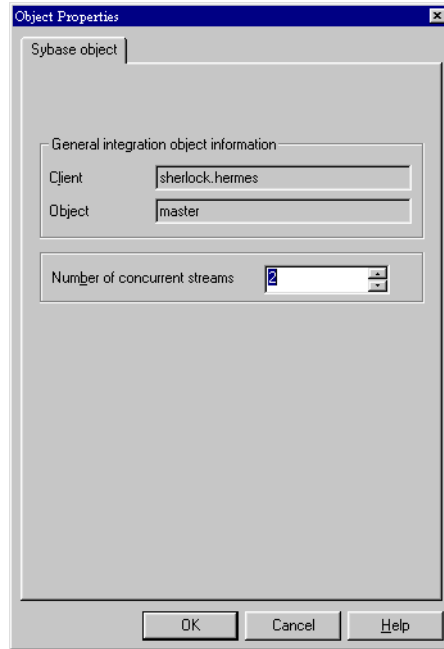
NOTE

Number of concurrent streams sets the number of Sybase database streams that are sent to backup devices. Depending on the device concurrency value set for each device, the streams are distributed among the backup devices.

IMPORTANT

Device concurrency values greater than 1 are possible only with the Sybase SQL Server 12.x version.

Figure 1-22 **Selecting the Number of Concurrent Streams**



The Sybase Backup Server splits the database into approximately equal portions and sends each portion to a different device. This is done concurrently on all devices, reducing the time required to back up an individual database or transaction log. This option is equivalent to Sybase *dump striping*. Refer to the *Sybase SQL Server Reference Manual* for more information.

TIP

To improve backup performance, back up your large databases to multiple streams.

You can now select individual objects and edit your backup specification options by clicking the `Properties` text box.

After defining the number of concurrent streams for all your objects, click `Next` and then save your backup specification. Click `Start Preview`, to test your backup specification.

**Editing Your
Backup
Specification**

Now you have created your backup specification and are ready to run your backups. You can always revert to your backup specification to edit it by selecting it by name in the Backup context. Click the appropriate tab and implement the changes you want. You need to save the backup specification afterwards.

What's Next?

Follow the steps in this section to configure other backup specifications you might need, for example, a backup specification to back up system databases.

Test your backup specification thoroughly before using it for production backups. See “Testing the Integration” on page 42.

Testing the Integration

Test your backup specifications thoroughly by previewing them, then running them on file devices and then finally on the actual devices you intend to use. To test your backup specifications, you can use either the Data Protector GUI or the Data Protector CLI.

Using the Data Protector GUI

To check if a backup specification has been properly configured, proceed with the following steps in the main HP OpenView Storage Data Protector Manager:

- Testing Procedure**
1. In the `Context List`, select `Backup`.
 2. In the `Scoping Pane`, expand `Backup`, and then `Backup Specifications`. Expand `Sybase Server` and then right-click the backup specification you want to preview.
 3. Click `Preview Backup` to open the `Start Preview` dialog box. Select the type of backup you want to run as well as the network load. For a description of these options, press **F1**.

Observe the generated messages. The “Session completed successfully” message is displayed at the end of a successful backup session of the *FullSybase* backup specification.

Using the Data Protector CLI

You can check if a Data Protector Sybase backup specification is properly configured using the Data Protector `omnib` command:

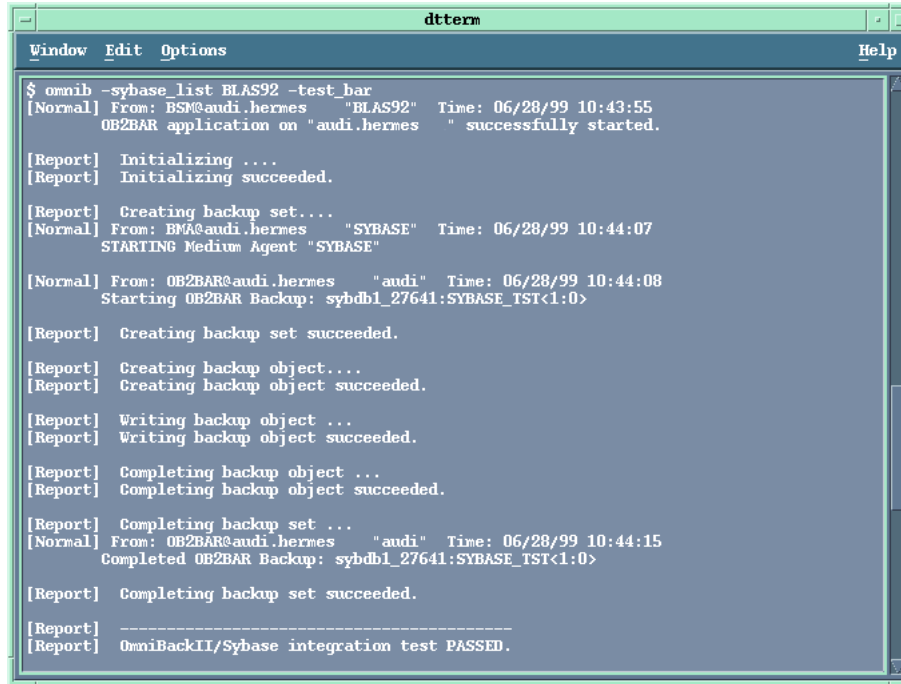
```
omnib -sybase_list <backup_specification_name> -test_bar
```

UNIX Example

In the following example, the backup specification is called BLAS92.

Figure 1-23

Testing the Configuration of the BLAS92 Backup Specification



```
dtterm
Window Edit Options Help
$ omni -sybase list BLAS92 -test bar
[Normal] From: BSM@audi.hermes "BLAS92" Time: 06/28/99 10:43:55
OB2BAR application on "audi.hermes" successfully started.

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

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

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

[Report] Creating backup set succeeded.

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

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

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

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

[Report] Completing backup set succeeded.

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

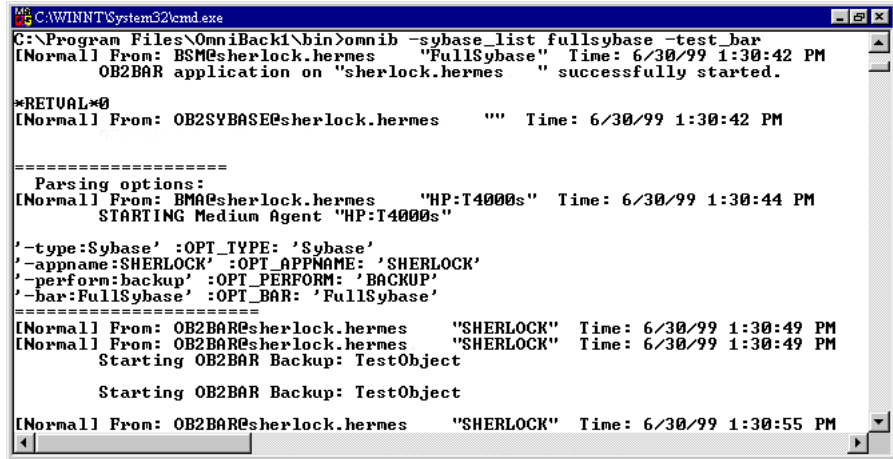
Upon successful configuration, a dialog box with the message *RETVAL *0 is returned.

In case of an error, the error number is displayed in the form *RETVAL*<error number> and description.

To get the error description, run the following command:
/opt/omni/sbin/omnigetmsg 12 <error_number>

Windows Example In the following example, the backup specification is called FullSybase.

Figure 1-24 Testing the Configuration of the FullSybase Backup Specification



```
C:\Program Files\OmniBack\bin>omnib -sybase_list fullsybase -test_bar
[Normal] From: BSM@sherlock.hermes "FullSybase" Time: 6/30/99 1:30:42 PM
OB2BAR application on "sherlock.hermes" successfully started.

*RETVAL*0
[Normal] From: OB2SYBASE@sherlock.hermes "" Time: 6/30/99 1:30:42 PM

=====
Parsing options:
[Normal] From: BMA@sherlock.hermes "HP:T4000s" Time: 6/30/99 1:30:44 PM
STARTING Medium Agent "HP:T4000s"

'-type:Sybase' :OPT_TYPE: 'Sybase'
'-appname:SHERLOCK' :OPT_APPNAME: 'SHERLOCK'
'-perform:backup' :OPT_PERFORM: 'BACKUP'
'-bar:FullSybase' :OPT_BAR: 'FullSybase'
=====
[Normal] From: OB2BAR@sherlock.hermes "SHERLOCK" Time: 6/30/99 1:30:49 PM
[Normal] From: OB2BAR@sherlock.hermes "SHERLOCK" Time: 6/30/99 1:30:49 PM
Starting OB2BAR Backup: TestObject

Starting OB2BAR Backup: TestObject

[Normal] From: OB2BAR@sherlock.hermes "SHERLOCK" Time: 6/30/99 1:30:55 PM
```

Upon successful configuration, a dialog box with the message *RETVAL *0 is returned.

In case of an error, the error is described.

What Happens? The given procedure performs a backup preview that tests:

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

The command tests only the Data Protector part of the configuration.

Backing Up a Sybase Database

In case of system failure, you can make a useful restore of your databases if you have been making *regular* backups of the databases *and* transaction logs.

Before You Begin To be prepared for hardware or software failure on your server, the two most important tasks are:

- Performing frequent backups of the **system databases**.

Back up your master database each time you create, alter or delete any device or database.

Back up your model database each time you change it. In case of a system failure, restore the model database as you would a user database.

If you make changes to the Sybase system procedure database or add your own stored procedures to the database, back up the database regularly.

- Keeping a copy of the following systems tables:

sysusages
sysdatabases
sysdevices
sysloginroles
syslogins

Backup Methods To run a backup of a Sybase database, use any of the following methods:

- Schedule the backup of an existing Sybase backup specification using the Data Protector Scheduler. See “Scheduling an Existing Backup Specification” on page 47.
- Start an interactive backup of an existing Sybase backup specification. You can start a backup using the Data Protector GUI or the Data Protector CLI. See “Running an Interactive Backup” on page 50.
- Start a backup using the Sybase CLI. See “Backing Up Using Sybase Commands” on page 52.

Backup Types

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

Table 1-3 Sybase Backup Types

Type	Description
Full	The backup of selected databases and transaction logs.
Transaction	The backup of transaction logs that have been modified since the last backup, providing a record of any changes made since the last full or transaction backup.

Refer to the *Sybase SQL Server Administration Guide* for more details on the backup types.

What Happens?

The following happens when you start a Sybase backup:

1. Data Protector executes the `ob2sybase` command on the Sybase Server. This command starts the Data Protector `util_sybase.exe` command to check the configuration of the integration. Then the `ob2sybase` command starts the `sybackup_<SYBASESERVERNAME>.sh` scripts (UNIX systems) or the `syback.exe` (Windows systems) programs in parallel. Each script starts an `isql` backup command.
2. The Sybase `isql` backup command initiates a backup session on Sybase Backup Server. During the backup session, Sybase Backup Server reads data from the disk and sends it to Data Protector for writing to devices.

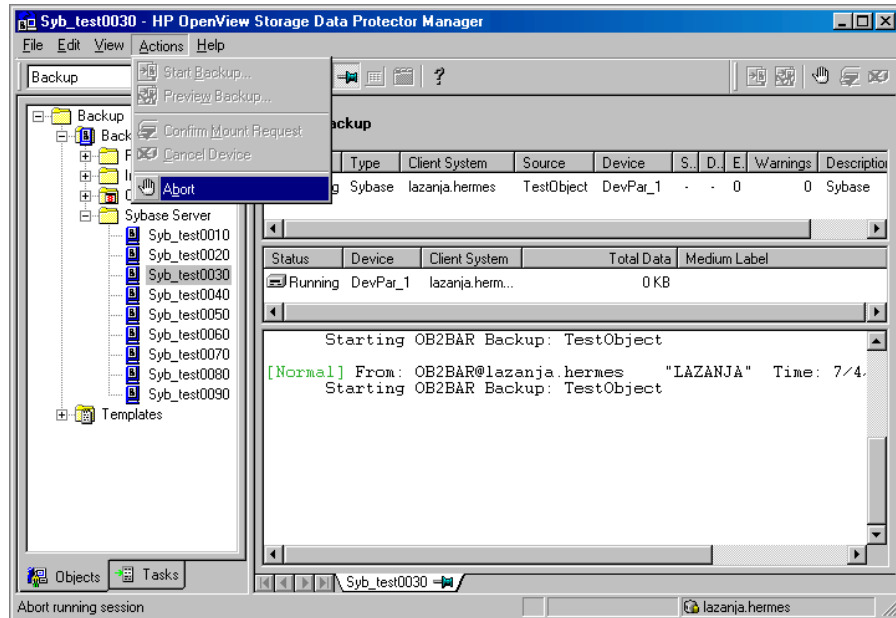
Messages from the Data Protector backup session and messages generated by Sybase are logged to the IDB. Upon successful completion of the backup, the “Session completed successfully” message is displayed in the Session Information dialog box.

Aborting a Running Session

In the Actions menu, click **Abort**, to abort a running Sybase backup session, and then confirm the action.

In the example shown in Figure 1-25, the backup session of the backup specification `FullSybase` is being aborted.

Figure 1-25 **Aborting a Sybase Backup Session**



Scheduling an Existing Backup Specification

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

Data Protector allows you to run unattended backups at specific times or periodically. The powerful Data Protector Scheduler can highly influence the effectiveness and performance of your backup.

To schedule a new Sybase backup specification, follow the steps described in “Creating a Data Protector Sybase Backup Specification” on page 31.

To schedule an existing backup specification, perform the following steps in the HP OpenView Storage Data Protector Manager:

Scheduling Procedure

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click Sybase Server.

A list of backup objects is displayed in the Results Area.

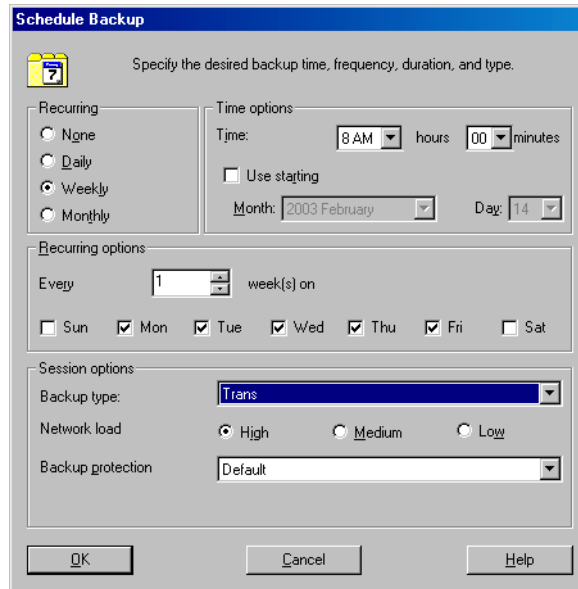
3. Double-click the backup specification you want to schedule and click the Schedule tab to open the Schedule property page.
4. In the Schedule property page, select a date in the calendar click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options. See Figure 1-26 on page 49.
6. Click OK to return to the Schedule property page.
7. Click Apply to save the changes.

**Scheduling
Example**

To schedule a backup specification called *FullSybase* so as to back up transaction logs at 8.00 a.m., and then at 1.00 p.m. and at 6.00 p.m. during week days, open the Schedule property page of the backup specification as described in the above procedure, and then proceed as follows:

1. In the Schedule property page, click Add to open the Schedule Backup dialog box.
2. Under Recurring, select Weekly. Under Time options, select the time 8 AM. Under Recurring Options, select Mon, Tue, Wed, Thu, and Fri. Under Session options, select the Trans backup type. Click OK. See Figure 1-26 on page 49.

Figure 1-26 Scheduling the FullSybase Backup Specification



3. Repeat steps 1 and 2 to schedule another backup. Specify options as described, except the time, which should be set to 1 PM.
4. Repeat steps 1 and 2 to schedule another backup. Specify options as described, except the time, which should be set to 6 PM.
5. Click Apply to save the changes.

After scheduling your backup, you can have it run unattended or you can still run it interactively, as shown in the next section.

Refer to the online Help or the *HP OpenView Storage Data Protector Administrator's Guide* for scheduling details.

NOTE

When creating a Sybase backup specification, you access the Data Protector Scheduler through the Backup Wizard. See “Creating a Data Protector Sybase Backup Specification” on page 31 for information about accessing the Backup Wizard.

Running an Interactive Backup

Interactive backups, as opposed to unattended scheduled backups, are run on demand. They are useful to test your scheduled backups, in case of failure of scheduled backups and to back up clients that need to be backed up urgently, before the regular scheduled periodic backup. You can run your interactive backups using the Data Protector GUI or the Data Protector CLI.

Using the Data Protector GUI

To start an interactive backup of a Sybase database, perform the following steps in the HP OpenView Storage Data Protector Manager:

Running an Interactive Backup

1. In the `Context List`, select `Backup`.
2. In the `Scoping Pane`, expand `Backup`, then `Backup Specifications`, and then `Sybase Server`.
3. Select the backup specification you want to back up and click `Start Backup` in the `Actions` menu.

The `Start Backup` dialog box is displayed.

Select the backup type {`Full`|`Trans`} and network load {`High`|`Medium`|`Low`}. For a description of these options, press **F1**.

TIP

You can also start a backup by right-clicking the Sybase backup specification you want to back up and then clicking `Start Backup`.

4. Click `OK`.

Observe the generated messages. the “Session completed successfully” message displayed at the end of a successful backup session of the `FullSybase` backup specification.

Using the Data Protector CLI

You can also start an interactive backup of a Sybase database using the `omnib` command located in the `/opt/omni/bin/` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory from any client in the Data Protector cell.

The syntax of the omnib command is as follows:

```
omnib -sybase_list <backup_specification_name>
                    [-barmode <SybaseMode>]
                    [<List_options>]
```

where:

- *<SybaseMode>*={full|trans}

A Sybase backup can be either of the following types:

Table 1-4

Sybase Backup Types

Type	Description
Full	The backup of selected databases and transaction logs.
Transaction	The backup of transaction logs that have been modified since the last backup, providing a record of any changes made since the last full or transaction backup.

- *<List_options>* can be any of the following:

```
protect {none | weeks <n> | days <n> | until <date> |
permanent}
```

This option enables you to set the period of protection for the data you back up to prevent the backup media from being overwritten for the specified period. The default is permanent.

```
load {low | medium| high}
```

This option enables you to set the network load during your backup. Set it to high for maximum performance and to low to reduce network load at busy times. The default is high.

```
crc
```

Set this option on to have Data Protector calculate the cycle redundancy check when a backup is run. This option enables you to later confirm using the Verify option whether data has been correctly written to the medium. The default is off.

Integrating Sybase and Data Protector

Backing Up a Sybase Database

`no_monitor`

By default, the command monitors the session and displays the status of the session.

`test_bar`

Tests the backup specification as described in “Testing the Integration” on page 42.

The following are some common backup examples:

Example 1

To start a full backup of the Sybase backup specification called FullSybase, execute the following command in the `/opt/omni/bin/` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory:

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

You can observe backup messages in the Data Protector Monitor.

Example 2

To start a transaction backup of a Sybase backup specification called TransSybase, execute the following command in the `/opt/omni/bin/` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory:

```
omnib -sybase_list TransSybase -barmode trans
```

Backing Up Using Sybase Commands

To start a backup of a database from the client where the database is located, using the Sybase `isql` command interface, proceed as follows:

- Backup Procedure**
1. Check if the devices used for the backup contain formatted (Windows systems) or initialized (UNIX systems) media with sufficient free space.
 2. Verify the backup options of the Data Protector Sybase backup specification.
 3. Log into the Sybase Server as a Sybase SQL Server Administrator and run the following command in the Sybase Backup Server home directory:

Windows Example

```
bin\isql -U<SA> -S<SYBASESERVERNAME> -P<SA_PASSWORD>  
dump database <TARGET_DATABASE> to "ob2syb:::<SYBASELISTNAME>"
```


UNIX Example

```
bin/isql -U<SA> -S<SYBASESERVERNAME> -P<SA_PASSWORD>  
dump database <TARGET_DATABASE> to "ob2syb::<SYBASELISTNAME>"
```

where,

<SA> is the Sybase user.

<SYBASESERVERNAME> is the name of Sybase SQL Server.

<SA_PASSWORD> is the password of the Sybase System Administrator, for example, sa.

<TARGET_DATABASE> is the name of the Sybase database that will be backed up, for example, database2.

<SYBASELISTNAME> is the name of the Data Protector Sybase backup specification, for example, FullSybase.

Restoring a Sybase Database

Restoring of a Sybase database consists of the following steps:

- Restore Procedure**
1. Restoring a full backup of the Sybase database.
 2. Restoring subsequent transaction backups, if they exist.

To restore a corrupted database, you need to find the right media and the sessionID of the last backup session with a full backup. If you have backed up a database with several streams, you need to know the number of streams. This information can be found using the Data Protector `omnidb` command. See “Finding Information Needed for Restore” on page 54 for more information. You can also use the Data Protector `syb_tool` command to create an `isql load` command that you then use to restore a database on a specified date. See “The Data Protector `syb_tool` Command” on page 65 for more information. Note that this tool is not used to restore your data, but just to return `load` commands that you then use for restore.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before starting a restore procedure from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

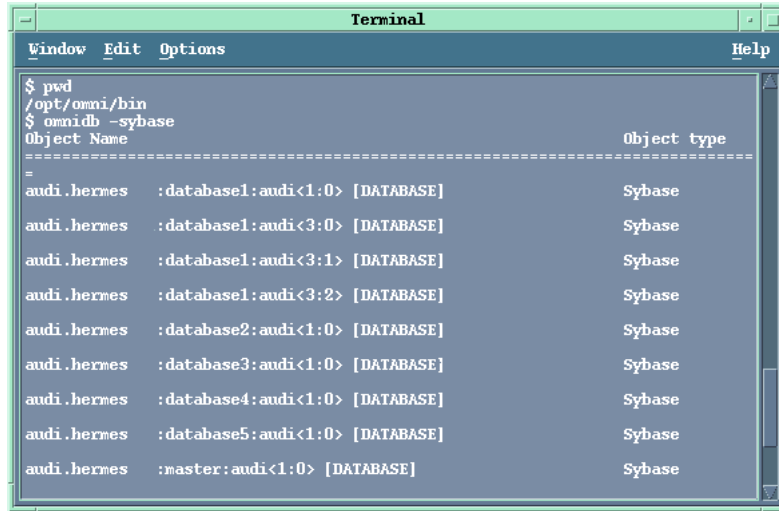
- On UNIX: `export OB2BARHOSTNAME=<virtual_hostname>`
- On Windows: `set OB2BARHOSTNAME=<virtual_hostname>`

Finding Information Needed for Restore

To find the information needed to restore your data, execute the following commands in the `/opt/omni/bin/` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory:

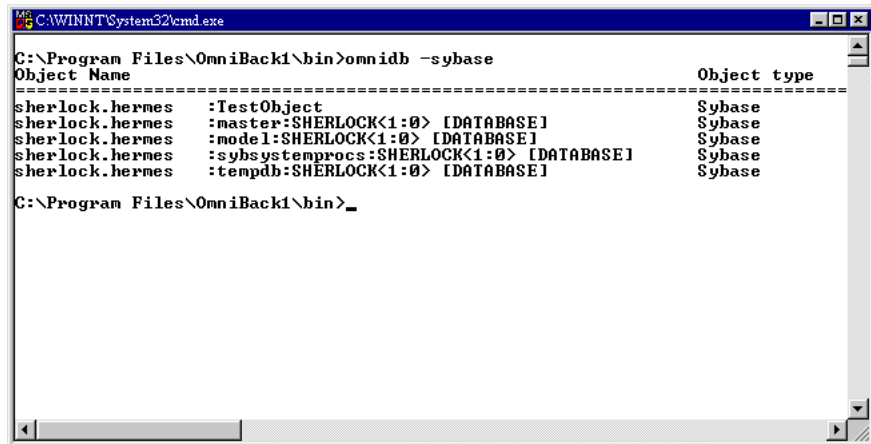
1. `omnidb -sybase`
to get a list of Sybase objects.

Figure 1-27 List of Sybase Objects (UNIX Example)



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

Figure 1-28 List of Sybase Objects (Windows Example)

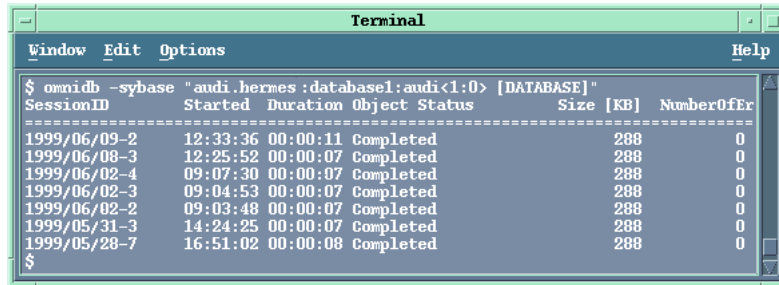


```
Ms-DOS C:\WINNT\System32\cmd.exe
C:\Program Files\OmniBack1\bin>omnidb -sybase
Object Name                                     Object type
-----
sherlock.hermes :TestObject                      Sybase
sherlock.hermes :master:SHERLOCK<1:0> [DATABASE]  Sybase
sherlock.hermes :model:SHERLOCK<1:0> [DATABASE]  Sybase
sherlock.hermes :sybsystemprocs:SHERLOCK<1:0> [DATABASE] Sybase
sherlock.hermes :tempdb:SHERLOCK<1:0> [DATABASE] Sybase
C:\Program Files\OmniBack1\bin>_
```

2. omnidb -sybase "<object_name>"

to get details on a specific object, including the SessionID of the backup session. In case of object copies, do not use the copy session ID for the restore, but the object's backup ID, which equals the object's backup session ID. Figure 1-29 shows how you get details about the object called `audi.hermes:database1:audi<1:0> [DATABASE]`.

Figure 1-29 Details about a Specific Session



3. `omnidb -session <SessionID> -media`

to display media needed for restore. In Figure 1-30 and in Figure 1-31, media used for session 1999/06/09-2 (UNIX example) or 1999/06/30-8 (Windows example) are displayed.

Figure 1-30 Finding Media Needed for Restore (UNIX Example)

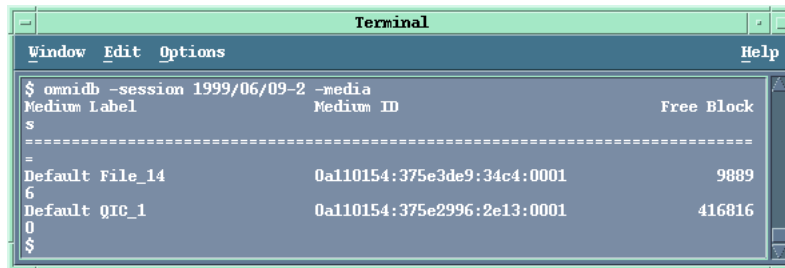
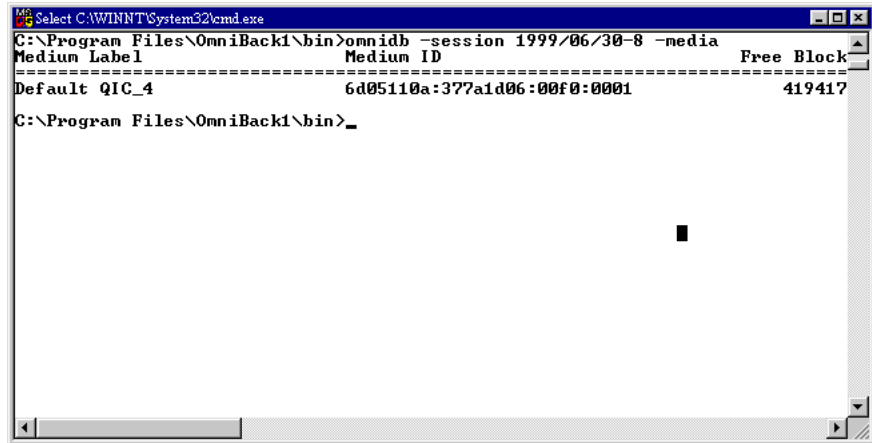


Figure 1-31 Finding Media Needed for Restore (Windows Example)



For detailed information on the omnidb command, refer to the omnidb man page.

Restore

A Sybase restore can only be started from a Sybase Server by using the isql command. To run the isql command, proceed as follows:

1. On UNIX, log on to your Sybase SQL Server as user sybase.
2. Type in the following command in the Sybase SQL Server home directory:

UNIX

```
bin/isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>
```

Windows

```
isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>
```

3. In the first line, type in the appropriate load command. To execute your command(s), type go in the last line and press Enter. See “Restore Examples” on page 60 for more information on running the load command.

Before you restore, you want to find out the databases on Sybase SQL Server. Execute the util_sybase command to list Sybase databases on the defined Sybase SQL Server, as shown in Figure 1-32 and Figure 1-33.

Cluster-Aware Clients

In a cluster environment, the environment variable OB2BARHOSTNAME must be defined as the virtual hostname before running the util_sybase command from the command line (on the client). The OB2BARHOSTNAME variable is set as follows:

- On UNIX: export OB2BARHOSTNAME=<virtual_hostname>
- On Windows: set OB2BARHOSTNAME=<virtual_hostname>

Running the util_sybase Command

The util_sybase command is run as follows:

- On UNIX: /opt/omni/lbin/util_sybase -OBS0 <SYBASESERVERNAME>
- On windows: <Data_Protector_home>\bin\util_sybase.exe -OBS0 <SYBASESERVERNAME>

Windows

Figure 1-32

List Sybase Database Names (UNIX Example)

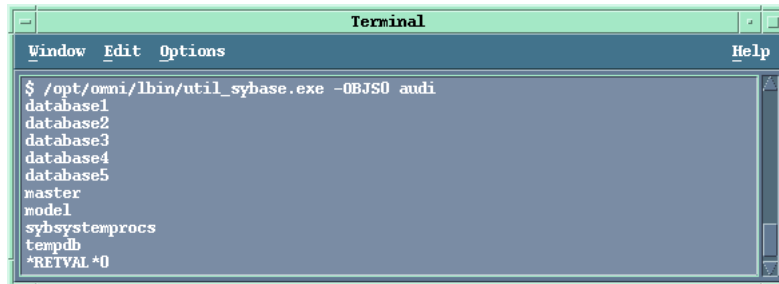
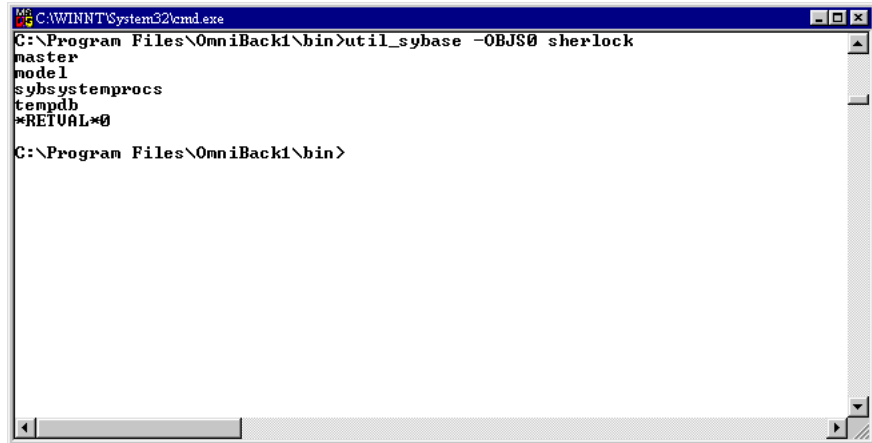


Figure 1-33 List Sybase Database Names (Windows Example)



The Sybase load database and load trans commands are covered in detail in the *Sybase SQL Server System Administration Guide*. In this guide, only a brief description of these commands will be given. Restore examples will also be provided.

The syntax of the Sybase load command is as follows:

```
load {database|transaction} <new_db_name>
from
"ob2syb::[::[::[::[::
```

database|transaction defines the backup of databases or transaction logs

ob2syb is the Data Protector Database Library

<version> can either be the SessionID of the backup session with the data you want to restore or the latest version keyword to restore the latest version of backup

<new_db_name> is the name of the new database to be restored

Restoring a Sybase Database

`<old_db_name>` is the name of the original database

`<old_db_server_name>` is the name of the original Sybase SQL Server.

Restore Examples

The following are examples of using the Sybase `load` command for restore. Before restoring, you need information, which you can find using the Data Protector `omnidb` command, as described in “Finding Information Needed for Restore” on page 54.

Run the `omnidb -sybase` command to get a list of Sybase objects and the `omnidb -sybase "Object_Name"` command to get details about the backed up object.

To run the `load` command, first start the Sybase `isql` command as described in “Restore” on page 57.

Example 1

To restore a database named `database2` (UNIX example) or a database named `sybssystemprocs` (Windows example), backed up in a session with `sessionID 1999/06/09-2` (UNIX example) or with `sessionID 1999/08/10-4` (Windows example), start the Sybase `isql` command and then execute the following command, also shown in Figure 1-34 and Figure 1-35:

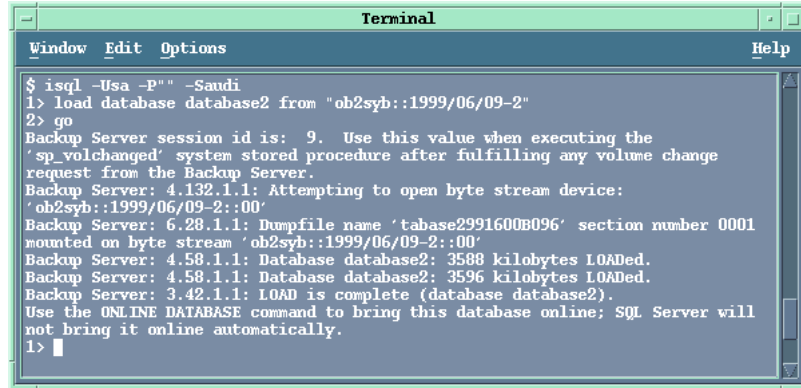
UNIX Example

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

Windows Example

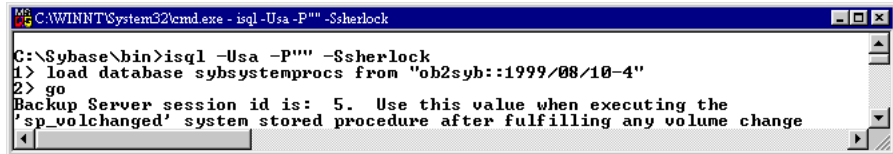
```
1>load database sybssystemprocs from "ob2syb::1999/08/10-4"  
2>go
```


Figure 1-34 Restoring database2, Backed Up in Session 1999/06/09-2 (UNIX Example)



```
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> █
```

Figure 1-35 Restoring sybssystemprocs, Backed Up in Session 1999/08/10-4 (Windows Example)



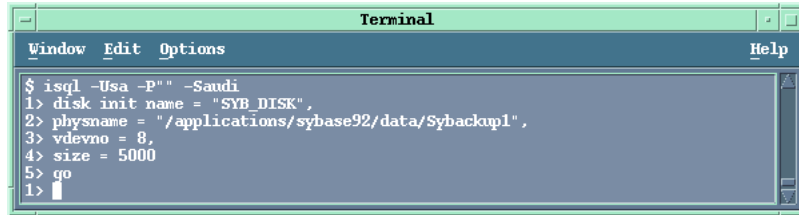
```
C:\WINNT\System32\cmd.exe - isql -Usa -P"" -Ssherlock
C:\Sybase\bin>isql -Usa -P"" -Ssherlock
1> load database sybssystemprocs from "ob2syb::1999/08/10-4"
2> go
Backup Server session id is: 5. Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
```

Example 2 To restore a database to a new database, first create an empty database, and then perform the restore.

To create an empty database with a defined layout, proceed as follows:

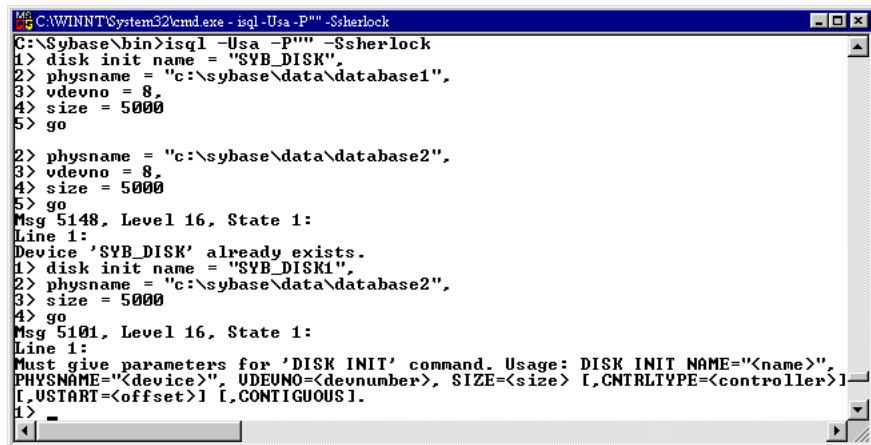
1. Create a “database device” as shown in Figure 1-36 and Figure 1-37:

Figure 1-36 **Creating a Database Device (UNIX Example)**



```
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>
```

Figure 1-37 **Creating a Database Device (Windows Example)**



```
C:\WINNT\System32\cmd.exe - isql -Usa -P"" -Ssherlock
C:\Sybase\bin>isql -Usa -P"" -Ssherlock
1> disk init name = "SYB_DISK",
2> physname = "c:\sybase\data\database1",
3> vdevno = 8,
4> size = 5000
5> go

2> physname = "c:\sybase\data\database2",
3> vdevno = 8,
4> size = 5000
5> go
Msg 5148, Level 16, State 1:
Line 1:
Device 'SYB_DISK' already exists.
1> disk init name = "SYB_DISK1",
2> physname = "c:\sybase\data\database2",
3> size = 5000
4> go
Msg 5101, Level 16, State 1:
Line 1:
Must give parameters for 'DISK INIT' command. Usage: DISK INIT NAME=<name>,
PHYSNAME=<device>, UDEVNO=<devnumber>, SIZE=<size> [,CNTRLTYPE=<controller>]
[,USTART=<offset>] [,CONTIGUOUS].
1>
```

2. Create the empty database using the create database command, as shown in Figure 1-38 and Figure 1-39. It should have the same layout as an existing database from which you then want to restore.

NOTE

A new database cannot be smaller than the model database.

Figure 1-38 Creating an Empty Database (UNIX Example)


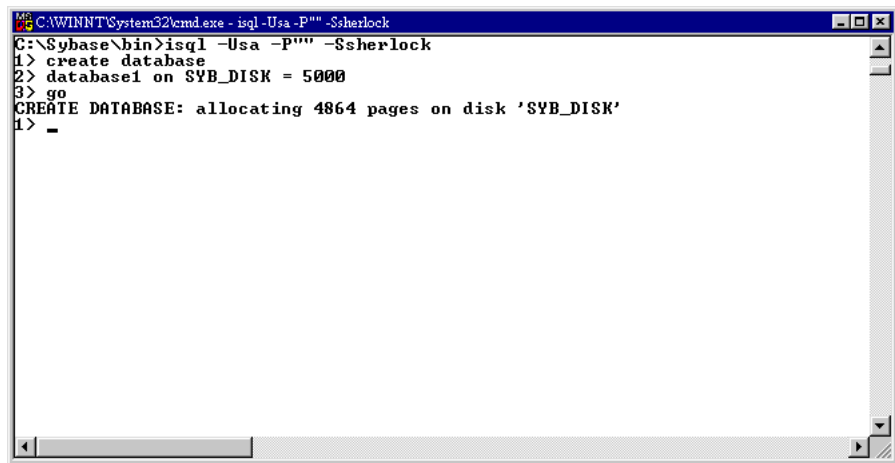


Figure 1-39 Creating an Empty Database (Windows Example)



3. To finally restore a database in a new database, proceed as in the following example:

Windows Example

A database named databas1 was created with the same layout as the database named model, which had backed up before.

To restore model into databas1, start the Sybase isql utility and execute the following commands:

```
1>load database databas1 from "ob2syb::latest version::model"
2>go
```

UNIX Example

In this example, a database named Sybdata1 was created with the same layout as the database named Sybdata, which had backed up before.

Integrating Sybase and Data Protector

Restoring a Sybase Database

To restore Sybdata into Sybdata1, start the Sybase isql utility and execute the following commands:

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

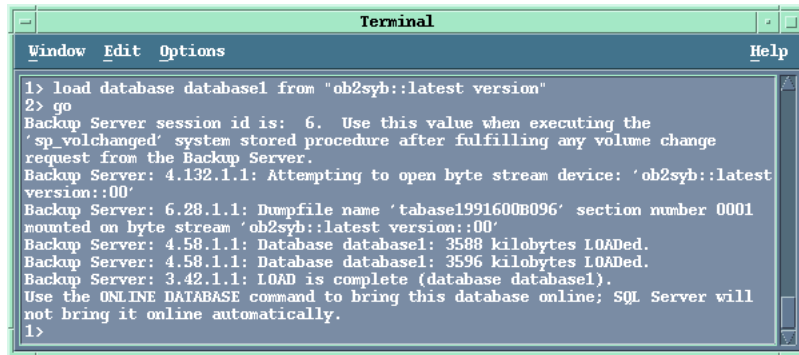
Example 3

To restore the latest version of a database named database1, start the Sybase isql utility and execute the following command:

```
1>load database database1 from "ob2syb::latest version"  
2>go
```

Figure 1-40

Restoring the Latest Version of database1



```
Terminal  
Window Edit Options Help  
1> load database database1 from "ob2syb::latest version"  
2> go  
Backup Server session id is: 6. 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::latest  
version::00'  
Backup Server: 6.28.1.1: Dumpfile name 'tabase1991600B096' section number 0001  
mounted on byte stream 'ob2syb::latest version::00'  
Backup Server: 4.58.1.1: Database database1: 3588 kilobytes LOAded.  
Backup Server: 4.58.1.1: Database database1: 3596 kilobytes LOAded.  
Backup Server: 3.42.1.1: LOAD is complete (database database1).  
Use the ONLINE DATABASE command to bring this database online; SQL Server will  
not bring it online automatically.  
1>
```

Example 4

To restore a database backed up with several streams, add the appropriate number of stripe commands. You can get the number of streams in the Data Protector Monitor.

For example, to restore the latest version of a database named database3, backed up with three streams, start the Sybase isql utility and execute the following commands:

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

The Data Protector `syb_tool` Command

The Data Protector `syb_tool` command, located in the `<Data_Protector_home>\bin` (Windows systems), `/opt/omni/bin` (UNIX systems) directory, creates an `isql` load command that is used to restore a database on a specified date.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `syb_tool` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

- On UNIX: `export OB2BARHOSTNAME=<virtual_hostname>`
- On Windows: `set OB2BARHOSTNAME=<virtual_hostname>`

To start the `syb_tool` command, log in as either a Data Protector administrator or as the Sybase user. Note that this tool is not used to restore your data, but to return `load` commands that you then need to use for restore. The command has the following syntax:

The `syb_tool` Command Syntax

```
syb_tool <dbname> <servername>  
-date <YYYY/MM/DD.hh:mm:ss>  
[ -new_db <dbname> ]  
[ -new_server <servername> ]  
[ -file <filename> ]  
[ -media      ]
```

where:

- | | |
|---------------------------------|--|
| <code><dbname></code> | is the name of the Sybase database to be restored (required) |
| <code><servername></code> | is the hostname of the Sybase SQL Server system; in cluster environment, specify the virtual hostname |
| <code><date></code> | represents a date after which the first new backup version of the database will be restored (required) |
| <code><new_db></code> | is the destination database name (optional, used for renaming the database to be restored) |
| <code><new_server></code> | is the destination Sybase SQL Server name (optional, used for changing Sybase SQL Server) |
| <code><file></code> | is the file containing an <code>isql</code> command or command sequence that should be used to restore the specified data (optional) |

Integrating Sybase and Data Protector

Restoring a Sybase Database

`media` lists the media necessary to perform the restore (optional)

A global options file variable `OB2SybaseTransLogDelay` is used to define the time between the points when the transaction logs are closed and the backup session is started. The default value is 20 seconds.

Example 1

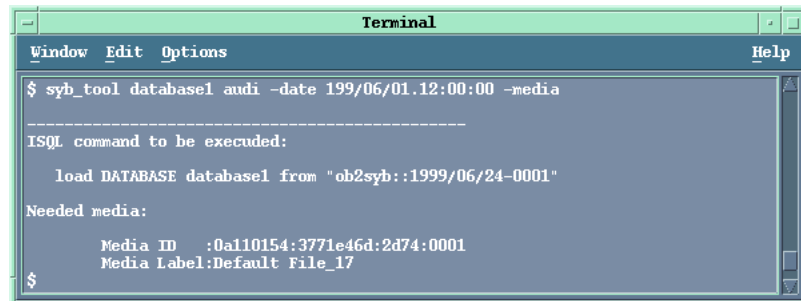
To return an `isql` command to restore the first backup of the database `database1` on Sybase Server `audi` that was made after 12:00 June 1, 1999 and also the media on which the backup was made, type in the following command in the `/opt/omni/bin` (UNIX systems) or in the `<Data_Protector_home>\bin` (Windows systems) directory:

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

The required `isql` command sequence as well as the media that you need for the restore are returned, as shown in Figure 1-41.

Figure 1-41

The load Command Including the Required Media



```
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 return the results of the above command to a file, `/tmp/isqlfile` (Windows systems) or to `c:\tmp\isqlfile` (UNIX systems), type in the following command:

Windows Example

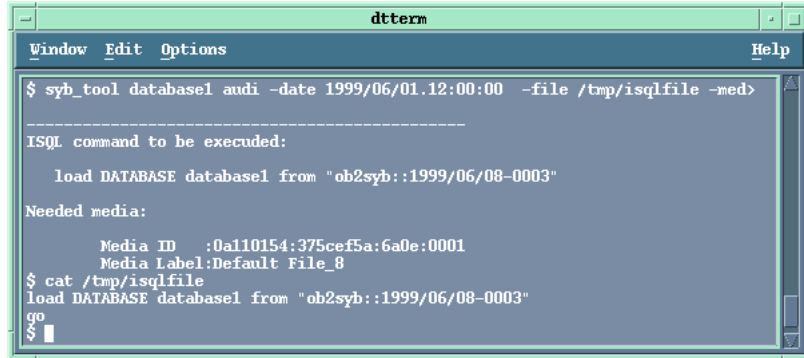
```
syb_tool database1 sherlock -date 1999/06/01.12:00:00 -file
c:\tmp\isqlfile -media
```

UNIX Example

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

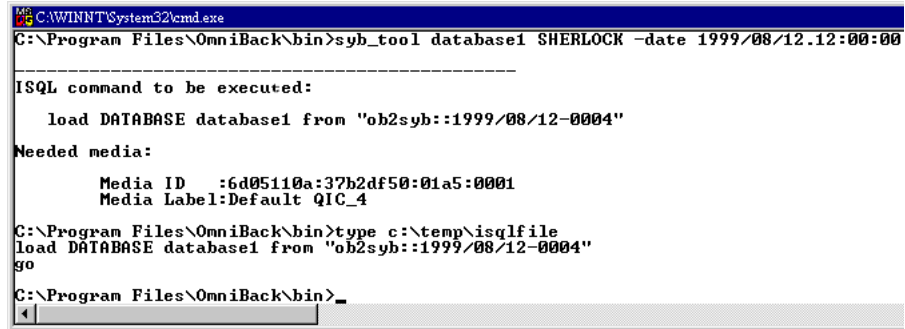
The required `isql` command sequence, the media that you need for the restore as well as the file to which the command sequence was loaded, are returned, as shown in Figure 1-42 and Figure 1-43.

Figure 1-42 The load Command Including the Required Media to a File (UNIX Example)



```
dtterm
Window Edit Options Help
$ syb_tool database1 audi -date 1999/06/01.12:00:00 -file /tmp/isqlfile -med>
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/06/08-0003"
Needed media:
    Media ID      :0a110154:375cef5a:6a0e:0001
    Media Label:Default File_8
$ cat /tmp/isqlfile
load DATABASE database1 from "ob2syb::1999/06/08-0003"
go
$
```

Figure 1-43 The load Command Including the Required Media to a File (Windows Example)



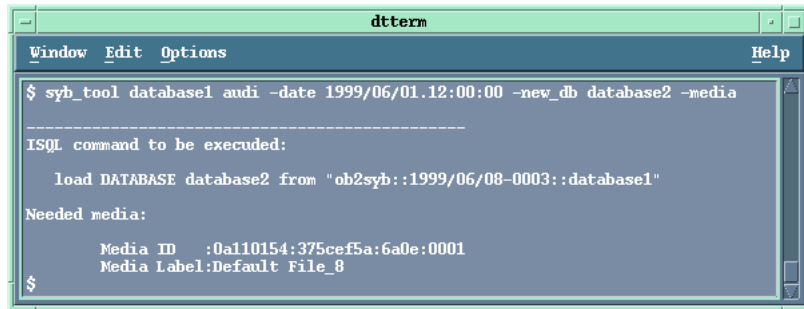
```
C:\Program Files\OmniBack\bin>syb_tool database1 SHERLOCK -date 1999/08/12.12:00:00
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/08/12-0004"
Needed media:
    Media ID      :6d05110a:37b2df50:01a5:0001
    Media Label:Default QIC_4
C:\Program Files\OmniBack\bin>type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

Example 3 To return the load command that restores a database database1 to database2, perform the following command:

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

The required isql command sequence as well as the media that you need for the restore are returned, as shown in Figure 1-44.

Figure 1-44 **The load Command to a Different Database**



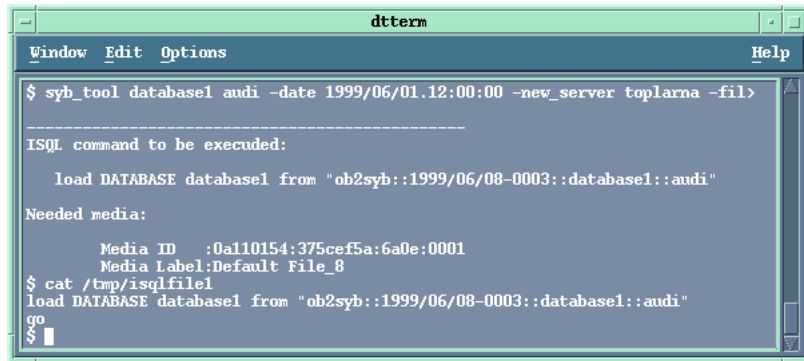
Example 4 To return the load command that restores a database database1 backed up using the server audi (UNIX example) sherlock (Windows example) to server toplarna, perform the following command:

Windows Example `syb_tool database1 sherlock -date 1999/06/01.12:00:00 -new_server toplarna -file c:\tmp\isql -media`

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

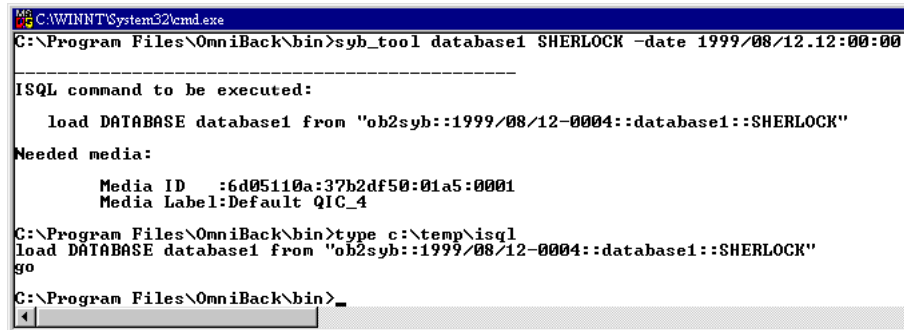
The required isql command sequence, the media that you need for the restore as well as the file to which the command sequence was loaded, are returned, as shown in Figure 1-45 and Figure 1-46.

Figure 1-45 The load Command to a Different Server (UNIX Example)



```
dtterm
Window Edit Options Help
$ syb_tool database1 audi -date 1999/06/01.12:00:00 -new_server toplama -fil>
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/06/08-0003::database1::audi"
Needed media:
    Media ID   :0a110154:375cef5a:6a0e:0001
    Media Label:Default File_8
$ cat /tmp/isqlfile1
load DATABASE database1 from "ob2syb::1999/06/08-0003::database1::audi"
go
$
$
```

Figure 1-46 The load Command to a Different Server (Windows Example)



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

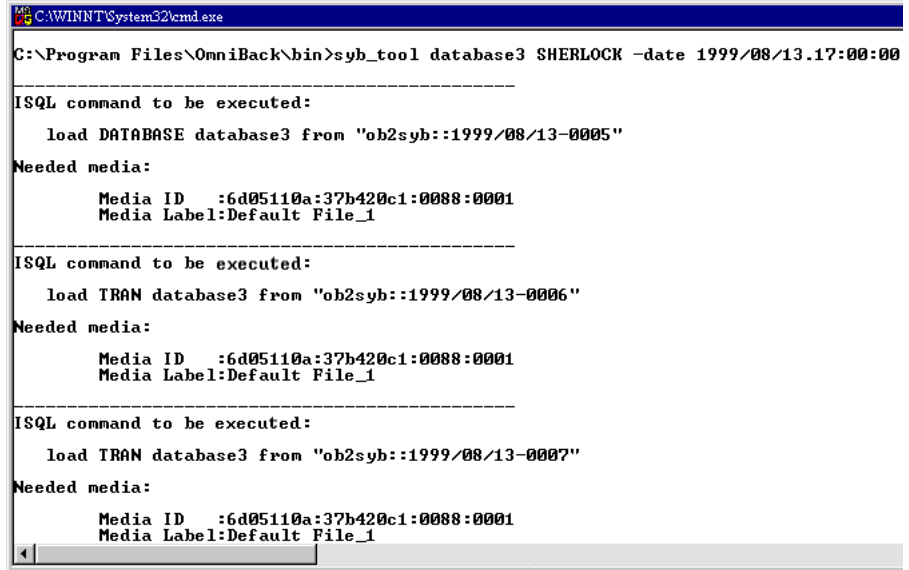
**Example 5
(Windows)**

To return the load command that restores a full backup and two transaction backups, backed up using the server, sherlock, including the required media, perform the following command:

```
syb_tool database3 sherlock -date 1999/08/13.17:00:00 -media
```

The required isql command sequence and the file to which the command sequence was loaded, are returned, as shown in Figure 1-47.

Figure 1-47 Loading Transaction Logs



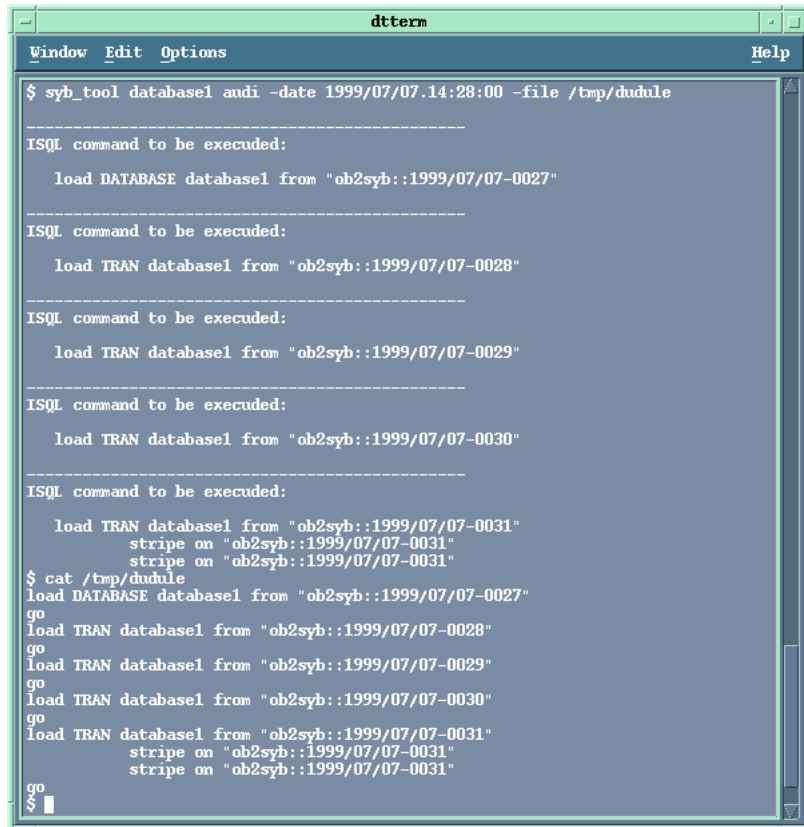
```
C:\WINNT\System32\cmd.exe
C:\Program Files\OmniBack\bin>syb_tool database3 SHERLOCK -date 1999/08/13.17:00:00
-----
ISQL command to be executed:
    load DATABASE database3 from "ob2syb::1999/08/13-0005"
Needed media:
    Media ID   :6d05110a:37b420c1:0088:0001
    Media Label:Default File_1
-----
ISQL command to be executed:
    load TRAN database3 from "ob2syb::1999/08/13-0006"
Needed media:
    Media ID   :6d05110a:37b420c1:0088:0001
    Media Label:Default File_1
-----
ISQL command to be executed:
    load TRAN database3 from "ob2syb::1999/08/13-0007"
Needed media:
    Media ID   :6d05110a:37b420c1:0088:0001
    Media Label:Default File_1
```

Example 5 (UNIX) To return the load command that restores a full backup and three transaction backups with concurrency one and a transaction backup with concurrency 3 backed up using the server, audi, to a file, /tmp/dudule perform the following command:

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

The required isql command sequence and the file to which the command sequence was loaded, are returned, as shown in Figure 1-47.

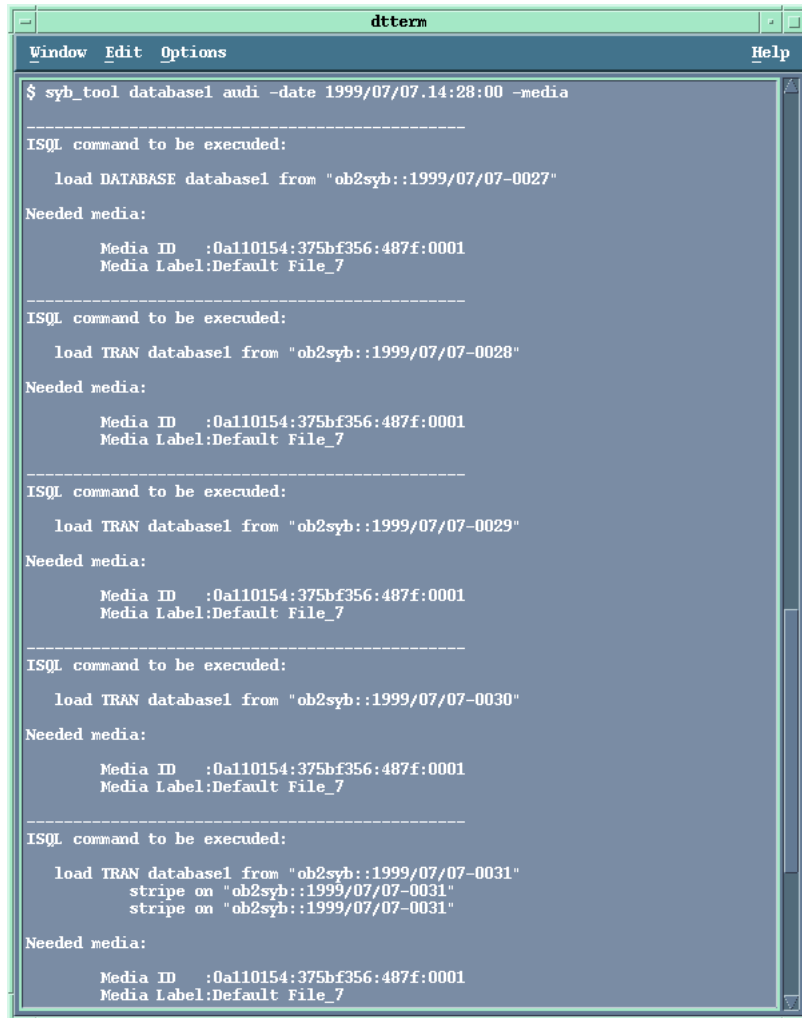
Figure 1-48 Loading Transaction Logs from Multiple Devices



```
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
$
```

The required media for the example are shown in Figure 1-49.

Figure 1-49 Required Media



```
dtterm
Window Edit Options Help
$ syb_tool database1 audi -date 1999/07/07.14:28:00 -media
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/07/07-0027"
Needed media:
    Media ID      :0a110154:375bf356:487f:0001
    Media Label:Default File_7
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0028"
Needed media:
    Media ID      :0a110154:375bf356:487f:0001
    Media Label:Default File_7
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0029"
Needed media:
    Media ID      :0a110154:375bf356:487f:0001
    Media Label:Default File_7
-----
ISQL command to be executed:
    load TRAN database1 from "ob2syb::1999/07/07-0030"
Needed media:
    Media ID      :0a110154:375bf356:487f:0001
    Media Label:Default File_7
-----
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"
Needed media:
    Media ID      :0a110154:375bf356:487f:0001
    Media Label:Default File_7
```

Restoring Using Another Device

Data Protector supports restore using a device other than the one that was used at backup time.

Specify these devices in the

<Data_Protector_home>\Config\server\Cell\restoredev
(Windows systems) or in the /etc/opt/omni/server/cell/restoredev
(UNIX systems) file in the following format:

```
"DEV 1" "DEV 2"
```

where,

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

Note that this file should be deleted after it is used.

On Windows, the file has to be in the UNICODE format.

Example

Suppose you have Sybase objects backed up on a device called DAT1. To restore them from a device named DAT2, specify the following in the restoredev file:

```
"DAT1" "DAT2"
```

Disaster Recovery

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

Check the instructions of the database/application vendor on how to prepare for the disaster recovery. Also refer to the Disaster Recovery chapter in the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on how to approach system disaster recovery using Data Protector.

This is a general procedure on how to recover an application:

1. Complete recovery of the operating system.
2. Installing, configuring, and initializing the database/application so that data on the Data Protector media can be loaded back to the system. Consult the documentation of the database/application vendor for a detailed procedure and steps needed to prepare the database.

Restoring a Sybase Database

3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and the procedures in the troubleshooting section.
4. Start the restore. When the restore is complete, follow the instructions of the database/application vendor for any additional steps required to bring the database back online.

Monitoring a Sybase Backup and Restore Session

Data Protector enables you to monitor currently running and view previous backup and restore sessions. When you run an interactive backup or restore session, a monitor window displays showing you the progress of the session. You can monitor the session from any Data Protector client in the network that has the Data Protector User Interface component installed. Note, however, that the session continues even with the User Interface closed.

Monitoring Current Sessions

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

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

All actions and messages of a session are logged to both the Data Protector and Sybase log files. Mount requests are displayed on the Data Protector monitor.

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

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

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

NOTE

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

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

Viewing Previous Sessions

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

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

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

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

Sybase Character Sets

Sybase SQL server supports various language environments. Refer to the *Sybase SQL Server Utility Programs* for more information.

To enable this support in the Data Protector Sybase integration, you need to specify the environmental variable `OB2_ISQL_OPTS` in the Data Protector Sybase configuration file as follows:

```
util_cmd -putopt[ion] Sybase <instance_name> OB2_ISQL_OPTS \  
="-J<char_set>" -sublist Environment
```

where `<char_set>` is the character set to be used. For example:

```
util_cmd -putopt[ion] Sybase <instance_name> OB2_ISQL_OPTS \  
="-Jsjis" -sublist Environment
```

The above command will enable the Sybase Japanese language environment support.

What Happens?

The environmental variable `OB2_ISQL_OPTS` is added to the `Environment` `sublist` of the Data Protector Sybase configuration file. For more information on Data Protector Sybase configuration file, see “Data Protector Sybase Configuration File” on page 9.

Consequently, every time Data Protector starts the `isql` command, it is started with the `-Jsjis` option.

Troubleshooting

This section describes procedures you should follow to troubleshoot your configuration, back up, or restore problems.

Before You Begin

1. Ensure that the latest official Data Protector patches are installed. Refer to “Verifying Which Data Protector Patches Are Installed” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or

http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

2. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a description of general Data Protector limitations as well as known problems and workarounds.

Follow the given procedures to troubleshoot your configuration, backup, or restore problems, respectively.

Troubleshooting on Windows Systems

Cluster-Related Troubleshooting

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before performing some procedures from the command line (on the client). When the GUI is used, this is not required. The `OB2BARHOSTNAME` variable is set as follows:

```
set OB2BARHOSTNAME=<virtual_hostname>
```

Configuration Problems

If you have problems configuring the Data Protector Sybase integration, proceed as follows:

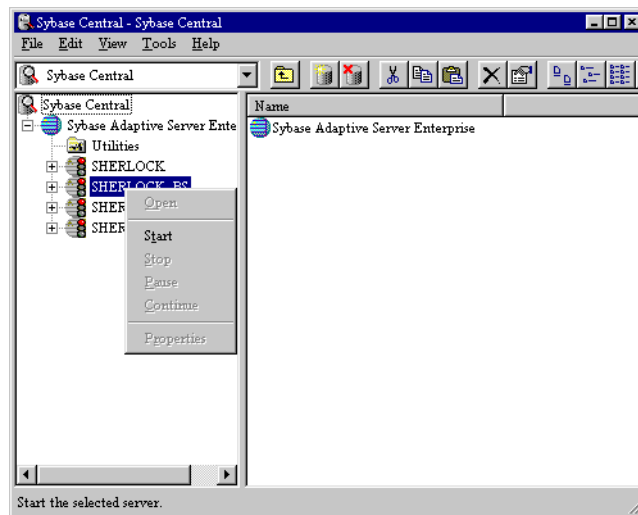
1. Make a Data Protector filesystem backup of the problematic client.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.

2. Ensure that Sybase SQL Server and the Sybase Backup Server are up and running.

To start a Sybase server from Sybase Central, without connecting to it, right-click the server, and then click *Start* in the pop-up menu, as shown in the Figure 1-50.

Figure 1-50 Starting Sybase Servers Using Sybase Central on Windows Systems



3. Examine system errors reported in the `<Data_Protector_home>\log\debug.log` file on the Sybase server.
4. If you have any non-default Sybase settings, ensure that they are registered in the System Properties dialog box, which you access by selecting System in the Control Panel.

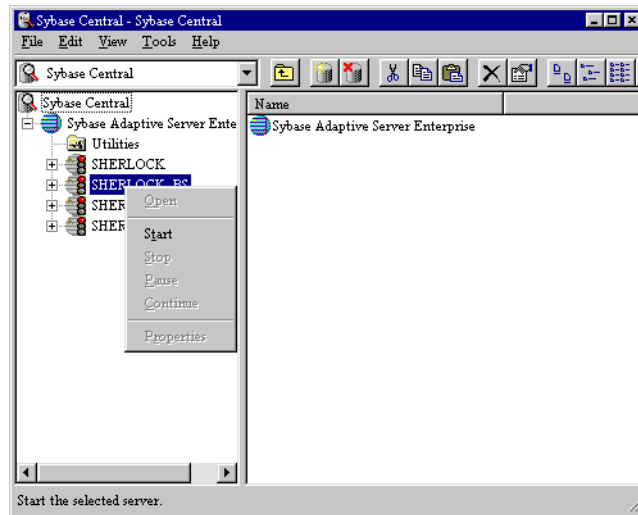
Backup Problems

If you have problems backing up Sybase databases, proceed as follows:

1. Make a Data Protector filesystem backup of the problematic client.
Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.
2. Ensure that Sybase SQL Server and Sybase Backup Server are up and running.

To start a Sybase server from Sybase Central, without connecting to it, right-click the server, and then click **Start** in the pop-up menu:

Figure 1-51 Starting Sybase Servers Using Sybase Central on Windows Systems



3. Verify the configuration of your Sybase server.

Using the Data Protector GUI

You can also check the configuration of your Sybase Server by performing the following steps in the HP OpenView Storage Data Protector Manager:

- a. In the Context List, select Backup.
- b. In the Scoping Pane, expand Backup, then Backup Specifications, and then Sybase Server.

- c. Click a configured Sybase backup specification you want.
The Sybase Server is displayed in the Results Area.
 - d. Right-click the client and then click Check Configuration.
A message is returned confirming that the integration is properly configured.
4. Test the Data Protector Sybase configuration as per instructions in “Testing the Integration” on page 42.

Example

Run the following command in the `<Data_Protector_home>\bin\` directory, to test the configuration of the backup specification called FullSybase:

```
omnib -sybase_list FullSybase -test_bar
```

- If the Data Protector part of the test fails then 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.
Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.
 - If the test succeeds, start a backup directly from a Sybase Server. See “Backing Up Using Sybase Commands” on page 52 for instructions.
If this backup succeeds, then the problem may be that the client on which the Data Protector User Interface runs does not have enough memory, disk space, or other operating system resources.
5. Test Data Protector data transfer using the testbar utility. Proceed as follows: in the `<Data_Protector_home>\bin` directory:

```
testbar  
-type:Sybase  
-appname:<SYBASESERVERNAME>  
-bar:<backup_specification_name>  
-perform:backup
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server and `<backup_specification_name>` the name of the Data Protector backup specification.

If the test is successful, then proceed to the next step, otherwise proceed as follows:

- a. Troubleshoot errors reported by the testbar utility using the Data Protector troubleshooting file,
`<Data_Protector_home>\docs\trouble.txt.`
 - b. Examine system errors reported in the
`<Data_Protector_home>\log\debug.log` file on the Sybase Server.
6. If you have any non-default Sybase settings, ensure that they are registered in the System Properties dialog box, which you access by selecting System in the Control Panel.

Restore Problems

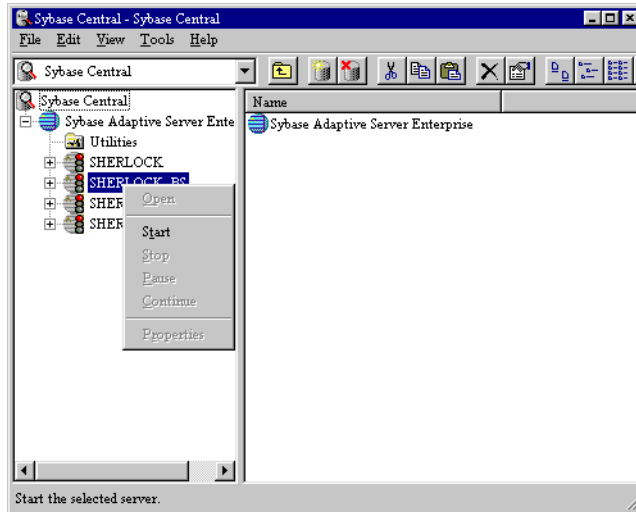
If you have problems restoring Sybase databases, proceed as follows:

1. Make a Data Protector filesystem backup and restore of the problematic client.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.
2. Examine system errors reported in the
`<Data_Protector_home>\log\debug.log` file on the Sybase server.
3. Ensure that Sybase SQL Server and Sybase Backup Server are up and running.

To start a Sybase server from Sybase Central, without connecting to it, right-click the server, and then click **Start** in the pop-up menu:

Figure 1-52 Starting Sybase Servers Using Sybase Central



4. Test Data Protector data transfer using the testbar utility. Proceed as follows:

```
testbar
-type: Sybase
-appname: <SYBASESERVERNAME>
-bar: <backup_specification_name>
-perform: backup
```

where *<SYBASESERVERNAME>* is the name of Sybase SQL Server and *<backup_specification_name>* the name of the Data Protector backup specification.

If the test is successful, then proceed to the next step, otherwise proceed as follows:

- a. Troubleshoot errors reported by the testbar utility using the Data Protector troubleshooting file,
<Data_Protector_home>\docs\trouble.txt.

- b. Examine system errors reported in the
`<Data_Protector_home>\log\debug.log` file on the Sybase
Server.

Troubleshooting on UNIX Systems

Cluster Related Troubleshooting

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before performing some procedures from the command line (on the client). When the GUI is used, this is not required. The `OB2BARHOSTNAME` variable is set as follows:

```
export OB2BARHOSTNAME=<virtual_hostname>
```

Configuration Problems

If you have problems configuring the Data Protector Sybase integration, proceed as follows:

1. Make a Data Protector filesystem backup of the problematic client.
Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.
2. Ensure that Sybase SQL Server is up and running.


The simplest way to find whether Sybase SQL Server is running is to try and log on to the server using the `isql` command:

- a. Log on to Sybase SQL Server as user `sybase`
- b. Type in the following command in the Sybase SQL Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server; `audi` in Figure 1-53, `<PASSWORD>` is the Sybase Administrator password and `<SA>` is the Sybase user.

Figure 1-53 **Checking if Sybase SQL Server Is Running on UNIX Systems**



If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<SYBASESERVERNAME>
```

3. Ensure that Sybase Backup Server is up and running

The simplest way to find whether Sybase Backup Server is running is to try and log on to the server using the `isql` command

- a. Log on to Sybase Backup Server as user `sybase`
- b. Type in the following command in the Sybase Backup Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<BACKUPSERVERNAME>
```

where `<BACKUPSERVERNAME>` is the name of Sybase Backup Server; `audi_back` in Figure 1-54, `<PASSWORD>` is the Sybase Administrator password and `<SA>` is the Sybase user.

Figure 1-54 **Checking if Sybase Backup Server Is Running on UNIX Systems**



If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<BACKUPSERVERNAME>
```

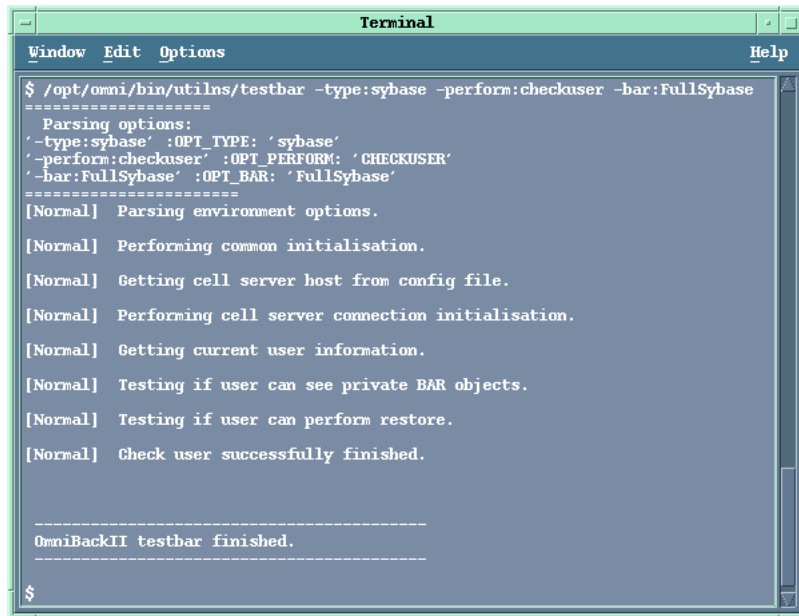
4. Examine system errors reported in the `/var/opt/omni/log/debug.log` file on the Sybase server.

5. If you have any non-default Sybase settings, ensure that they are registered in the Data Protector Sybase configuration file. For more information on Data Protector Sybase configuration file, see “Data Protector Sybase Configuration File” on page 9.
6. Test if the Sybase user has the right privileges in Data Protector. Log in as the Sybase user, for example, as user `sybase`, and run the following command on the Sybase Server:

```
/opt/omni/bin/utilns/testbar -type:Sybase  
-perform:checkuser -bar:FullSybase
```

Figure 1-55

Checking the Sybase User



```
Terminal  
Window Edit Options Help  
$ /opt/omni/bin/utilns/testbar -type:sybase -perform:checkuser -bar:FullSybase  
=====  
Parsing options:  
'-type:sybase' :OPT_TYPE: 'sybase'  
'-perform:checkuser' :OPT_PERFORM: 'CHECKUSER'  
'-bar:FullSybase' :OPT_BAR: 'FullSybase'  
=====  
[Normal] Parsing environment options.  
[Normal] Performing common initialisation.  
[Normal] Getting cell server host from config file.  
[Normal] Performing cell server connection initialisation.  
[Normal] Getting current user information.  
[Normal] Testing if user can see private BAR objects.  
[Normal] Testing if user can perform restore.  
[Normal] Check user successfully finished.  
  
-----  
OmniBackII testbar finished.  
-----  
$
```

In the example depicted in Figure 1-55, the user has all the appropriate rights.

If a user `ana` on Sybase Server `nyasha.zim.com`, does not have the appropriate rights, you get an error message like the following:

```
[Critical] From: OB2BAR@nyasha.zim.com "" Time: 08/06/99  
17:35:37
```

[131:53] User "ana.users@nyasha.zim.com" is not allowed to perform a restore.

See “Configuring a Sybase User in Data Protector” on page 17 for information about the right privileges.

Backup Problems

If you have problems backing up Sybase databases, proceed as follows:

1. Make a Data Protector filesystem backup of the problematic client.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.

2. Ensure that Sybase SQL Server is up and running.

The simplest way to find whether Sybase SQL Server is running is to try and log on to the server using the `isql` command:

- a. Log on to Sybase SQL Server as user `sybase`.
- b. Type in the following command in the Sybase SQL Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>,
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server; `audi` in Figure 1-53, `<PASSWORD>` is the Sybase Administrator password `<SA>` is the Sybase user.

Figure 1-56

Checking if Sybase SQL Server Is Running



If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<SYBASESERVERNAME>
```

3. Ensure that Sybase Backup Server is up and running by logging on to the server using the `isql` command:
 - a. Log on to Sybase Backup Server as user `sybase`

- b. Type in the following command in the Sybase Backup Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<BACKUPSERVERNAME>
```

where *<BACKUPSERVERNAME>* is the name of Sybase Backup Server; *audi_back* in Figure 1-54, *<PASSWORD>* is the Sybase Administrator password *<SA>* is the Sybase user.

Figure 1-57

Checking if Sybase Backup Server Is Running



If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<BACKUPSERVERNAME>
```

4. Verify the configuration of your Sybase server. using the following command:

```
util_sybase.exe -CHKCONF <SYBASESERVERNAME>
```

where *<SYBASESERVERNAME>* is the name of Sybase SQL Server.

In case of an error, the error number is displayed in the form **RETVAL*<error number>*.

To get the error description, start the command,
/opt/omni/lbin/omnigetmsg 12 <error_number>.

Using the Data Protector CLI

Using the Data Protector GUI

You can also check the configuration of your Sybase Server by performing the following steps in the HP OpenView Storage Data Protector Manager:

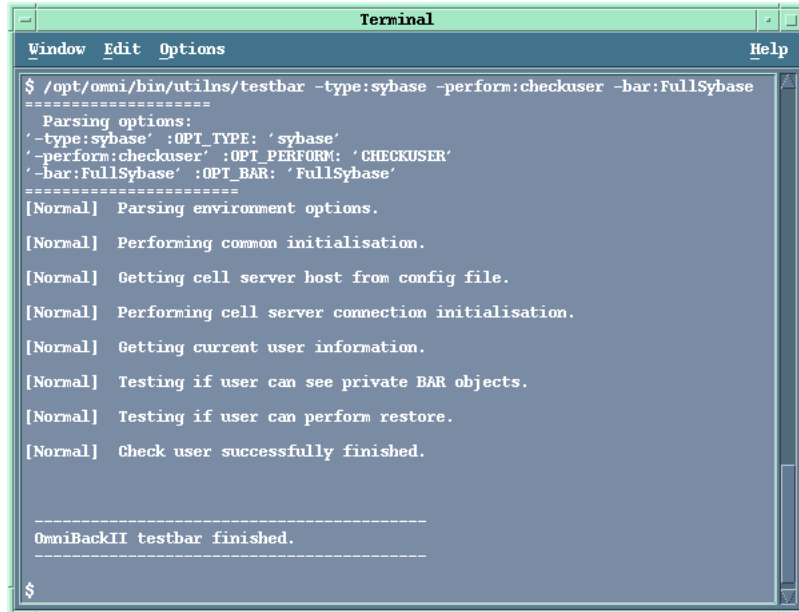
- a. In the Context List, select Backup.
- b. In the Scoping Pane, expand Backup, then Backup Specifications, and then Sybase Server.
- c. Click a configured Sybase backup specification you want.

The Sybase Server is displayed in the Results Area.

- d. Right-click the client and then click Check Configuration.
 A message is returned confirming that the integration is properly configured.
5. Test if the Sybase user has the right privileges in Data Protector. Log in as the Sybase user, for example, as user `sybase`, and run the following command on the Sybase Server:

```
/opt/omni/bin/utilns/testbar -type:Sybase
-perform:checkuser -bar:FullSybase
```

Figure 1-58 Checking the Sybase User



In the example depicted in Figure 1-58, the user has all the appropriate rights for the backup specification named `FullSybase`.

If a user `andrea` on Sybase Server `cool.shon.com`, does not have the appropriate rights, you get an error message like the following:

```
[Critical] From: OB2BAR@cool.shon.com "" Time: 08/06/99 17:51:41
[131:53] User "andrea.users@cool.shon.com" is not allowed to
perform a restore.
```

See “Configuring a Sybase User in Data Protector” on page 17 for information about the right privileges.

6. Test the Data Protector Sybase configuration as per instructions in “Testing the Integration” on page 42.

Example

Run the following command to test the configuration of the backup specification called FullSybase:

```
/opt/omni/bin/omnib -sybase_list FullSybase -test_bar
```

- If the Data Protector part of the test fails then 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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

- a. Verify that the owner of the backup specification is the Sybase user, and if they are in the Data Protector operator or admin group.
- 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.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

- If the test succeeds, start a backup directly from a Sybase Server. See “Backing Up Using Sybase Commands” on page 52 for instructions.

If this backup succeeds, then the problem may be that the client on which the Data Protector User Interface runs does not have enough memory, disk space, or other operating system resources.

7. Test Data Protector data transfer using the testbar utility. Log in as the Sybase user on the Sybase Server and proceed as follows:

```
/opt/omni/bin/utilns/testbar  
-type:Sybase  
-appname:<SYBASESERVERNAME>  
-bar:<backup_specification_name>  
-perform:backup
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server and `<backup_specification_name>` the name of the Data Protector backup specification.

If the test is successful, then proceed to the next step, otherwise proceed as follows:

- a. Troubleshoot errors reported by the testbar utility using the Data Protector troubleshooting file,
`/opt/omni/gui/help/Trouble.txt`.
 - b. Examine system errors reported in the
`/var/opt/omni/log/debug.log` file on the Sybase Server.
8. If you have any non-default Sybase settings, ensure that they are registered in the Data Protector Sybase configuration file. For more information on Data Protector Sybase configuration file, see “Data Protector Sybase Configuration File” on page 9.

Restore Problems

If you have problems restoring Sybase databases, proceed as follows:

1. Make a Data Protector filesystem backup and restore of the problematic client.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.

2. Examine system errors reported in the
`/var/opt/omni/log/debug.log` file on the Sybase server.
3. Ensure that Sybase SQL Server is up and running.

The simplest way to find whether Sybase SQL Server is running is to try and log on to the server using the `isql` command

- a. Log on to Sybase SQL Server as user `sybase`
- b. Type in the following command in the Sybase SQL Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<SYBASESERVERNAME>
```

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server; `audi` in Figure 1-53, `<PASSWORD>` is the Sybase Administrator password and `<SA>` is the Sybase user.

Figure 1-59 Checking if Sybase SQL Server Is Running



If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<SYBASESERVERNAME>
```

4. Ensure that Sybase Backup Server is up and running

The simplest way to find whether Sybase Backup Server is running is to try and log on to the server using the `isql` command

- a. Log on to Sybase Backup Server as user `sybase`
- b. Type in the following command in the Sybase Backup Server home directory:

```
bin/isql -U<SA> -P<PASSWORD> -S<BACKUPSERVERNAME>
```

where `<BACKUPSERVERNAME>` is the name of Sybase Backup Server; `audi_back` in Figure 1-60, `<PASSWORD>` is the Sybase Administrator password and `<SA>` is the Sybase user.

Figure 1-60 Checking if Sybase Backup Server Is Running



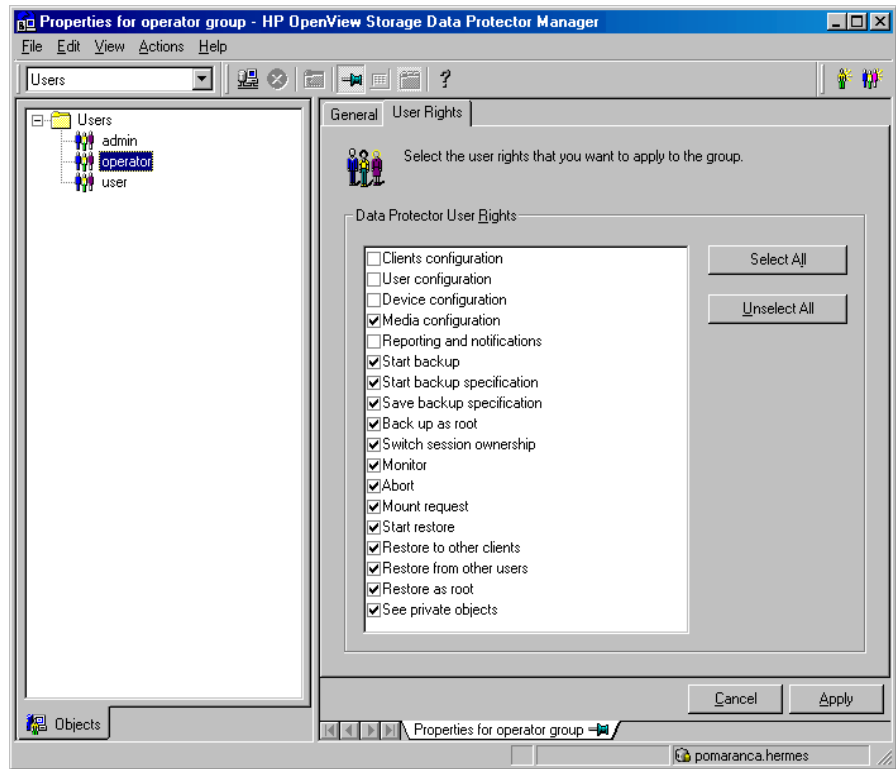
If the server is not running, then perform the following command in the Sybase home directory, to start it:

```
./install/RUN_<BACKUPSERVERNAME>
```

5. Verify that the user specified for the restore session. is the Sybase user, and that they are in the Data Protector operator or admin group.

6. Ensure that the See private objects user right of the Data Protector operator group is selected
 - a. In the Context List, select Users.
 - b. In the Results Area, right-click Operator and click Properties.

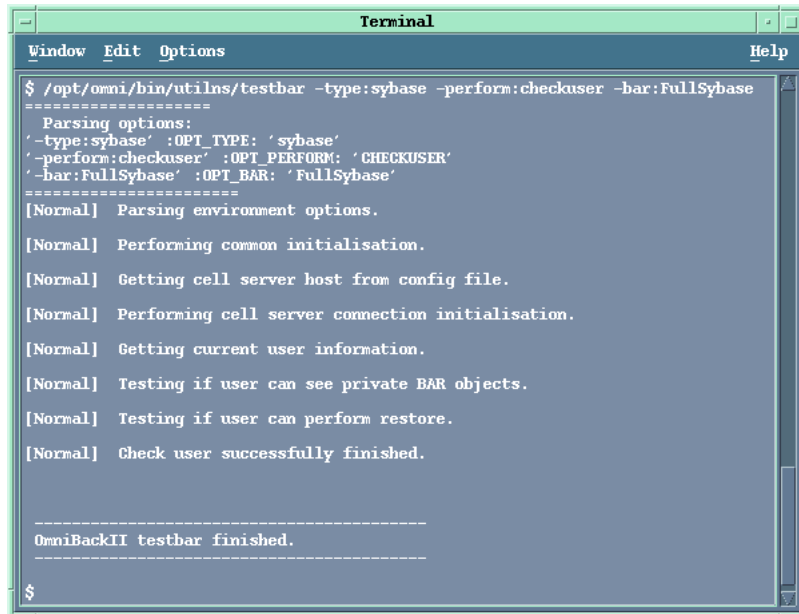
Figure 1-61 Setting the See Private Objects User Right



- c. If the See private objects user right is selected, click Apply.
7. Test if the Sybase user has the right privileges in Data Protector. Log in as the Sybase user, for example as user sybase, and run the following command on the Sybase Server:

```
/opt/omni/bin/utilns/testbar -perform:checkuser
```

Figure 1-62 Checking the Sybase User



```
Terminal
Window Edit Options Help
$ /opt/omni/bin/utilns/testbar -type:sybase -perform:checkuser -bar:FullSybase
=====
Parsing options:
'-type:sybase' :OPT_TYPE: 'sybase'
'-perform:checkuser' :OPT_PERFORM: 'CHECKUSER'
'-bar:FullSybase' :OPT_BAR: 'FullSybase'
=====
[Normal] Parsing environment options.
[Normal] Performing common initialisation.
[Normal] Getting cell server host from config file.
[Normal] Performing cell server connection initialisation.
[Normal] Getting current user information.
[Normal] Testing if user can see private BAR objects.
[Normal] Testing if user can perform restore.
[Normal] Check user successfully finished.

-----
OmniBackII testbar finished.
-----
$
```

In the above example, the user has all the appropriate rights.

If a user ana on Sybase Server nyasha.zim.com is not in the operator or admin group, you get an error message like the following:

```
[Critical] From: OB2BAR@nyasha.zim.com "" Time: 08/06/99
17:35:37

[131:53] User "ana.users@nyasha.zim.com" is not allowed to
perform a restore.
```

See “Configuring a Sybase User in Data Protector” on page 17 for information about the right privileges.

8. Test Data Protector data transfer using the testbar utility. Log in as the Sybase user on the Sybase Server and proceed as follows:

```
/opt/omni/bin/utilns/testbar
-type:Sybase
-appname:<SYBASESERVERNAME>
-bar:<backup_specification_name>
-perform:backup
```

where *<SYBASESERVERNAME>* is the name of Sybase SQL Server and *<backup_specification_name>* the name of the Data Protector backup specification.

If the test is successful, then proceed to the next step, otherwise proceed as follows:

- a. Troubleshoot errors reported by the testbar utility using the Data Protector troubleshooting file,
/opt/omni/gui/help/Trouble.txt.
- b. Examine system errors reported in the
/var/opt/omni/log/debug.log file on the Sybase Server.

2

Integrating Network Node Manager and Data Protector

In This Chapter

This chapter explains how to configure and use the HP OpenView Network Node Manager (NNM) integration.

It is organized into the following sections:

“Overview” on page 99

“Prerequisites and Limitations” on page 101

“Integration Concept” on page 102

“Configuring an NNM Backup” on page 104

“Backing Up an NNM Database” on page 109

“Restoring NNM” on page 113

“Monitoring an NNM Backup and Restore” on page 115

“Troubleshooting” on page 117

Overview

Data Protector offers online backup of NNM. The online backup concept is widely accepted. It addresses the business requirements for high application availability, as opposed to the offline concept.

You can perform online backup of the whole database or parts of it using the Data Protector integration:

Using the Data Protector NNM integration, you can restore of the whole database or parts of it.

Using the Data Protector NNM integration offers several advantages:

- Media Management

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

- Backup Management

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

- Scheduling

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

- Device Support

Data Protector supports a wide range of devices, from files and standalone drives to complex multiple drive libraries. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported backup devices.

- Monitoring

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

All backup sessions are logged in the embedded IDB.

Overview

The administrator is thus provided with a history of activities that can be queried at a later time.

Prerequisites and Limitations

Prerequisites

- Before you begin, ensure that you have correctly installed and configured NNM and the Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Software Release Notes* for an up-to-date list of supported versions, devices, platforms, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector Network Node Manager integration.
 - *HP OpenView Storage Data Protector Administrator's Guide* for instructions on how to configure and run backups.
 - *Reporting and Data Analysis with HP OpenView Network Node Manager* for NNM concepts and backup and recovery strategies.
- It is assumed that you are familiar with NNM administration and basic Data Protector functionality.

Limitation

Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a list of general Data Protector limitations. This section describes limitation specific to this integration.

The Data Protector pre-exec script checks to see if the NNM database is currently backing up. If it is, the pre-exec script aborts.

Integration Concept

The Data Protector NNM integration links the NNM SOLID database with Data Protector. From the NNM point of view, Data Protector represents a media management utility. On the other hand, the NNM database management system can be seen as a data source for backup, using media controlled by Data Protector.

Backup

The method of backup uses very small Perl scripts that are contained within Data Protector. The NNM Perl compiler is used. Backup is performed as follows:

1. The Data Protector Cell Manager invokes the Disk Agent for the NNM backup.
2. The Disk Agent runs the pre-backup Perl script, which informs the NNM embedded database to back itself up to a specified location. The script pauses eight NNM processes after the embedded database performs its backup.

NOTE

The files created by the embedded database backup remain on the system. Future backups simply overwrite any contents in the specified location. The NNM administrator should remove the contents manually to conserve disk space.

3. Data Protector starts the backup of the NNM directory upon successful termination of the pre-backup script.
4. The Disk Agent runs the post-backup Perl script, which resumes the eight paused processes.
5. The backup session ends upon successful termination of the post-backup script.

Restore

To restore information from Data Protector to NNM, the Data Protector administrator does a restore in accordance with Data Protector online Help or the *HP OpenView Storage Data Protector Administrator's Guide* and the NNM administrator follows the instructions contained in the *NNM Reporting and Data Analysis Manual*. The two administrators must communicate and work in concert.

Components

There are two primary components for the integration: `NNMpre.ovpl`, a Perl script that prepares NNM for backup, and `NNMpost.ovpl`, a Perl script that returns NNM to a normal state. On Windows, a third component `NNMScript.exe` is needed.

The software components involved in backup and restore processes are:

- `NNMpre.ovpl`, a script with no arguments that:
 - ✓ Starts the NNM embedded database backup. The embedded database makes a direct copy of itself to a location specified in the `solid.ini` file.
 - ✓ Pauses eight NNM processes
- `NNMpost.ovpl`, a script with no arguments that:
 - ✓ Resumes the eight processes paused by `NNMpre.ovpl`.
- `NNMScript.exe` (on Windows only)
 - ✓ The NNM Perl compiler is used for `NNMpre.ovpl` and `NNMpost.ovpl` and the compiler path must be supplied to Windows on the command line. `NNMScript.exe` finds the location of the NNM Perl compiler and the location of the scripts. The directory location is found via the registry and the location of the compiler and scripts are relative to this location. `NNMScript.exe` also starts the scripts using the NNM Perl compiler.
 - ✓ An argument of `pre` or `post` is given to `NNMScript.exe` to specify which script to run.

Configuring an NNM Backup

To configure an NNM backup, perform the following steps:

1. Configure the backup devices, media, and media pools.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instructions.

2. Create a backup specification specifying the data that you want to back up, the media and devices to which you want your data to be backed up, as well as Data Protector backup options that define the behavior of your backup or restore session.

Once the backup specification is created and saved, it can be scheduled to perform unattended backups. You may use the default backup template for NNM objects, or you can create a new, custom template.

NOTE

If you choose to back up NNM from Data Protector, you must deactivate the default scheduled backup of the embedded database (`solid.ini`). If both NNM and Data Protector are backing up `solid.ini`, the processes may conflict, causing backups that may not restore properly.

Tasks for the NNM Administrator

It's important for the NNM administrator and the Data Protector administrator to work together throughout the integration and backup process. During configuration the NNM administrator must:

- Communicate the location of the NNM backup directory as specified in the NNM embedded database file, `solid.ini`.
- Comment out the line in `solid.ini` that schedules a nightly backup of the NNM embedded database. (The line begins `At=.`)

Creating a New Template

You can use backup templates to apply the same set of options to a number of backup specifications. By creating your own template, you can specify the options exactly as you want them to be.

This allows you to apply all the options to a backup specification with a few mouse clicks, rather having to specify all the options over and over again. This task is optional, as you can use the default template, as well.

If you prefer to use the predefined template, see “Creating a Backup Specification” on page 105 for a detailed explanation.

To create a new backup template, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup and then Templates, and then right-click Filesystem.
3. Click Add Template. Follow the wizard to define the appropriate backup options in your template.

NOTE

If you create your own template and intend to use the NNM integration module, you *must* use the pre- and post-exec scripts exactly as they are used in the default NNM template.

Creating a Backup Specification

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

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

4. Double-click `NT_NNM Backup` or `Unix_NNM Backup` to create backup specifications with predefined options on Windows or UNIX clients, respectively. You can also double-click `Blank Filesystem Backup` to create a backup specification without pre-defined options or use the pre-defined template.
5. Select the `Local` and `network` backup type and click `OK`. Select the appropriate client and the directories to be backed up on this client. Click `Next`.

Once you have entered the required information, the `Backup Wizard` is started, provided that the respective NNM device has already been configured. If not, you must configure the client at this stage by entering the appropriate connection strings.

NOTE

If you still have not configured your devices and media, do so now. Refer to online Help or the *HP OpenView Storage Data Protector Administrator's Guide*.

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

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

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

Click `Next`.

7. Follow the wizard to define options and the schedule to be used.

IMPORTANT

Although you can configure the backup options as you like, the pre- and post-exec options are set by default and *must not* be changed.

Refer to the online Help or the *HP OpenView Storage Data Protector Administrator's Guide* for backup options common to all objects. See “NNM Backup Options” on page 107 for details about NNM specific options.

Once you have defined all backup options, you must save and name your NNM backup specification under a name of your choice. It is recommended that you save all NNM backup specifications in the NNM group.

8. If the backup is to be scheduled, specify the dates and times that the backup should be performed. Skip this step if the backup is to be manual. Click **Next**.
9. Modify the backup specification as necessary. Click **Next**.
10. Save, preview, or start your backup.
11. You can examine the just-created and saved backup specification in the Backup context, under the specified group of backup specifications. The backup specification itself is stored in the following file on the Cell Manager:
 - on Windows:
`<Data_Protector_home>\Config\server\datalists\<Backup_Specification_Name>`
 - on UNIX:
`/etc/opt/omni/server/datalists/<Backup_Specification_Name>`

NNM Backup Options

The NNM backup options are specified in the Data Protector GUI in the Application Specific Options window.

This window can be accessed from the Options property page of an NNM backup specification by clicking the **Advanced** button.

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a live backup. The test verifies that the pre- and post-exec scripts are functioning and that the configuration is valid.

Integrating Network Node Manager and Data Protector
Configuring an NNM Backup

The procedure consists of running a test backup. If the scripts execute, the test is a success.

Backing Up an NNM Database

There are two strategies for backing up a database. These are an offline or consistent database backup, and an online or inconsistent database backup. The latter is also known as a hot backup.

NNM Offline

An offline backup of a database is a backup of the datafiles and control files that are consistent at a certain point in time. The only way to achieve this consistency is to cleanly shut down the database and then back up the files while the database is closed.

The issue with offline backup of NNM is that the NNM and Data Protector administrators must work carefully together to synchronize the events on the NNM and Data Protector systems.

The offline database backup is performed as follows:

1. Shut down the database cleanly by typing `ovstop` at the command line of the NNM machine.
2. Use Data Protector to back up the complete NNM tree.
3. Restart the database by typing `ovstart` at the command line of the NNM machine.

NNM Online

As opposed to the offline backup, the online backup is performed when a database is open.

The backup of an open database is generally thought to be inconsistent, because portions of the database are being modified and written to disk while the backup is progressing. With the NNM integration module, however, all changes to the database are entered into temporary files, as well. When the database is removed from its pause state, the information that has been accumulating in the temporary files is written to the database.

To run an online backup of NNM, use any of the following methods:

- Schedule the backup of a saved NNM backup specification using the Data Protector Scheduler. See “Scheduling a Backup” on page 110.
- Start an interactive backup of the NNM backup specification. See “Starting an Interactive Backup” on page 111.

Backup Procedure This is what happens when you start an NNM backup using the Data Protector User Interface:

Windows

1. Data Protector executes `NNMScript.exe`, which finds the location of the NNM Perl compiler and the `NNMpre.ovpl` script and then starts the script.
2. The Data Protector backup commences. Data Protector extracts data from the client and writes it to the backup device.
3. When the backup is complete, Data Protector executes `NNMScript.exe`, which finds the location of the NNM Perl compiler and the `NNMpost.ovpl` script and then starts the script.

UNIX

1. Data Protector executes `NNMpre.ovpl` on the client. This script starts the backup of the NNM embedded database, which makes a direct copy of itself to a location specified in the `solid.ini` file. The script also pauses eight NNM processes.
2. The Data Protector backup commences. Data Protector extracts data from the client and writes it to the backup device.
3. When the backup is complete, Data Protector executes `NNMpost.ovpl`, a script with no arguments that resumes the eight processes paused by `NNMpre.ovpl`.

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

Scheduling a Backup

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

A backup schedule can be tailored according to your business needs. If you have to keep the database online continuously, then you should back it up frequently.

For example, you may decide to schedule backups of production databases like this:

- Weekly full backup
- Daily incremental backup

To schedule an NNM backup specification, proceed as follows:

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

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

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

The backup type can be full or incremental, with the incremental level as high as level 4.

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

Starting an Interactive Backup

You are most likely to run an interactive backup after creating a new backup specification or when you need a backup immediately and the corresponding backup specification is scheduled at a later time.

The interactive backup can be started using the Data Protector GUI or Data Protector CLI.

When you start a backup, Data Protector invokes `NNMpre.ovpl` (UNIX systems) or `NNMScript.exe` (Windows systems) on the NNM system and the Media Agents on the client system on which backup devices are configured.

Running a Backup Interactively Using the Data Protector GUI

Follow the procedure below to start an interactive backup of an NNM backup specification:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then Filesystem.
3. Right-click the backup specification, then click Start Backup.
Select the backup type and network load in the Start Backup window.
4. Click OK to execute the backup. Upon successful completion of the backup session, a Session Completed message appears.

Restoring NNM

Data Protector acts as a media management utility for the NNM system. Therefore, NNM utilities must be used for a restore. Basically, the Data Protector administrator does a restore according to the instructions in the *HP OpenView Storage Data Protector Administrator's Guide*, and the NNM administrator follows the instructions contained in the *NNM Reporting and Data Analysis* manual. The two administrators must communicate and work in concert.

The basic restore process follows this model:

1. The NNM administrator stops all NNM processes.
2. The Data Protector administrator restores data from the specified backup.
3. The NNM administrator carries out NNM recovery procedures.
4. The NNM administrator restarts all NNM processes.

Disaster Recovery

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

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

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.
2. Install, configure, and initialize the database/application so that data on the Data Protector media can be loaded back to the system.
Consult the documentation from the database/application vendor for a detailed procedure and the steps needed to prepare the database.

Restoring NNM

3. Ensure that the database/application server has the required Data Protector client software installed and configured for the database/application.
4. Start the restore. When the restore is complete, follow the instructions from the database/application vendor for any additional steps required to bring the database back online.

Monitoring an NNM Backup and Restore

During a backup, system messages are sent to the Data Protector. You can monitor the backup session from any Data Protector client on the network where the Data Protector User Interface is installed.

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

Acceptable Warnings on the Windows Systems

The following messages are likely to be generated during an NNM backup. The warnings listed are informational and do not impact the reliability or validity of the backup. Any warnings or errors other than these warrant closer inspection. See “Error and Warning Messages” on page 117 for more information about other messages.

```
[Warning] From: <session_owner> Time: <mm/dd/yy hr:mn:sc> [<error code>] <path>\HP OpenView\NNM\bin\tcl7.5.dll  
Cannot preserve time attributes: ([5] Access is denied.).
```

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

```
[Warning] From: <session_owner> Time: <mm/dd/yy hr:mn:sc> [<error code>] <path>\HP OpenView\NNM\databases\analysis\default\solid.db  
Cannot open: ([33] The process cannot access the file ....).
```

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

```
[Warning] From: <session_owner> Time: <mm/dd/yy hr:mn:sc> [<error code>] <path>\HP OpenView\NNM\databases\openview\topo\netmon.lock  
Cannot open: ([33] The process cannot access the file ....).
```

```
[Warning] From: <session_owner> Time: <mm/dd/yy hr:mn:sc> [<error code>] <path>\HP OpenView\NNM\databases\snmpCollect\dblock  
Cannot open: ([33] The process cannot access the file ....).
```

Monitoring an NNM Backup and Restore

```
[Warning] From: <session_owner> Time: <mm/dd/yy hr:mn:sc> [<error  
code>] <path>\HP OpenView\NNM\databases\snmpCollect\snmpCollectPid  
Cannot open: ([33] The process cannot access the file ....).
```

These three zero-byte files are not significant to Data Protector.

Troubleshooting

Before you start troubleshooting the Data Protector NNM integration, check the following:

1. Ensure that the latest official Data Protector patches are installed. Refer to “Verifying Which Data Protector Patches Are Installed” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* or

http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

2. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a description of general Data Protector limitations, supported versions, problems, and workarounds, and a list of related Data Protector patches.

The following sections provide some checking procedures you should perform before you call Data Protector support. In this way you may either resolve the problem yourself or identify the area where the difficulties are occurring.

Should you fail when performing a troubleshooting procedure, an action is proposed to help you work around the problem.

Error and Warning Messages

Error

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

This message is output if NNM is already paused. In this case the script `NNMpre.ovpl` fails. Check to see if the NNM administrator has manually paused the NNM processes.

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

This message is output if a Data Protector NNM session has completed and Data Protector tries to execute `NNMpost.ovpl`. If the NNM processes are active (not paused), the script fails and Data Protector returns the message that the backup completed with errors (meaning `NNMpre.ovpl` executed without errors, the backup executed without errors, but `NNMpost.ovpl` failed). The backup should be considered unreliable because the processes must have been restarted manually sometime during the Data Protector backup. It is also possible that all NNM processes have been completely stopped.

Make sure that the NNM administrator knows when Data Protector sessions are taking place to avoid this problem. Run the backup again.

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

This output occurs if not all the NNM processes are running. The pre-backup script, `NNMpre.ovpl`, is not able to connect to the NNM embedded database, so the script fails. Check to see if the NNM administrator has manually stopped the NNM processes. NNM must be running for the script to succeed.

Error **Embedded database is currently in the backup process.
Aborting Data Protector backup.**

This output occurs if the NNM embedded database has an active backup in progress. Make sure that the NNM administrator has commented out the default scheduled backup in the `solid.ini` file.

Windows-Specific Error and Warning Messages

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

This output occurs if the wrong number of arguments is given to `NNMScript.exe`. Change the pre/post backup options to use the correct argument.

Error **Couldn't find Network Node Manager key in registry.**

This error occurs if NNM is not installed on the target host or the registry key has been deleted. Make sure that NNM is installed on the target host. If so, ensure the registry entry for NNM under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView` exists, and has a key named `Network Node Manager`. Any other name will result in failure.

Error **Couldn't find the Network Node Manager PathName in registry.**

This error occurs if the string value `PathName` doesn't exist under the `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\Network Node Manager` key. Ensure that a string value for the NNM path exists under this key. If so, make sure that it is named `PathName`. If not, contact the NNM administrator.

Error **Couldn't find OmniBack II key in registry.**

This error occurs if Data Protector is not installed on the target host or the registry key has been deleted. Make sure that a Disk Agent exists on the target host. If so, ensure that the registry entry for Data Protector under `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView` exists, and has a key named `OmniBackII`. Any other name will result in failure. Reinstallation of the Disk Agent may be necessary.

Error **Couldn't find the Data Protector HomeDir in registry.**

This error occurs if the string value `HomeDir` does not exist under the `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Common` key. Make sure that the string value for the Data Protector path exists under this key. If so, ensure it is named `HomeDir`. If not, then either create it or reinstall the Disk Agent.

Error **Incorrect argument. Use "pre" or "post".**

This error occurs if the argument given to `NNMScript.exe` is incorrect. Change the pre/post backup options to use the correct argument. The arguments are not case sensitive.

Error **"Failure starting "NNM_perl_compiler_path Data Protector_Home_Dir\bin*.ovpl".**

This error occurs if `NNMScript.exe` can't execute `NNMpre.ovpl` or `NNMpost.ovpl`. Contact the NNM administrator to make sure that NNM's Perl compiler has not been removed. Ensure that the paths for Data Protector and NNM in the registry are correct.

Error

Execution of "NNM_perl_compiler_path Data Protector_Home_Dir\bin*.ovpl failed.

This error occurs if the command to run either `NNMpre.ovpl` or `NNMpost.ovpl` fail. Make sure that `<path>\HP OpenView\NNM\bin` is in the PATH. Also, make sure that the scripts are in the `<<Data_Protector_home>\bin` directory.

Backup and Restore Problems

If the pre-exec script fails, it is possible that the following errors exist:

- Some or all of the NNM processes are already paused or stopped. Contact the NNM administrator.
- The NNM embedded database was already in a backup state.
 - ✓ Ensure that the default scheduled backup in the `SOLID.ini` file is commented out. Contact the NNM administrator to verify this.

If the post-exec script fails, the NNM processes may have been running during backup. *Backup should be considered unreliable.* The processes may have been resumed manually during backup.

In This Chapter

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

This chapter is organized into the following sections:

- “Overview” on page 123
- “Prerequisites and Limitations” on page 125
- “Integration Concept” on page 127
- “Network Data Management Protocol (NDMP)” on page 130
- “Configuring the Integration” on page 135
- “Network Appliance Configuration” on page 146
- “EMC Celerra Configuration” on page 148
- “Backing Up the NDMP Server Data” on page 149
- “Restoring the NDMP Server Data” on page 154
- “NDMP Environment Variables” on page 158
- “The NDMP Related omnirc File Variables” on page 160
- “Media Management” on page 164
- “Troubleshooting” on page 165

Overview

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

The Data Protector NDMP server integration supports filesystem backups and the following types of restore:

- Filesystem restore
- Direct access restore

Integrating Data Protector with the NDMP server offers the following features:

- Central management for all backup operations:

The administrator can manage backup operations from a central point.

- Media management:

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

- Scheduling:

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

- Reporting:

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

Overview

- **Monitoring:**

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

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

Prerequisites and Limitations

The following is a list of prerequisites and limitations that are specific to this integration:

Prerequisites

- You need a special license to use the Data Protector NDMP server integration. For more information, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- Before you begin, ensure that you have correctly installed and configured the NDMP server and the Data Protector Cell Manager system. Refer to the:
 - *HP OpenView Storage Data Protector Software Release Notes* for an up-to-date list of supported versions, devices, platforms, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector and how to install the Data Protector NDMP clients.
 - *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instruction on how to configure and run backups.

Limitations

- Only filesystem backups are possible.
- A limited set of standard Data Protector media management functions is supported.
- The NDMP integration can handle backups of up to 20 million files if up to 10% of the total number of backed up files are directories, for an average directory name length of 25 characters, and average filename length of 10 characters. In such a case, the NDMP integration allocates up to 1.9 GB of system memory and 2.8 GB of disk space.

For optimal performance the recommended number of files and directories for an NDMP backup specification is 10 million.

The default upper limit for the number of files for an NDMP backup specification is 5 million. To enable higher values, the `OB2NDMPMEMONLY` omnirc file variable must be set to 0. Refer to “The NDMP Related omnirc File Variables” on page 160 for more information on omnirc file variables.

- Only static backup specifications are supported.

Prerequisites and Limitations

- It is not possible to have an NDMP backup session and a normal Data Protector session on the same medium.
- Maximum device concurrency is 1.
- Only the devices supported by Data Protector and the NDMP Server are supported. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for the device support matrix.
- Device as well as filesystem browsing is not possible.
- Device block size is limited to 64K.
- NDMP devices must use special dedicated media pools.
- Only FULL and INC1 backup levels are supported.
- It is not possible to deselect a subtree or a file of the tree selected for restore.
- Object copying, object mirroring and media copying is not supported for NDMP backup.

NetApp NAS Device Limitations

- Direct access restore for files is supported only on the NDMP Server ONTAP v6.1.x and higher.
- Direct access restore for directories is supported only on the NDMP ONTAP Server v6.4.x and higher. With direct access restore for directories, if you select both directories and files for a restore in the Data Protector Restore context, only files are restored.

Celerra NAS Device Limitations

- If you select directories to be restored in the Data Protector Restore context on the Celerra NAS device using the direct access restore for files, only the selected directory without its contents is restored.
- Direct access restore for directories is not supported on Celerra NAS Device.


Refer to the *HP OpenView Storage Data Protector Software Release Notes* for an up-to-date list of supported versions, devices, platforms, and other information.

Integration Concept

Data Protector uses the NDMP interface for backing up and restoring the NDMP Server data using the Data Protector NDMP Server integration.

From the Data Protector perspective, NetApp NAS device and Celerra NAS device are NDMP Servers with its own specifics on the execution level. To support the backup and restore of both NAS devices within the NDMP framework, the backup application has to implement the NDMP client. This client controls backup and restore operations of the NDMP Server through the NDMP protocol. The Data Protector NDMP Media Agent software component must be installed on the NDMP client. The client resides on a host in the LAN. In such a configuration, the NDMP Server disks are typically backed up on locally attached tape devices, so that data does not flow through the LAN. Data movement is local on the NDMP Server with the NDMP client (Data Protector NDMP Media Agent client) controlling the whole process. See Figure 3-1 on page 128.

Figure 3-1 **The NDMP Environment Configuration**



The important consequence of such a configuration is that the actual writing and reading of backed up data is not done by Data Protector, but by the NDMP Server itself in its own format. Since Data Protector cannot recognize this data format, some limitations apply for management of the NDMP media. For more information, see “Prerequisites and Limitations” on page 125.

Data Protector does not influence the speed of backup or restore. It only initiates the backup or restore session. The only performance concern for Data Protector is the CPU load and memory consumption on the NDMP client because of the processing of catalog information. Due to the NDMP catalog handling design, Data Protector caches the entire catalog

information on the NDMP client system before storing it into the Data Protector internal database (IDB). Since the catalog can grow quite big, depending on the number of files backed up, the NDMP client caches parts of the catalog into **file history swap files** on disk, rather than keeping it in memory.

By default, Data Protector writes these files in the `<Data_Protector_home>\tmp` (Windows systems) or in the `/var/opt/omni/tmp` (UNIX systems) directory on the NDMP client.

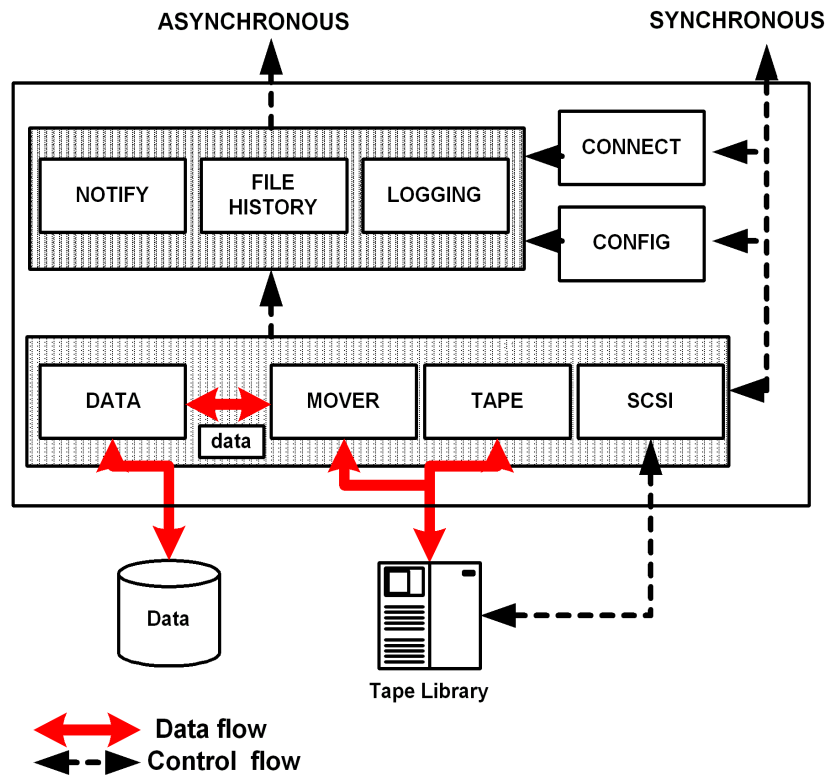
NOTE

The file history swap files can grow very large. See “The NDMP Related omnirc File Variables” on page 160 for more information on how to change the location of the file history swap files, other related parameters, and the expected size of file history swap files.

Network Data Management Protocol (NDMP)

NDMP is a protocol used for communication between a data management client and a data management server in the environment. The NDMP client (Data Protector NDMP Media Agent client) is used to initiate, monitor and control the data management operations. The NDMP server is used to actually execute those operations. Figure 3-2 shows the NDMP interfaces:

Figure 3-2 **The NDMP Interfaces**



Functionality, available for the NDMP client, is inherently defined by the available NDMP interfaces:

- **CONNECT** interface

This interface is used after establishing the connection to the NDMP Server. The **CONNECT** interface allows the NDMP Server to authenticate the client and negotiate the version of protocol used.

- **CONFIG** interface

This interface allows the NDMP client to discover the configuration of the NDMP Server. The **CONFIG** interface can be used to discover the NDMP Server configuration and attributes.

- **SCSI** interface

This interface is used to pass SCSI CDBs through to a SCSI device and retrieve the resulting SCSI status. The NDMP client uses the **SCSI** interface to control a locally attached library. Software on the NDMP client constructs SCSI CDBs and interprets the returned status and data. The **SCSI** interface can also be used to exploit special features of SCSI backup devices.

- **TAPE** interface

This interface supports both tape positioning and tape read and write operations. The NDMP client typically uses the **TAPE** interface to write tape volume header and trailer files. The NDMP client also uses the **TAPE** interface to position the tape during backup and restore sessions.

- **MOVER** interface

This interface is used to control reading and writing of backup data to and from the tape device. During a backup, the **MOVER** reads the data from the data connection, buffers the data into tape records, and writes the data to the tape device. During a restore, the **MOVER** reads the data from the tape device and writes the data to the data connection. The **MOVER** is responsible for handling tape exceptions and notifying the NDMP client.

- **NOTIFY** interface

The NDMP Server uses this message to notify the NDMP client that the NDMP Server requires attention.

- FILEHISTORY

These messages allow the NDMP Server to make entries in the file history for the current backup. The NDMP client uses the file history to select files for retrieval.

NOTE

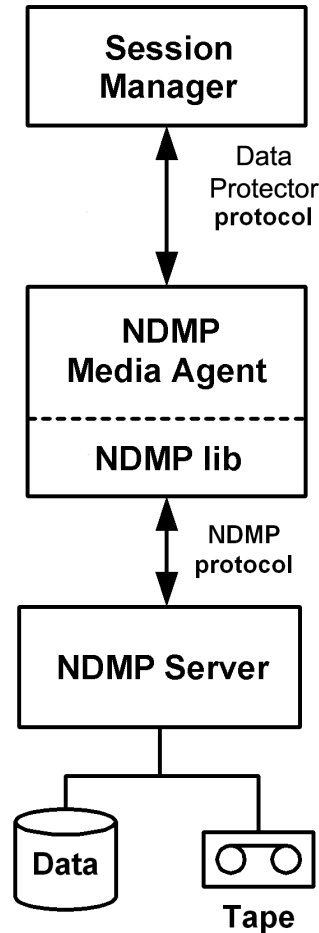
Using Data Protector, you can configure the NDMP Server not to create the file history information. See “NDMP Environment Variables” on page 158 and “The NDMP Related omnirc File Variables” on page 160. Note that you cannot enable the direct access restore if the file history information is not created (is disabled).

- LOG interface

These messages allow the NDMP Server to make entries in the backup log. The operator uses the backup log to monitor the progress and completion status of the backup. The log is also used to diagnose problems.

See Figure 3-3 for the schematic view of the backup environment controlled by Data Protector where the NDMP client functionality is implemented within the NDMP Media Agent.

Figure 3-3 Schematic View of the Data Protector Controlled NDMP Environment



Main modules of such a configuration are:

- Session Manager:

Session Manager controls the backup and restore operations as in the standard Data Protector session. The main difference is that there are no Disk Agents involved because the functionality, needed for the operation, is implemented within the NDMP Media Agent.

- The NDMP Media Agent:

The NDMP Media Agent implements the NDMP client functionality. It is linked with the Data Protector NDMP library that enables it to communicate with the NDMP Server using the NDMP specified interfaces. The NDMP Media Agent handles the tape header, starts a backup or restore session, monitors the operation execution, and handles the catalog dumping.

It is important to understand that the NDMP client in general is not involved in any data moving and does not access the device directly. It controls the operations and accesses the devices only indirectly through the NDMP interface.

Configuring the Integration

Prerequisite

The NDMP Server must always have a tape drive connected. The library robotics can be connected either to the NDMP Server or to a Data Protector client (Data Protector NDMP Media Agent or Data Protector General Media Agent client) depending on the chosen configuration. For more information on sharing libraries with multiple systems, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

Library Devices

Library devices can be controlled in two ways:

- The library robotics is attached to a Data Protector client (Data Protector NDMP or General Media Agent client).

In this case, all library devices (including ADIC/GRAU and StorageTek ACS libraries) supported on the Data Protector client with the robotics attached can be used.

Refer also to Figure 3-4 on page 136.

- The library robotics is attached to the NDMP Server.

If the library robotics is attached to the NDMP Server, the library must be supported by the NDMP Server as well as by Data Protector.

Refer also to Figure 3-5 on page 137.

For information on library devices supported by Data Protector, refer to support matrices in the *HP OpenView Storage Data Protector Software Release Notes*.

Supported Drives

The support of different drives also depends on the NDMP Server. Backup operations can be done with all drives which are supported by the NDMP Server and by Data Protector.

Library Configurations


Data Protector provides flexible library configuration by enabling the sharing of drives in the library with other applications and among multiple Data Protector client systems, regardless of the platform used.

Data Protector NDMP clients (Data Protector clients that control the NDMP backup) must have the Data Protector NDMP Media Agent installed.

Integrating the NDMP Server and Data Protector
Configuring the Integration

For more information on sharing libraries with multiple systems, refer to the *HP OpenView Storage Data Protector Concepts Guide*. Figure 3-4 and Figure 3-5 show examples of a library configuration for the Data Protector NDMP integration.

Figure 3-4 Library Configuration—II



The Data Protector client that controls the backup of the NDMP Server must have the NDMP Media Agent installed. Several drives can be connected to the NDMP Server.

In Figure 3-4, the library robotics is controlled by either a Data Protector General Media Agent client or by a Data Protector NDMP Media Agent client, since the Data Protector client does not control the NDMP backup.

The Data Protector client that has a drive attached can have either the General Media Agent or the NDMP Media Agent installed.

The support of different library devices depends on Data Protector.

Figure 3-5 Library Configuration—I




Figure 3-5 shows a configuration where library drives are controlled through the NDMP Server TAPE interface and the library robotics is controlled through the NDMP Server SCSI interface. The Data Protector client that controls the backup of NDMP Servers must have the NDMP Media Agent installed. Several drives can be connected to each NDMP Server.

NOTE

If a drive is not attached to the NDMP Server that controls the library robotics, the Data Protector client configured for the same NDMP Server can have the General Media Agent installed, provided there is at least one Data Protector NDMP Media Agent client configured in the same cell.

The support of different library devices depends on the NDMP Server and on Data Protector.

Configuration Procedure

Configuring the Data Protector NDMP Server integration consists of the following steps:

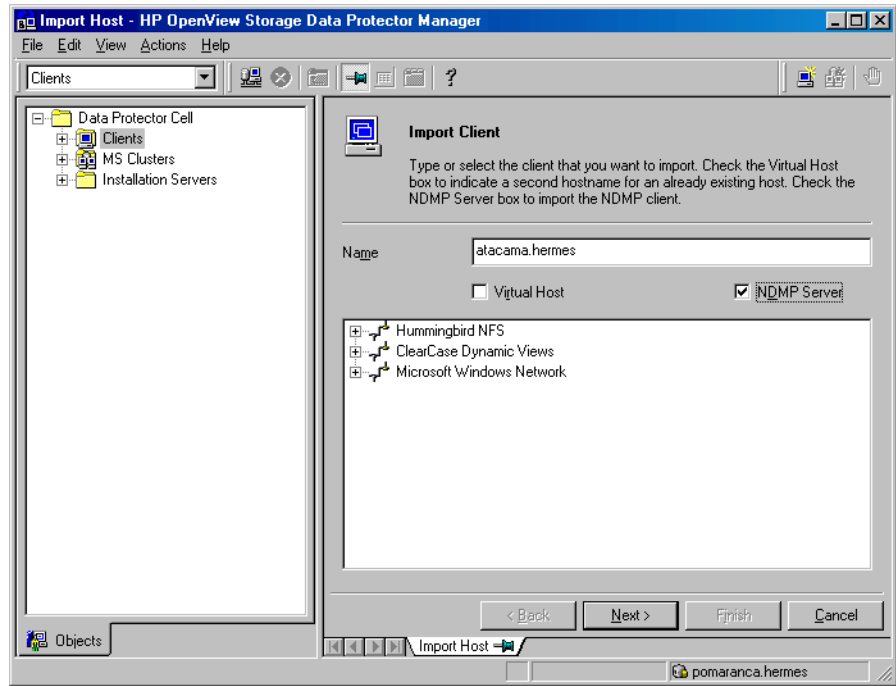
1. Importing the NDMP Server host as a client into the Data Protector cell.
2. Creating a media pool.
3. Configuring a backup device.

Importing the NDMP Server Host

To import the NDMP Server host, perform the following steps:

1. In the HP OpenView Storage Data Protector Manager, switch to the Clients context.
2. In the Scoping Pane, expand Data Protector Cell, and then right-click Clients.
3. Click Import Client.
4. In the Import Client window, enter the name of the NDMP host you want to import, and select the NDMP Server option.

Figure 3-6 Importing the NDMP Server Host

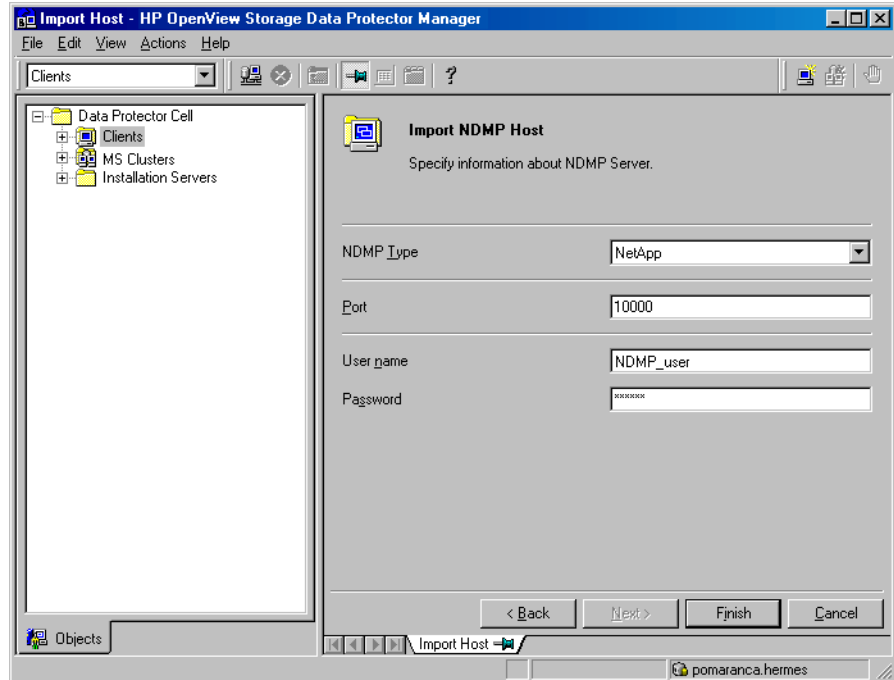


- Click Next.
5. In the Import NDMP Host window, enter the specific import parameters:
In the NDMP Type drop-down list, select the type of your NAS device.
Select the TCP/IP port number of the NDMP Server. The default port number is 10000.
Enter the user name and password that Data Protector will use to establish a connection to the NDMP Server.

Integrating the NDMP Server and Data Protector Configuring the Integration

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

Figure 3-7 Entering the Specific Import Parameters



6. Click **Finish** to import the NDMP host.

What Happens?

When the NDMP host is imported, the NDMP server is also imported to the Data Protector environment.

Creating a Media Pool

NDMP media can only be managed using the NDMP devices (devices configured using the NDMP "data format"). For Data Protector to work properly, the NDMP devices have to use special dedicated media pools. These pools should be used only by the NDMP devices.

Before you create an NDMP device, create a special dedicated media pool to be used for this device:

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for the information on how to create media pools.

NOTE

Data Protector free pools are not supported with the Data Protector NDMP integration.

After you have created the media pool, proceed with device configuration.

Configuring an NDMP Backup Device

Data Protector NDMP Server integration supports standalone and tape library unit backup devices.

An NDMP device (in case it is a device with robotics - tape library unit) can have its robotics attached to the NDMP Server or to a Data Protector NDMP Media Agent or General Media Agent client. Drives can be shared between several NDMP Servers and *Data Protector General Media Agent clients*.

Tape Library Units (TLU) Configuration Procedure

TLU Connected to a Data Protector NDMP or General Media Agent Client To configure a TLU with robotics attached to a Data Protector NDMP or General Media Agent client and drives attached to the NDMP Server, refer to the online Help index keyword “configuring SCSI libraries” and configure the library robotics as described there. Then configure the drives as described in the steps 8 - 11 on page 143.

TLU Connected to an NDMP Server To configure a tape library unit with robotics attached to the NDMP Server, perform the following steps in the Data Protector GUI:

1. Switch to the `Devices and Media` context.
2. Expand the `Environment` item, right-click `Devices`, and then select `Add Device`. The `Add Device` wizard appears.
3. Specify the device name. Optionally, enter the description for the device.

In the `Device Type` drop-down list, select `SCSI Library`.

Integrating the NDMP Server and Data Protector Configuring the Integration

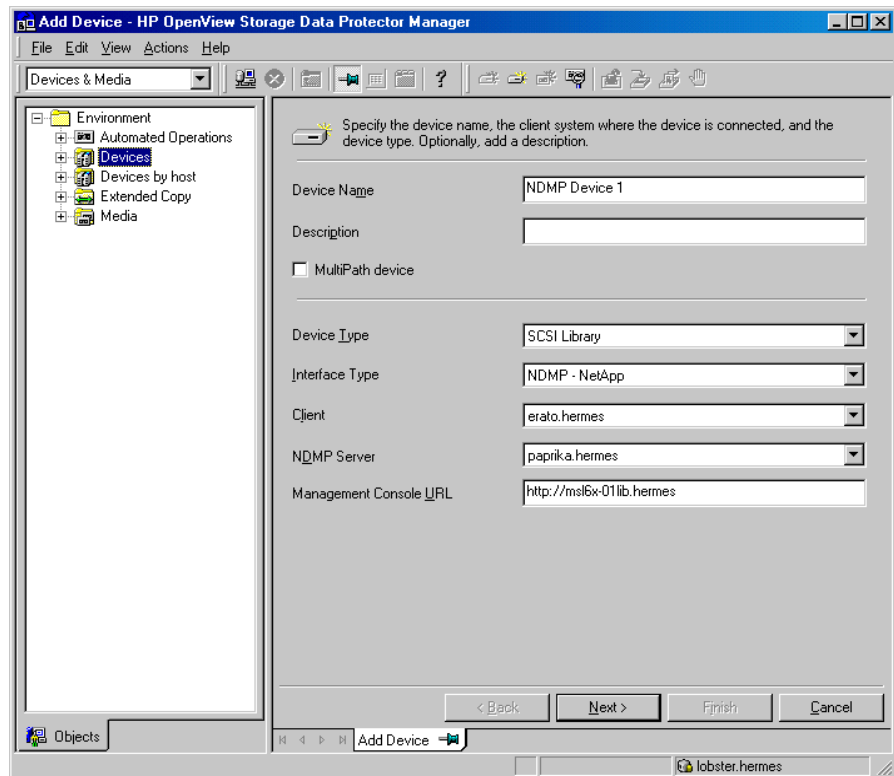
In the `Interface Type` drop-down list, select the NAS device used.

In the `Client` drop-down list, select the client system used for controlling the library through NDMP (the client system running the NDMP or General Media Agent, not the NDMP Server) and then select the NDMP Server with the robotics attached to it from the NDMP Server drop-down list.

Optionally, enter a valid URL of the library management console in the `Management Console URL` text box. This will enable you to invoke a web browser and load the management console interface directly from the Data Protector GUI.

See Figure 3-8 for details.

Figure 3-8 Library Configuration



- Click Next.
4. Following the wizard, enter the required information about library SCSI ID and drive handling. See “Network Appliance Configuration” on page 146 and “EMC Celerra Configuration” on page 148 for more information on library SCSI ID. Click Next.
 5. Specify the slots you want to use with Data Protector. Click Next.
 6. Select the type of media used in the library. Click Next.
 7. Click **Finish** to configure your device, and then click **Yes** to configure the drives in the library.
 8. Specify the drive name. Optionally, enter the description for the drive.
Select the NAS device used in the **Interface Type** drop-down list.
Click Next.
 9. Follow the wizard and enter the information about the drive’s SCSI address. See “Network Appliance Configuration” on page 146 and “EMC Celerra Configuration” on page 148 for more information on obtaining the drive’s SCSI address.
Do not change the drive index number in the **Drive Index** text box.
Click Next.
 10. In the next page of the wizard, specify the information about media and media pools.

NOTE

Multiplexing data streams is not supported by NDMP, therefore the device concurrency is limited to 1.

11. Click **Yes** to create another drive or **NO** to finish creating drives for the library.

Refer to the *HP OpenView Storage Data Protector Administrator’s Guide* or online Help for more information on creating a drive.

Standalone Devices Configuration Procedure

To configure a standalone device proceed as follows:

1. Switch to the **Devices and Media** context.

Integrating the NDMP Server and Data Protector Configuring the Integration

2. Expand the Environment item, right-click Devices, and then select Add Device. The Add Device wizard appears.
3. In the first page of the wizard, specify the device name. Optionally, enter the description.

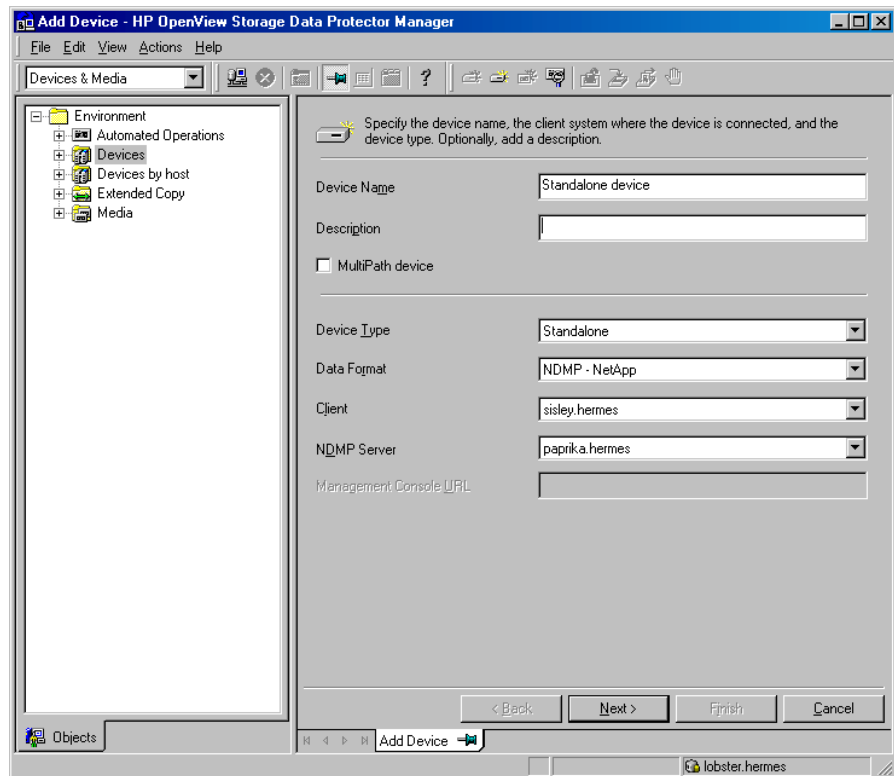
In the Device Type drop-down list, select Standalone.

In the Data Format drop-down list, select the NAS device used.

In the Client drop-down list, select the NDMP client (Data Protector NDMP Media Agent client).

In the NDMP Server drop-down list, select the NDMP Server to which the backup device is connected.

Figure 3-9 Standalone Device Configuration



Click Next.

4. Follow the wizard and enter the information about the standalone device's SCSI address. See "Network Appliance Configuration" on page 146 and "EMC Celerra Configuration" on page 148 for more information. Click Next.
5. In the next page of the wizard, specify the information about media and media pools.

NOTE

Multiplexing data streams is not supported by NDMP, therefore the device concurrency is limited to 1.

6. Click `Finish` to create the drive and exit the wizard.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for more information on creating a standalone device.

Network Appliance Configuration

Before you begin with Network Appliance configuration, consider the following:

- The NDMP Server must be up and running.
- The following NDMP versions are supported: v2, v3, and v4.

To configure the Data Protector part of the integration, you should perform the following tasks:

- Get information about the standalone tape devices or drives in the tape library unit connected to your NDMP Server, which are used to configure the device for the backup session:

Run the `sysconfig -t` command on the NDMP Server. The physical device name always consists of the following parts:

`rst` - always present, means raw SCSI tape.

prefix `n`, `u` - stand for (respectively) no rewind and unload/reload.

NOTE

Data Protector supports only the no rewind devices.

first suffix 0, 1, 2 ... represents the number of the device.

second suffix `l`, `m`, `h`, `a` - represents the data density and compression.

Example - Standalone Device or Drive SCSI Address

Example output for a DLT 4000 drive:

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

When configuring standalone tape devices or drives in the tape library unit connected to your NDMP Server, enter the first part of the output as the SCSI address of the data drive or of the standalone device. In the above example, enter `nrst0m`.

- Get information about the library devices connected to your NDMP Server, which is used to configure the device for the backup session.

To get the information, run the `sysconfig -m` command on NDMP Server. The typical physical device name always consists of the following parts:

`mc` - always present, means SCSI media changer device.

suffix `0, 1, 2 ...` - number of the device.

**Example - Library
Robotics SCSI
Address**

Example output for a DLT 4000 library:

```
>sysconfig -m
```

```
mc0
```

```
>
```

When configuring tape library units connected to your NDMP Server, enter the output as the library's robotics SCSI address. In the above example, enter `mc0`.

- Get information about the filesystems exported from the NDMP Server. This information will be needed when creating the backup specification.

To do this, run the `exportfs` command.

EMC Celerra Configuration

Before you begin with EMC Celerra configuration, consider the following:

- The NDMP Server must be up and running.
- The following NDMP versions are supported: v2, and v3.

To configure EMC Celerra you need to retrieve backup device information. Perform the following steps:

1. Log on to the Celerra control station.
2. To retrieve a list of all SCSI devices attached to the server, run the following command:

```
server_devconfig <server_name> -list -scsi -all
```

3. Use the device addresses to configure the backup device using the Data Protector GUI.

Example

The following is an example of a list of SCSI devices attached to the Celerra NAS device. When configuring the device for use with Data Protector, the device address `c2t010` is used for the library robotics, and the device addresses `c2t310` and `c2t210` are used for the DLT drives.

Table 3-1

Example of a List of SCSI Devices

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

Backing Up the NDMP Server Data

It is assumed that you are familiar with the Data Protector backup procedure. For more information, refer to the online Help or the *HP OpenView Storage Data Protector Administrator's Guide*.

Limitations

- Only filesystem backup is supported.
- It is not possible to have an NDMP backup session and a normal Data Protector session on the same medium.
- Only static backup specifications are supported. No load balancing is supported.
- Maximum device concurrency is 1.
- Device browsing is not possible in the standard Data Protector way.
- Filesystem browsing is not possible.
- Backup level: only FULL and INC1 backups are supported.
- It is not possible to deselect a subtree or a file of the tree selected for backup. This implies that the following is not supported:
 - the GUI Tree/Filters set of options; this involves the Trees, Excludes, Skips and Onlys GUI options
 - the CLI omnib command options `-trees`, `-exclude`, `-skip` and `only`

Backup Procedure Before you start the backup procedure make sure that your media have been formatted. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for the instructions.

To back up a filesystem, perform the following steps:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand the Backup item, and then double-click Backup Specifications.
3. In the Results Area, right click Filesystem and then click Add Backup. The Create New Backup dialog box appears.
4. In the Create New Backup dialog box, select a template to apply to the backup.

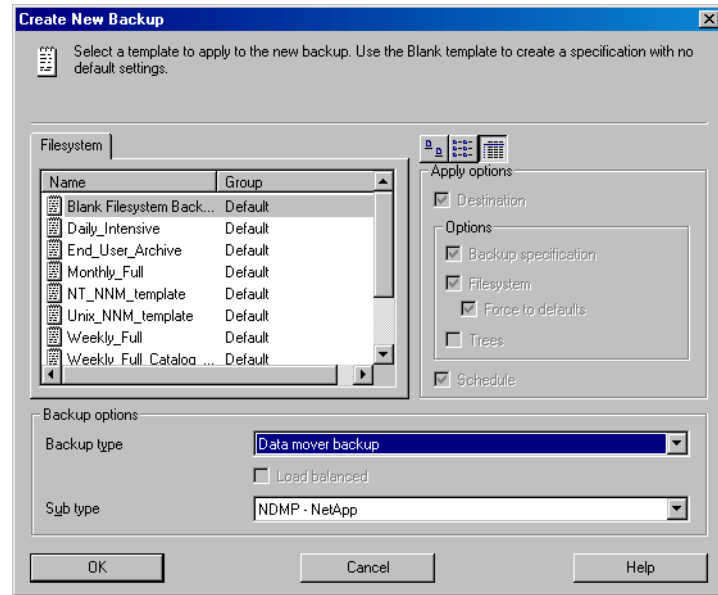
Integrating the NDMP Server and Data Protector

Backing Up the NDMP Server Data

In the Backup type drop-down list, select Data mover backup. In the Sub type drop-down list, select NDMP-NetApp or NDMP-Celerra.

See Figure 3-10 on page 150.

Figure 3-10 Creating a New Backup



NOTE

Load balancing is not supported, so the Load balanced check box is disabled.

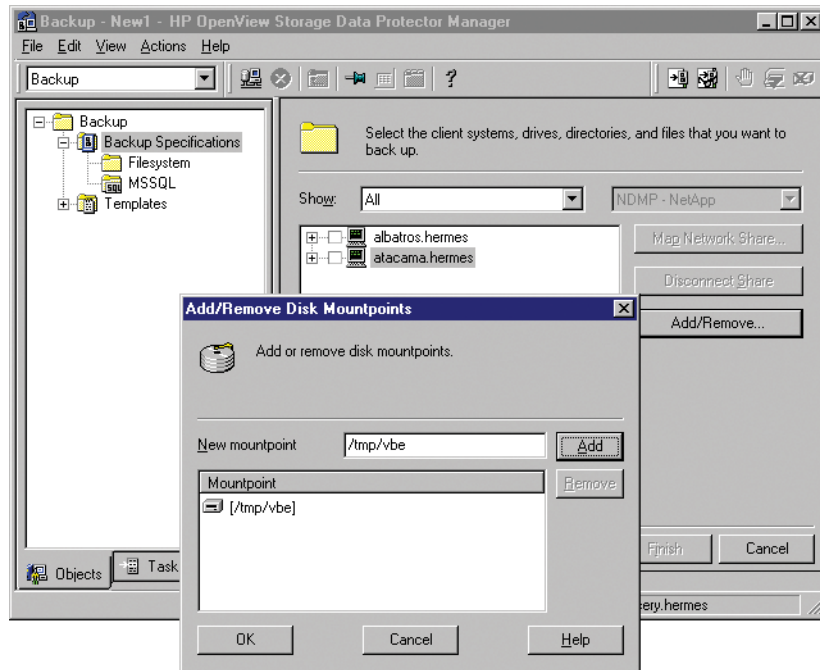
Click OK to start the Backup wizard.

5. In the first page of the wizard, select the data you want to back up. File browsing is not possible. Select the NDMP host and click Add/Remove.

The Add/Remove Disk Mountpoints dialog box appears. Specify the filesystem mount point manually.

- Specify the new directory (with the full path). After you have selected what you want to back up, click **Add** and then **OK**. See Figure 3-11 on page 151.

Figure 3-11 Add/Remove Mountpoint of the NDMP Server



- Click **Next** to display the next page of the Backup wizard.
- Select a device you want to use for the backup, and then click **Next** to proceed.
 - Specify the **Filesystem Options** and **Backup Specification Options**. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on these options. Click **Next**.
 - In the next page of the wizard, you can specify the dates and time you want the backups to be performed. If you want to run an interactive backup, click **Next** to proceed. Refer to the online Help for more information about scheduling your backups.

10. Review the summary of the backup specification. Select the backed up object and click the `Properties` button if you want the `Object Properties` page to be displayed.
11. In the `Object Properties` page, click the `NDMP` tab and specify the `NDMP NetApp` specific options for selected objects.

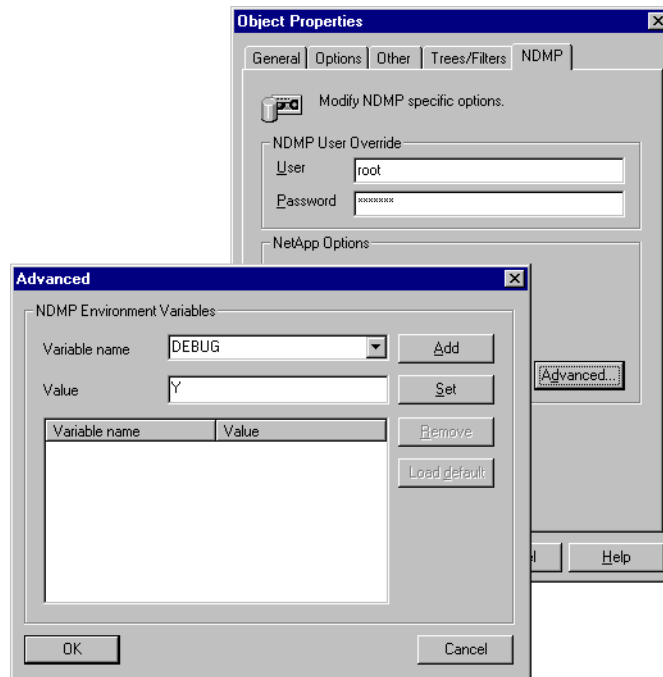
For each object that will be backed up, you can specify a user name and password, which will override the user name and password values entered in the `Import NDMP Host` dialog box. Access rights must be set properly on the `NetApp` or `Celerra` host in order to use user name and password overrides.

You have an option to specify `NDMP` environment variables for specific `NDMP` implementations in the `Advanced` dialog box. See [Figure 3-12](#). See “`NDMP Environment Variables`” on page 158 for more information on available `NDMP` environment variables and their values.

NOTE

Using `Data Protector`, you can configure the `NDMP Server` not to create the file history information. See “`NDMP Environment Variables`” on page 158 and “`The NDMP Related omnirc File Variables`” on page 160. Note that you cannot enable the direct access restore if the file history information is not created (is disabled).

Figure 3-12 Specifying the Advanced Options



Click OK to close the Object Properties dialog box, and then click Next.

12. Save your backup specification, and click Start Backup. The Start Backup dialog box appears.
13. Select the backup type and network load, and then click OK to start the backup session.

Restoring the NDMP Server Data

It is assumed that you are familiar with the Data Protector restore procedure. For more information, refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help.

Limitations

- It is not possible to deselect a subtree or a file of the tree selected for restore. This implies that the following is not supported:
 - the GUI options `Restore only` and `Skip`; deselecting a subtree in the `Restore` context
 - the CLI `omnirc` command options `-only`, `-skip` and `-exclude`
- Direct access restore on NetApp NAS device requires ONTAP v6.1.x or higher.
- Direct access restore is possible only if the file history is switched on on the NDMP Server during the backup session (default). See “The NDMP Related `omnirc` File Variables” on page 160 for more information on how to switch the NDMP Server file history on or off.

Restore Procedure To restore a filesystem backed up from a NAS device through NDMP, perform the following steps:

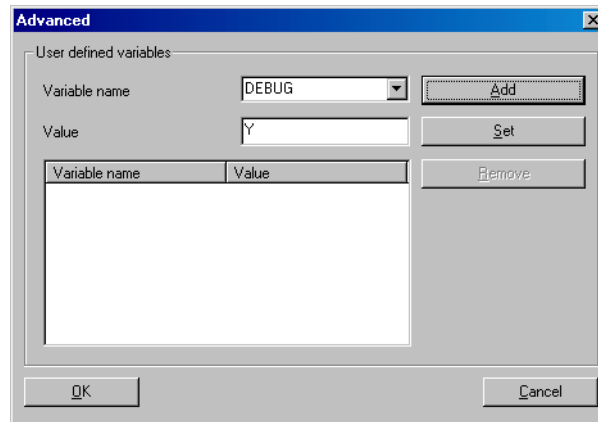
1. In the `Data Protector Manager` window, switch to the `Restore` context.
2. In the `Scoping Pane`, expand `Restore`, and then expand `Filesystem`. The `Restore` wizard appears.
3. Under the `Filesystem` item, browse for and select the backup object you want to restore.
4. For each selected object, enter the required information, such as target destination, as well as devices you want to use for the restore.

You can set the default time interval, which will be used when browsing object versions for restore in the Data Protector database by using the `Search interval`, `From`, `To` and `Update` buttons. Refer to the “Restore” chapter of the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for more information on the search interval.

In the Options property page, specify the user name and the password.

You can specify NDMP environment variables for specific NDMP implementations in the Advanced options dialog box. See Figure 3-13 on page 155 and “NDMP Environment Variables” on page 158 for more information.

Figure 3-13 NDMP Advanced Restore Options



5. Click Restore to open the Start Restore Session window where you can preview your selection.

Select the network load and report level, and then click Finish to exit the wizard and start the restore session.

When the session starts, the messages are displayed in the Results Area.

Direct Access Restore

Direct access restore is an optimized data recovery operation. This functionality enables Data Protector to directly access backed up data in the middle of the tape, without having to parse the tape set sequentially. This is achieved by partitioning the backed up data into segments that are written to tape and recording their location on the tape, together with their start and end addresses relative to the start of the backup data stream. Data Protector computes which segment contains the

starting point of the requested file or directory and the restore process is started from the beginning of the tape containing that segment. The mover then moves across segments and starts reading through the particular segment to locate the beginning of the file or directory.

Prerequisite

Direct access restore is possible only if the file history is switched on the NDMP Server during the backup session. See “NDMP Environment Variables” on page 158 or to “The NDMP Related omnirc File Variables” on page 160 for more information on how to switch the NDMP Server file history on or off.

The procedure for the NDMP direct access restore is the same as for normal NDMP restore, with the exception that you select a single file or directory or more files or directories in the Results Pane of the Data Protector Restore context.

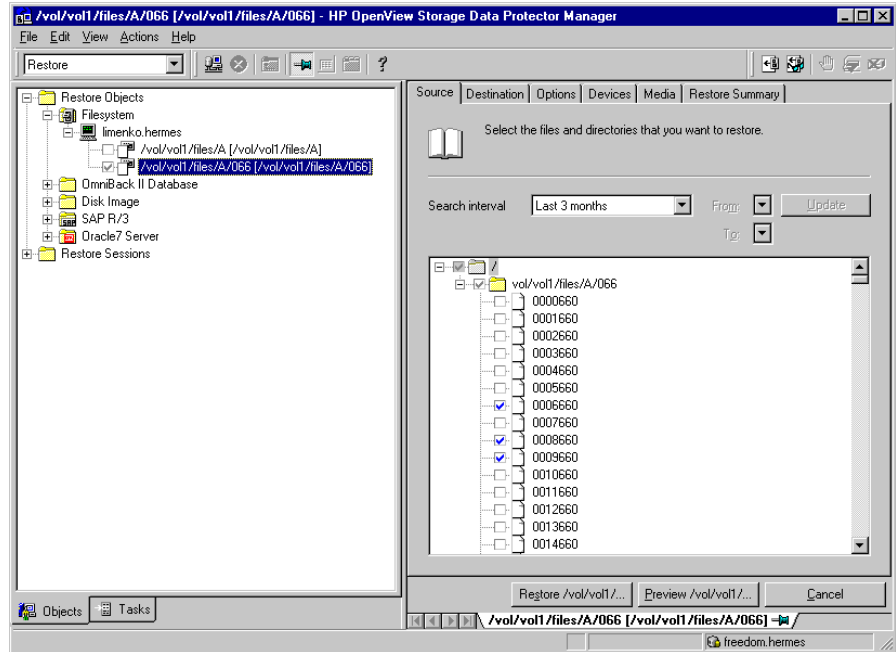
**NetApp NAS
Device Limitations**

- Direct access restore for files is supported only on the NDMP Server ONTAP v6.1.x and higher.
- Direct access restore for directories is supported only on the NDMP ONTAP Server v6.4.x and higher. With direct access restore for directories, if you select both directories and files for a restore in the Data Protector Restore context, only files are restored, the selected directories are not restored.

**Celerra NAS
Device Limitations**

- If you select directories to be restored in the Data Protector Restore context on the Celerra NAS device using the direct access restore for files, only the selected directory without its contents is restored.
- Direct access restore for directories is not supported on Celerra NAS Device.

Figure 3-14 NDMP Direct Access Restore



Restore Using Another Device

Data Protector supports restore using a different device than the original one, which was used at backup time. Refer to the “Restoring Under Another Device” section of the *HP OpenView Storage Data Protector Administrator’s Guide* for more information on how to perform a restore using another device.

NDMP Environment Variables

NDMP environment variables for the NetApp and Celerra NAS device can be set using the Data Protector GUI (as in Figure 3-12 or in Figure 3-13).

The following tables represent the supported user defined NDMP environment variables for the NetApp and Celerra NAS device:

Table 3-2

NDMP Variables for NetApp NAS Device

Variable	Value	Function
HIST	y/n	Enable or disable file history
DIRECT	y/n	Restore using direct access restore
LEVEL	0, 1, 2, ... 9	Backup level (0=full)

Table 3-3

NDMP Variables for Celerra NAS Device

Variable	Value	Function
HIST	y/n	Enable or disable file history
DIRECT	y/n	Restore using direct access restore
LEVEL	0, 1, 2, ... 9	Backup level (0=full)
BASE_DATE	<32bit level><32bit date>	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

NOTE

Some of the stated NDMP environment variables for the NetApp and Celerra NAS device can also be set using the `omnirc` file variables. See “The NDMP Related `omnirc` File Variables” on page 160 for more information on the `omnirc` file variables. The setting in the Data Protector GUI overrides the setting in the `omnirc` file.

The NDMP Related omnirc File Variables

For more information on the location and usage of the `omnirc` file refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE

Some of the variables can also be set using the Data Protector GUI. See Figure 3-12 on page 153 and “NDMP Environment Variables” on page 158 for more information on setting these variables using the Data Protector GUI. The setting in the Data Protector GUI overrides the setting in the `omnirc` file.

The NDMP related `omnirc` file variables are:

- **OB2NDMPFH**

Default value: Y

When the NDMP related variable `OB2NDMPFH` is set to Y, the NDMP Server file history is switched on. If it is switched on, Data Protector can restore single files. Note that the processing of the NDMP data catalog at backup time may take quite a significant amount of time.

When it is set to N, the NDMP Server file history is switched off. If it is switched off, you cannot browse and restore single files. This Data Protector variable overrides the file history settings on the NDMP Server every time a backup session is started.

- **OB2NDMPDIRECT**

Default value: Y

If this variable is set to Y, Data Protector uses the direct access restore functionality.

If this variable is set to N, the direct access restore functionality is disabled.

The `OB2NDMPDIRECT` variable, which defines the use of the direct access restore functionality, is used only if file history is switched on at backup time.

- **OB2NDMPMEMONLY**

Default value: 1

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

If the variable is set to 1, the NDMP Media Agent uses system physical memory only.

If the variable is set to 0, the NDMP Media Agent will store part of catalog information in file history swap files. For more information on file history swap files see the `OB2NDMPFHFILEOPT` variable description.

The `OB2NDMPMEMONLY` variable must be set to 0 if the number of files in a backup specification exceeds 5 million. Setting the variable to 0 will allow the NDMP Media Agent to handle backups of up to 20 million of files (in one backup specification) if the system has enough resources; for backup of 20 million files, approximately 1.9 GB of system memory and 2.8 GB of disk space is needed.

- **OB2NDMPCATQUESIZE**

Default value: 5

This variable is used to fine tune memory and disk consumption, and overall performance of the NDMP backup.

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

If enough system memory is available and the number of files in a backup specification is less than 20 million, the variable can be set to higher values. Increasing the value of the variable will increase the performance of NDMP backups to a certain extent.

The value of the variable represents the number of internal buffers that hold catalog information prior to writing it to file history swap files.

Memory allocation overhead in kilobytes can be calculated by multiplying the value of the variable by 512.

- **OB2NDMPFHFILEOPT**

This variable can be used to fine tune the file history swap files usage. See “Integration Concept” on page 127 for more information on file history swap files.

Default value:

- on Windows: `<Data_Protector_home>\tmp, 32, 1024, 10`
- on UNIX: `/var/opt/omni/tmp, 32, 1024, 10`

The first parameter sets the path to the file history swap files, the second parameter sets the maximum number of the file history swap files Data Protector creates on the NDMP client's disk, the third parameter sets the maximum size for a file history swap file (in MB), and the fourth parameter sets the minimum amount of disk space that must be left free on the NDMP client's disk (in MB).

The parameters are separated by commas. Note that it is also possible to specify several paths; in such a case semicolon is to be used to separate them. When all files for a specified path are full, the integration writes data to the files in the next specified directory (path). If the specified amount of disk space is reached during the backup, the backup session fails.

For example:

- on Windows: `C:\tmp, 32, 1024, 10; D:\tmp\tmp_1, 10, 1024, 40`
- on UNIX: `/tmp, 10, 1024, 50; /var/tmp, 5, 60, 20`

The file history swap files can grow very large. The formula below can be used to calculate approximate disk consumption by file history swap files:

$$EstConsumption = (NumofFiles + NumofDirs) \times (136 + AverageFileNameSize)$$

Where `NumofFiles` is the number of backed up files and `NumofDirs` is the number of backed up directories.

Table 3-4 on page 163 shows approximate disk consumption by file history swap files for 5, 10 and 20 million backed up files, when 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 3-4 **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

What Is Supported?

Only a limited set of standard Data Protector media management functions is supported. These functions are:

- Importing and exporting of media.
- Scanning of media.
- Initializing of media.

What Is Not Supported?

The following Data Protector media management functions are not supported:

- Verifying of backed up data.
- Copying of media.
- Dirty drive detection.

For more information, refer to the *HP OpenView Storage Data Protector Administrator's Guide* or online Help.

Troubleshooting

Error Messages

There are some Data Protector backup and restore error messages that explain the nature of errors that occur in the phase of establishing connection to the NDMP Server:

```
"NDMP:Error creating connection to NDMP server on <host>  
<port>"
```

```
"NDMP:Error connecting to NDMP server on <host> <port>"
```

```
"NDMP:Error authorizing to the server. User/Passwd"  
<user> <passwd>"
```

In addition to the Data Protector native error messages, messages reported by the NDMP Server are reported separately as described in the NDMP Messages section. For more information, see “Network Data Management Protocol (NDMP)” on page 130.

Catalog Data Does Not Fit

After the backup has finished on the NDMP Server, Data Protector writes catalog data to the media. The size of the catalog depends on the number of files that have been backed up - more files mean a bigger catalog. Since Data Protector does not control the flow of the backed up data, it is unknown how much space is left on the media. Therefore, the End of Media error can occur during the writing of catalog data. In this case, the catalog will still be stored in the IDB and restores will work as usual. However, the import of the medium will not be possible anymore.

Importing NDMP Media

Importing the media with an NDMP backup is not possible if the devices are attached to a standard Data Protector host. An error message is reported. To import the NDMP media, use an NDMP device.

Use of Media on Different Types of NDMP Servers

If the medium is used with one type of the NDMP Server, you cannot use it with another type of the NDMP Server. Data that was backed up to the first host can therefore not be restored with the other type of the NDMP Server.

Use of NDMP Dedicated Media Pools with Standard Non-NDMP Devices

NDMP-dedicated media pools cannot be used by standard Data Protector (non-NDMP) devices. If these media pools are used by such devices, Data Protector returns an error message and aborts the session. The same holds for the NDMP devices that use the media pools not dedicated to NDMP.

A Tape Remains in the Drive After the Scan Operation

If a tape remains in the drive after Data Protector performed a scan operation for this drive and reported its success, you should eject the tape manually and set the `OB2SCTLMOVETIMEOUT` omnirc file variable on the NDMP client to a higher value (for example, set it to 360000 or higher). Refer to *HP OpenView Storage Data Protector Administrator's Guide* for more information on how to set the omnirc file variables.

Glossary

access rights

See **user rights**.

ACSLS (*StorageTek specific term*)

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

Active Directory (*Windows specific term*)

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

AML (*EMASS/GRAU specific term*)

Automated Mixed-Media library.

application agent

A component needed on a client to back up or restore online database integrations.

See also **Disk Agent**.

application system (*ZDB specific term*)

A system the application or database runs on. The application or database data is located on source volumes.

See also **backup system** and **source volume**.

archived redo log (*Oracle specific term*)

Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- **ARCHIVELOG** - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “hot” backup can be performed only when the database is running in this mode.
- **NOARCHIVELOG** - The filled online redo log files are not archived.

See also **online redo log**.

archive logging (*Lotus Domino Server specific term*)

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

ASR Set

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

Glossary

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.

These files are stored as an ASR archive file on the Cell Manager (in `<Data_Protector_home>\Config\Server\dr\asr` on a Windows Cell Manager or in `/etc/opt/omni/server/dr/asr/` on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger

See **library**

autoloader

See **library**

BACKINT (*SAP R/3 specific term*)

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

backup API

The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The

interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain

This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device

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

backup generation

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

backup ID

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

Glossary

backup object

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

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.
- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines backup objects with identical client name and mount point.
- Type: backup object type (for example filesystem or Oracle).

backup owner

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

backup session

A process that creates a copy of data on storage media. The activities are

specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.

*See also **incremental backup** and **full backup**.*

backup set

A complete set of integration objects associated with a backup.

backup set (*Oracle specific term*)

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

backup specification

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

Glossary

Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system (*ZDB specific term*)

A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a 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 for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC (*EMC Symmetrix specific term*)

Business Continuity are processes that allow customers to access and manage

instant copies of EMC Symmetrix standard devices.

See also **BCV.**

BC (*HP StorageWorks Disk Array XP specific term*)

The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system. *See also* **HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.**

BC Process (*EMC Symmetrix specific term*)

A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuity Volumes to protect data on EMC Symmetrix standard devices.

See also **BCV.**

BC VA (*HP StorageWorks Virtual Array specific term*)

Business Copy VA allows you to

Glossary

maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

See also **HP StorageWorks Virtual Array LUN, application system, and backup system.**

BCV (*EMC Symmetrix specific term*) Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.

See also **BC** and **BC Process.**

Boolean operators

The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a

multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

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

BRARCHIVE (*SAP R/3 specific term*)

An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.

See also **SAPDBA, 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 **SAPDBA, BRARCHIVE** and **BRRESTORE.**

BRRESTORE (*SAP R/3 specific term*)

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

- Database data files, control files, and online redo log files saved with BRBACKUP

Glossary

- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.

See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM

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

CA (*HP StorageWorks Disk Array XP specific term*)

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

See also **BC** (*HP StorageWorks Disk*

Array XP specific term), **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP (*StorageTek specific term*)

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

catalog protection

Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.

See also **data protection**.

CDB

The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.

See also **MMDB**.

CDF file (*UNIX specific term*)

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

Glossary

cell

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

Cell Manager

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

centralized licensing

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

See also MoM.

Centralized Media Management Database (CMMDB)

See CMMDB.

channel (*Oracle specific term*)

An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which

performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is 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 filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk

Glossary

discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or client system

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

cluster-aware application

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

CMD Script for OnLine Server

(Informix specific term)

Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the

robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Registration Database

(Windows specific term)

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

command-line interface

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

Command View (CV) EVA *(HP*

StorageWorks EVA specific term)

The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Glossary

Storage Management Appliance, and is accessed by a Web browser.

See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency

See **Disk Agent concurrency**.

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

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

CRS

The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.

CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM

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

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

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

data protection

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

See also **catalog protection**.

Data Protector Event Log

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group

Glossary

membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream

Sequence of data transferred over the communication channel.

database library

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

Dboject (*Informix specific term*)

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

DC directory

The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB,

which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the `<Data_Protector_home>\db40` directory on a Windows Cell Manager and in the `/var/opt/omni/server/db40` directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF

The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup.

delta backup

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

device

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

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one

Glossary

device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (*EMC Symmetrix specific term*)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup

An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.

See **incremental backup**.

differential backup (*MS SQL specific term*)

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

See also **backup types**.

differential database backup

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

direct backup

A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.

See also **XCOPY engine**.

Glossary

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.

Directory Store (DS) (*Microsoft Exchange specific term*)

A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.

See also **Information Store (MDB)**.

disaster recovery

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

Disk Agent

A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency

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

disk discovery

The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group (*Veritas Volume Manager specific term*)

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

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You

Glossary

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 System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network

(Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller

A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image

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

DR OS

A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

Glossary

drive

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

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)

See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*

An Informix configuration file that resides in the <INFORMIXDIR>\etc directory (on HP-UX) or <INFORMIXDIR>/etc directory (on Windows) and is called ixbar.<server_id>, where <INFORMIXDIR> is the OnLine 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.

Enterprise Backup Environment

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

Event Logs

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

exchanger

Also referred to as SCSI Exchanger. See also **library**.

exporting media

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

Glossary

Extensible Storage Engine (ESE)

(Microsoft Exchange Server 2000/2003 specific term)

A database technology used as a storage system for information exchange in Microsoft Exchange Server 2000/2003.

failover

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

FC bridge

See **Fibre Channel bridge**

Fibre Channel

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

Fibre Channel bridge

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

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

file depot

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

file jukebox device

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

file library device

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

File Replication Service (FRS)

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

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector

Glossary

retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

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

first level mirror (*HP StorageWorks Disk Array XP specific term*)

HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.

See also **Primary Volume**, and **MU numbers**.

fnames.dat

The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting

A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are

not formatted until the protection expires or the media are unprotected/recycled.

free pool

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

full backup

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

See also **backup types**.

full database backup

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

full mailbox backup

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

global options file

A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the

Glossary

<Data_Protector_home>\Config\Server\Options directory on Windows systems.

group (*Microsoft Cluster Server specific term*)

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

GUI

A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery (*Microsoft Exchange Server specific term*)

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

heartbeat

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

Hierarchical Storage Management (HSM)

A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to

less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file

A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and <Data_Protector_home>\Config\Server\holidays on the Windows Cell Manager.

host backup

See **client backup with disk discovery**.

hosting system

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

HP ITO

See **OVO**.

HP OpC

See **OVO**.

HP OpenView SMART Plug-In (SPI)

A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView

Glossary

SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO

See **OVO**.

HP StorageWorks Disk Array XP LDEV

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **replica**.

HP StorageWorks EVA Agent (legacy)

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.

See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.

See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.

See also **HP StorageWorks EVA SMI-**

Glossary

S Agent and Command View (CV) EVA.

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also BC VA and replica.

HP VPO
See OVO.

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

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media

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

incremental backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also backup types.

incremental backup (*Microsoft Exchange Server specific term*)

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

incremental mailbox backup

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

incremental1 mailbox backup

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

Glossary

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

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

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

incremental restore (*EMC Symmetrix specific term*)

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

written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

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 2000/2003 specific term*)

The Microsoft Exchange Server 2000/2003 service that is responsible for storage management. Information Store in Microsoft Exchange Server 2000/2003 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**.

Glossary

Information Store (*Microsoft Exchange Server 5.5 specific term*)

This is the default message store provider for the Microsoft Exchange Server 5.5. Information Store consists of the following stores:

- public information store
- private information store
- personal folder store
- offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within an Exchange Server 5.5 organization, even if multiple Exchange Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Exchange Server 5.5.

See also **Directory Store (DS)**.

initializing

See **formatting**.

Installation Server

A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is

used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)

A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.

See also **replica**, **zero downtime backup (ZDB)**, **ZDB to disk**, and **ZDB to disk+tape**.

integrated security (*MS SQL specific term*)

Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL

Glossary

Server are referred to as trusted connections. Only trusted connections are allowed.

integration object

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

Internet Information Server (IIS)

(Windows specific term)

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

IP address

Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*

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

ITO

See OVO.

jukebox

See library.

jukebox device

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

Key Management Service *(Microsoft Exchange Server 2000/2003 specific term)*

The Microsoft Exchange Server 2000/2003 service that provides encryption functionality for enhanced security. *See also Information Store and Site Replication Service.*

LBO *(EMC Symmetrix specific term)*

A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library

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

lights-out operation or **unattended operation**

A backup or restore operation that takes

Glossary

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.

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 UNIX specific term*)

A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the OnLine 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`.

Glossary

logging level

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

logical-log files

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

login ID (*MS SQL Server specific term*)

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

login information to the Oracle

Target Database (*Oracle and SAP R/3 specific term*)

The format of the login information is <user_name>/<password>@<service>, where:

- <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have been granted Oracle SYSDBA or SYSOPER rights.
- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to 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>/

Glossary

<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 has to be the owner of the Oracle Recovery (Oracle) Catalog.

Lotus C API (*Lotus Domino Server specific term*)

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

LVM

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

Magic Packet

See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term*)

The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of

personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server 2000/2003 specific term*)

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

Main Control Unit (MCU) (*HP StorageWorks Disk Array XP specific term*)

An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)

See **Enterprise Cell Manager**.

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup

Glossary

medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI (*Microsoft Exchange specific term*)

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

media allocation policy

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

media condition

The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors

The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID

A unique identifier assigned to a medium by Data Protector.

media label

A user-defined identifier used to describe a medium.

media location

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

media management session

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

media pool

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

media set

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

Glossary

media type

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

media usage policy

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

merging

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

MFS

The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated. *See also* **VBFS**.

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 7.0/2000

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

Glossary

sets.

See also **shadow copy, shadow copy provider, writer.**

mirror (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

See **target volume.**

mirror rotation (*HP StorageWorks Disk Array XP specific term*)

See **replica set rotation.**

MMD

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

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.

See also **CMMDB, CDB.**

MoM

Several cells can be grouped together and managed from a central cell. The

management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request

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

mount point

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

MSM

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

MU number (*HP StorageWorks Disk Array XP specific term*)

A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.

See also **first level mirror.**

multi-drive server

A license that allows you to run an unlimited number of Media Agents on a

Glossary

single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat

An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device

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

object

See **backup object**

object copy

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

object copy session

A process that creates an additional copy of the backed up data on a different media set. During an object copy

session, the selected backed up objects are copied from the source to the target media.

object copying

The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID (*Windows specific term*)

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

object mirror

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

object mirroring

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

offline backup

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

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use

Glossary

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. Recovery of the Cell Manager is always offline.

offline redo log

See **archived redo log**

OmniStorage

Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar (*Informix specific term*)

A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dobjects and track instances of dobjects through multiple backups.

onbar utility (*Informix specific term*)

The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG (*Informix specific term*)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file `<INFORMIXDIR>/etc/onconfig` (on HP-UX) or `<INFORMIXDIR>\etc\onconfig` (on Windows).

Glossary

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 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**.

OnLine Server (*Informix specific term*)

Refers to INFORMIX-OnLine Dynamic Server.

OpC

See **OVO**.

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

Glossary

files are restored from a backup even if they are older than existing files.

See also **merging**.

OVO

HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.

See also **merging**.

ownership

The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: `root.sys@<Cell Manager>`, and for the Windows Cell Manager, the user that was specified during the

installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file

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

`<Data_Protector_home>\Config\Server\dr\p1s` directory on a Windows Cell Manager or in `/etc/opt/omni/server/dr/p1s` directory on a UNIX Cell Manager with the filename `recovery.p1s`.

package (*MC/ServiceGuard and Veritas Cluster specific term*)

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

pair status (*HP StorageWorks Disk Array XP specific term*)

A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

Glossary

- **COPY** - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- **PAIR** - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- **SUSPENDED** - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism

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

physical device

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

post-exec

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

See also **pre-exec**.

pre- and post-exec commands

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

prealloc list

A subset of media in a media pool that specifies the order in which media are used for backup.

Glossary

pre-exec

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

See also **post-exec**.

Primary Volume (P-VOL) *(HP*

StorageWorks Disk Array XP specific term)

Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.

See also **Secondary Volume (S-VOL)**.

Private Information Store *(Microsoft*

Exchange Server 5.5 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.

protection

See **data protection** and also **catalog protection**.

public folder store *(Microsoft*

Exchange Server 2000/2003 specific term)

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

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

public/private backed up data

When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID

Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP*

StorageWorks Disk Array XP specific term)

The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks*

Disk Array XP specific term)
The RAID Manager XP application

Glossary

provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup

See **disk image backup**.

RCU (*HP StorageWorks specific term*)

The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS

Relational Database Management System.

RDF1/RDF2 (*EMC Symmetrix specific term*)

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

RDS

The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog (*Oracle specific term*)

A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database (*Oracle specific term*)

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

RecoveryInfo

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume,

Glossary

and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) (*Oracle specific term*)

An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle

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

redo log (*Oracle specific term*)

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

Remote Control Unit (RCU) (*HP StorageWorks Disk Array XP specific term*)

The Remote Control Unit (RCU) acts as a slave of an MCU in a CA

configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database (*Windows specific term*)

A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point (*Windows specific term*)

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

replica (*ZDB specific term*)

An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/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

Glossary

snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object is replicated.

See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

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 session

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

RMAN (*Oracle specific term*)

See **Recovery Manager**.

RSM

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

RSM (*Windows specific term*)

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

SAPDBA (*SAP R/3 specific term*)

An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan

A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the

Glossary

device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler

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

Secondary Volume (S-VOL) (*HP StorageWorks Disk Array XP specific term*)

Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. *See also* **Primary Volume (P-VOL)**.

session

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

session ID

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

session key

This environment variable for the Pre- and Post-exec script is a Data Protector

unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy (*MS VSS specific term*)

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

See also **Microsoft Volume Shadow Copy service**.

shadow copy provider (*MS VSS specific term*)

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

See also **shadow copy**.

shadow copy set (*MS VSS specific term*)

A collection of shadow copies created at the same point in time.

See also **shadow copy**.

Glossary

shared disks

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

SIBF

The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service (*Microsoft Exchange Server 2000/2003 specific term*)

The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.

See also **Information Store** and **Key Management Service**.

slot

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

SMB

See **split mirror backup**.

SMBF

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

snapshot (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.

See also **replica** and **snapshot creation**.

snapshot backup (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

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

snapshot creation (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created

Glossary

at one particular point-in-time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.

See also **snapshot**.

source (R1) device (*EMC Symmetrix specific term*)

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

See also **target (R2) device**.

source volume (*ZDB specific term*)

A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are:
-A matrix in which some or much of the data contains zeros
-files from image applications
-high-speed databases
If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone,

of the contents of the source volumes.

See also **replica** and **split mirror creation**.

split mirror backup (*EMC Symmetrix specific term*)

See **ZDB to tape**.

split mirror backup (*HP StorageWorks Disk Array XP specific term*)

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

split mirror creation (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

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

See also **split mirror**.

split mirror restore (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete

Glossary

sessions can be restored using this method.

See also **ZDB to tape, ZDB to disk+tape, and replica.**

sqlhosts file (*Informix specific term*)

An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file

The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF (*EMC Symmetrix specific term*)

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

SSE Agent (*HP StorageWorks Disk Array XP specific term*)

A Data Protector software module that

executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file

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

st.conf file

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

stackers

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

Glossary

standalone file device

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

standard security (*MS SQL specific term*)

Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.

See also **integrated security**.

Storage Group

(*Microsoft Exchange Server 2000/2003 specific term*)

A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library

(*StorageTek specific term*)

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

storage volume (*ZDB specific term*)

A storage volume represents an object that may be presented to an operating system or some other entity (for

example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover

See **failover**

Sybase Backup Server API (*Sybase specific term*)

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

Sybase SQL Server (*Sybase specific term*)

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

Symmetrix Agent (SYMA) (*EMC Symmetrix specific term*)

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

Glossary

System Backup to Tape (*Oracle specific term*)

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

system databases (*Sybase specific term*)

The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybssystemprocs)
- model database (model).

system disk

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

system partition

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

System State (*Windows specific term*)

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

system volume/disk/partition

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

SysVol (*Windows specific term*)

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

tablespace

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

Glossary

tapeless backup (*ZDB specific term*)
See **ZDB to disk**.

target database (*Oracle specific term*)
In RMAN, the target database is the database that you are backing up or restoring.

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

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

target volume (*ZDB specific term*)
A storage volume to which data is replicated.

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

thread (*MS SQL Server 7.0/2000 specific term*)
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder (*EMC Symmetrix specific term*)
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA (*Oracle and SAP R/3 specific term*)
A network configuration file that

Glossary

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

transaction

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

transaction backup

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

transaction backup (*Sybase and SQL specific term*)

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

transaction log backup

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

transaction log files

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

transaction logs (*Data Protector specific term*)

Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table (*Sybase specific term*)

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

transportable snapshot (*MS VSS specific term*)

A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.

See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file (*Novell NetWare specific term*)

A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

Glossary

unattended operation

See lights-out operation.

user account

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

user disk quotas

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

user group

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

user profile (*Windows specific term*)

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

user rights

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

vaulting media

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

VBFS (*OmniStorage specific term*)

A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute'

Glossary

information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify

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

Virtual Controller Software (VCS)

(HP StorageWorks EVA specific term)

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

Virtual Device Interface (MS SQL Server 7.0/2000 specific term)

This is a SQL Server 7.0/2000 programming interface that allows fast backup and restore of large databases.

virtual disk (HP StorageWorks EVA specific term)

A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array

snapshot functionality.

See also **source volume** and **target volume**.

virtual server

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

volser (ADIC and STK specific term)

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

volume group

A unit of data storage in an LVM system. A volume group can consist of 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

Glossary

mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service

See **Microsoft Volume Shadow Copy service**.

VPO

See **OVO**.

VSS

See **Microsoft Volume Shadow Copy service**.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN

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

Web reporting

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

wildcard character

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

Windows CONFIGURATION

backup

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

Windows Registry

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

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.

Glossary

writer

(MS VSS specific term)

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

XBSA interface *(Informix specific term)*

The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*

A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also **direct backup**.

ZDB

See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*

A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.

See also **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.

See also **zero downtime backup (ZDB)**, **ZDB to tape**, **ZDB to disk+tape**, **instant recovery**, and **replica set rotation**.

ZDB to disk+tape *(ZDB specific term)*

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored

Glossary

using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.

See also **zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.**

ZDB to tape (*ZDB specific term*)

A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.

See also **zero downtime backup (ZDB), ZDB to disk, instant recovery, ZDB to disk+tape, and replica.**

zero downtime backup (ZDB)

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

See also **ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.**

A

- advantages
 - NDMP integration, 123
 - NNM integration, 99
 - Sybase integration, 4
- architecture
 - Sybase integration, 7

B

- backing up NDMP, 149–153
 - backup options, 152
 - backup specification, creating, 149
- backing up NNM, 109–112
 - backup flow, 102
 - backup options, 107
 - backup specification, creating, 105
 - backup templates, creating, 105
 - backup types, 111
 - full backup, 111
 - incremental backup, 111
 - scheduled backup, 110
 - starting interactive backup, using GUI, 112
- backing up Sybase, 45–53
 - backup flow, 8
 - backup methods, 45
 - backup options, 34
 - backup specification, creating, 31
 - backup types, 3
 - database objects backup, 31
 - full backup, 3
 - problems on UNIX, 87
 - problems on Windows, 80
 - scheduled backup, 47
 - scheduled backup, example, 48
 - starting interactive backup, using CLI, 50
 - starting interactive backup, using GUI, 50
 - trans backup, 3
 - transaction logs backup, 50, 51
- backup flow
 - NNM integration, 102
 - Sybase integration, 8
- backup methods
 - Sybase integration, 45
- backup options
 - NDMP integration, 152
 - NNM integration, 107
 - Sybase integration, 34
- backup problems
 - Sybase integration, on UNIX, 87

- Sybase integration, on Windows, 80
- backup specification
 - NDMP integration, creating, 149
 - NNM integration, creating, 105
 - NNM integration, scheduling, 110
 - Sybase integration, creating, 31
 - Sybase integration, scheduling, 47
- backup templates
 - NNM integration, creating, 105
- backup types
 - NNM integration, 111
 - Sybase integration, 3

C

- Celerra NAS device *See* NDMP
- checking
 - Sybase configuration, 28
- concepts
 - NDMP integration, 127
 - NNM integration, 102
 - Sybase integration, 7
- configuration file
 - Sybase integration, 9–13
 - Sybase integration, parameters, 11
- configuration problems
 - Sybase integration, on UNIX, 84
 - Sybase integration, on Windows, 79
- configuring NDMP, 135–145
 - backup, 149
 - backup device, 141
 - backup specification, 149
 - creating media pool, 140
 - importing NDMP Server, 138
 - library configurations, 135
 - overview, 138
 - prerequisites, 135
- configuring NNM, 104–108
 - backup, 104
 - backup specification, 105
 - backup templates, 105
 - overview, 104
- configuring Sybase, 14–41
 - backup, 30
 - backup specification, 31
 - checking configuration, 28
 - overview, 14
 - prerequisites, 14
 - problems on UNIX, 84
 - problems on Windows, 79

Index

- server, 20
- users, 17
- conventions, ix
- creating
 - NDMP backup specification, 149
 - NNM backup specification, 105
 - NNM backup templates, 105
 - Sybase backup specification, 31

D

- disaster recovery
 - NNM integration, 113
 - Sybase integration, 73

E

- environment variables
 - NDMP integration, 158
- examples
 - Sybase integration, scheduling, 48
 - Sybase restore, 60

F

- file history swap file
 - NDMP integration, 129
- full backup
 - NNM integration, 111
 - Sybase integration, 3

I

- incremental backup
 - NNM integration, 111
- interactive backup
 - NNM integration, 111
 - Sybase integration, 50

L

- limitations
 - NDMP integration, 125
 - NNM integration, 101
 - Sybase integration, 6

M

- monitoring
 - NNM sessions, 115
 - Sybase sessions, 75

N

- NDMP, 129, 130
- NDMP backup, 149–153
 - backup options, 152
 - backup specification, creating, 149
- NDMP configuration, 135–145
 - backup, 149
 - backup device, 141
 - backup specification, 149
 - creating media pool, 140
 - importing NDMP Server, 138
 - library configurations, 135
 - overview, 138
 - prerequisites, 135
- NDMP integration
 - advantages, 123
 - backing up NDMP, 149–153
 - concepts, 127
 - configuring NDMP, 135–145
 - environment variables, 158
 - file history swap file, 129
 - limitations, 125
 - NDMP, 129, 130
 - omnirc file variables, 160
 - overview, 123
 - prerequisites, 125
 - restoring NDMP, 154–157
 - troubleshooting NDMP, 165–166
- NDMP restore, 154–157
 - direct access restore, 155
 - direct access restore, limitations, 156
 - direct access restore, prerequisites, 156
 - limitations, 154
 - procedure, 154
 - using another device, 157
- NDMP troubleshooting, 165–166
- NetApp NAS device *See* NDMP
- Network Data Management Protocol *See* NDMP
- NNM backup, 109–112
 - backup flow, 102
 - backup options, 107
 - backup specification, creating, 105
 - backup templates, creating, 105
 - backup types, 111
 - full backup, 111
 - incremental backup, 111
 - scheduling, 110
 - starting interactive backup, using GUI, 112

- NNM configuration, 104–108
 - backup, 104
 - backup specification, 105
 - backup templates, 105
 - overview, 104
 - NNM integration
 - advantages, 99
 - backing up NNM, 109–112
 - concepts, 102
 - configuring NNM, 104–108
 - disaster recovery, 113
 - limitations, 101
 - monitoring sessions, 115
 - NNMpost.ovpl, 103
 - NNMpre.ovpl, 103
 - NNMScript.exe, 103
 - overview, 99
 - prerequisites, 101
 - restoring NNM, 113–114
 - troubleshooting NNM, 117–120
 - NNM restore, 113–114
 - disaster recovery, 113
 - NNM troubleshooting, 117–120
 - NNMpost.ovpl
 - NNM integration, 103
 - NNMpre.ovpl
 - NNM integration, 103
 - Sybase integration, 103
 - NNMScript.exe
 - NNM integration, 103
 - O**
 - omnirc file variables
 - NDMP integration, 160
 - overview
 - NDMP integration, 123
 - NNM integration, 99
 - Sybase integration, 3
 - P**
 - parameters
 - Sybase integration, configuration file, 11
 - prerequisites
 - NDMP integration, 125
 - NNM integration, 101
 - Sybase integration, 5
 - R**
 - restore flow
 - Sybase integration, 8
 - restore problems
 - Sybase integration, on UNIX, 91
 - Sybase integration, on Windows, 82
 - restore types
 - Sybase integration, 3
 - restoring NDMP, 154–157
 - direct access restore, 155
 - direct access restore, limitations, 156
 - direct access restore, prerequisites, 156
 - limitations, 154
 - procedure, 154
 - using another device, 157
 - restoring NNM, 113–114
 - disaster recovery, 113
 - restoring Sybase, 54–74
 - disaster recovery, 73
 - examples, 60
 - finding needed information, 54
 - problems on UNIX, 91
 - problems on Windows, 82
 - procedure, 57
 - restore flow, 8
 - restore types, 3
 - syb_tool, 65
 - using another device, 72
 - running backup *See* starting backup
 - S**
 - scheduling backup
 - NNM integration, 110
 - Sybase integration, 47
 - server
 - NDMP integration, importing, 138
 - Sybase integration, configuring, 20
 - starting backup
 - NNM integration, interactively, 112
 - NNM integration, using GUI, 112
 - Sybase integration, interactively, 50
 - Sybase integration, using CLI, 50
 - Sybase integration, using GUI, 50
 - Sybase backup, 45–53
 - backup flow, 8
 - backup methods, 45
 - backup options, 34
 - backup specification, creating, 31
 - backup types, 3
 - database objects backup, 31
 - full backup, 3
-

- problems on UNIX, 87
- problems on Windows, 80
- scheduling, 47
- scheduling, example, 48
- starting interactive backup, using CLI, 50
- starting interactive backup, using GUI, 50
- trans backup, 3
- transaction logs backup, 50, 51
- Sybase configuration, 14–41
 - backup, 30
 - backup specification, 31
 - checking, 28
 - overview, 14
 - prerequisites, 14
 - problems on UNIX, 84
 - problems on Windows, 79
 - server, 20
 - users, 17
- Sybase configuration file, 9–13
- Sybase integration
 - advantages, 4
 - architecture, 7
 - backing up Sybase, 45–53
 - concepts, 7
 - configuration file, 9–13
 - configuring Sybase, 14–41
 - disaster recovery, 73
 - limitations, 6
 - monitoring sessions, 75
 - NNMpre.ovpl, 103
 - overview, 3
 - prerequisites, 5
 - restoring Sybase, 54–74
 - testing, 107
 - testing, using CLI, 42
 - testing, using GUI, 42
 - troubleshooting Sybase, 78–95
 - util_cmd, 11, 65
 - viewing sessions, 76
- Sybase restore, 54–74
 - disaster recovery, 73
 - examples, 60
 - finding needed information, 54
 - problems on UNIX, 91
 - problems on Windows, 82
 - procedure, 57
 - restore types, 3
 - syb_tool, 65
 - using another device, 72
- Sybase troubleshooting, 78–95
 - backup problems, on UNIX, 87
 - backup problems, on Windows, 80
 - configuration problems, on UNIX, 84
 - configuration problems, on Windows, 79
 - restore problems, on UNIX, 91
 - restore problems, on Windows, 82
- T**
- testing
 - Sybase integration, 107
 - Sybase integration, using CLI, 42
 - Sybase integration, using GUI, 42
- trans backup
 - Sybase integration, 3
- transaction logs backup
 - Sybase integration, 50, 51
- troubleshooting NDMP, 165–166
- troubleshooting NNM, 117–120
- troubleshooting Sybase, 78–95
 - backup problems, on UNIX, 87
 - backup problems, on Windows, 80
 - configuration problems, on UNIX, 84
 - configuration problems, on Windows, 79
 - restore problems, on UNIX, 91
 - restore problems, on Windows, 82
- typographical conventions, ix
- U**
- users
 - Sybase integration, configuring, 17
- util_cmd
 - Sybase integration, 11, 65
- V**
- viewing
 - Sybase sessions, 76