

HP OpenView Storage Data Protector UNIX Integration Guide

Manual Edition: May 2003



Manufacturing Part Number: B6960-90082

Release A.05.10

© Copyright Hewlett-Packard Development Company, L.P.2003.

Legal Notices

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

Warranty. A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend. All rights are reserved. No part of this document may be photocopied, reproduced, or translated to another language without the prior written consent of Hewlett-Packard Company. The information contained in this document is subject to change without notice.

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

Hewlett-Packard Company
United States of America

Copyright Notices. ©Copyright 1983-2003 Hewlett-Packard Development Company, L.P. all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

©Copyright 1979, 1980, 1983, 1985-93 Regents of the University of California

This software is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California.

©Copyright 1986-1992 Sun Microsystems, Inc.

©Copyright 1985-86, 1988 Massachusetts Institute of Technology

©Copyright 1989-93 The Open Software Foundation, Inc.

©Copyright 1986-1997 FTP Software, Inc. All rights reserved

©Copyright 1986 Digital Equipment Corporation

©Copyright 1990 Motorola, Inc.

©Copyright 1990, 1991, 1992 Cornell University

©Copyright 1989-1991 The University of Maryland

©Copyright 1988 Carnegie Mellon University

©Copyright 1991-1995 by Stichting Mathematisch Centrum,
Amsterdam, The Netherlands

©Copyright 1999, 2000 Bo Branten

Trademark Notices. UNIX® is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

X Window System is a trademark of the Massachusetts Institute of Technology.

Motif is a trademark of the Open Software Foundation, Inc. in the U.S. and other countries.

Windows NT™ is a U.S. trademark of Microsoft Corporation. Microsoft®, MS-DOS®, Windows® and MS Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle®, SQL*Net®, and Net8® are registered U.S. trademarks of Oracle Corporation, Redwood City, California. Oracle Reports™, Oracle8™, Oracle8 Server Manager™ and Oracle8 Recovery Manager™ are trademarks of Oracle Corporation, Redwood City, California.

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

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

ARM® is a registered trademark of ARM Limited.

X/Open® is a registered trademark, and the X device is a trademark of X/Open Company Ltd. in the UK and other countries.

VisiCalc® is a U.S. registered trademark of Lotus Development Corp.

HP-UX Release 11.00 and later (in both 32- and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

Netscape and Netscape Navigator are U.S. trademarks of Netscape Communications Corporation.

OpenView® is a registered U.S. trademark of Hewlett-Packard Company.

© 2003 Bristol Technology, Inc., Bristol Technology, Wind/U, HyperHelp and Xprinter are registered trademarks of Bristol Technology Inc.

Other reserved names are trademarks of the respective companies.

1. Integrating Oracle8/9 and Data Protector

In This Chapter	2
Overview	3
Prerequisites and Limitations	5
Prerequisites	5
Limitations	6
Integration Concept	8
Data Protector Oracle8/9 Configuration Files	13
Setting, Retrieving and Listing Data Protector Oracle8/9 Configuration Files’ Parameters Using the CLI	15
Installing and Upgrading the Oracle8/9 Integration	18
Configuring the Integration	20
Linking Oracle8/9 with the Data Protector Database Library	21
Configuring an Oracle8/9 User in Data Protector	27
Configuring the Oracle8/9 Server	29
Configuring an Oracle8/9 Backup	37
Creating a New Template	37
Creating a Backup Specification	38
Oracle8/9 Backup Options	42
Testing the Integration	49
Backing Up an Oracle8/9 Database	52
Scheduling a Backup	55
Starting an Interactive Backup	57
Using the Oracle8/9 Recovery Manager (RMAN)	60
Restoring Oracle8/9 Databases	66
Restoring Oracle8/9 Using the Data Protector Restore GUI for Oracle	66
Restore and Recovery Options	76
Restoring Oracle8/9 Using RMAN	80
Disaster Recovery	81
Monitoring an Oracle8/9 Backup and Restore	83
Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration	84
Removing the Data Protector Oracle8/9 Integration Link on HP-UX Systems	84
Removing the Data Protector Oracle8/9 Integration Link on Solaris and other UNIX Systems	85
Oracle8i and Oracle9i RMAN Metadata and Data Protector Media Management Database Synchronization	87
Configuring the Integration as Cluster-Aware	89
Installation and Configuration	89
Backup and Restore	90

Contents

Troubleshooting	92
Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration.....	92
Setting Up the Environment Variables	93
Checking Prerequisites Related to the Oracle8/9 Side of the Integration.....	93
Configuration Problems	97
Backup Problems	100
Restore Problems	103
Examples of an Oracle8/9 Database Restore	107
Preparing the Oracle8/9 Database for Restore	107
Examples of Oracle8/9 Database Restore	109

2. Integrating SAP R/3 and Data Protector

In This Chapter	118
Overview.....	119
Prerequisites and Limitations.....	121
Integration Concept	123
Data Protector SAP R/3 Configuration File	132
Setting, Retrieving, Listing and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI.....	135
Installing and Upgrading the Data Protector SAP R/3 Integration	138
Configuring the Integration	140
Configuring an SAP R/3 User in Data Protector.....	140
Configuring an SAP R/3 Database Server.....	142
Configuring an SAP R/3 Backup.....	150
Creating a New Template.....	150
Creating a Data Protector SAP R/3 Backup Specification	150
SAP R/3 Backup Options	155
Creating or Modifying the Parameter File on the SAP R/3 Database Server.....	160
Backing Up Using Recovery Manager.....	161
Manual Balancing of Files into Subsets	163
Creating an SAP /R3 Backup Specification for Manual Balancing	163
Testing the Integration.....	165
Backing Up an SAP R/3 Database	167
Scheduling a Backup	168
Starting an Interactive Backup	171
Using SAP R/3 Commands.....	173
Restoring an SAP R/3 Database	174
Finding Information Needed for Restore.....	174

Using the Data Protector GUI	175
Using the Data Protector CLI	176
Using SAP R/3 Commands	177
Using Another Device	177
Disaster Recovery	178
Monitoring an SAP R/3 Backup and Restore	180
Configuring the Integration as Cluster-Aware	181
Installation and Configuration	181
Backup and Restore	182
Troubleshooting	183
Using Oracle8 After Removing the Data Protector Oracle8 Integration	183
Prerequisites Concerning the Oracle Side of the Integration	184
Prerequisites on the SAP R/3 Side of the Integration	187
Configuration Problems	188
Backup Problems	190
Restore Problems	193
Example of SAP R/3 Database Restore	197
Preparing the SAP R/3 Database for Restore	197
Examples of SAP R/3 Database Restore	199

3. Integrating IBM DB2 UDB and Data Protector

In This Chapter	208
Overview	209
Prerequisites and Limitations	211
Integration Concept	213
Installing the DB2 Integration	219
Configuring the Integration	220
Configuring a DB2 User	220
Configuring a DB2 Instance	220
Configuring a DB2 Backup	222
DB2 Specific Backup Options	226
Testing the Integration	228
Backing Up a DB2 Database	230
Scheduling an Existing Backup Specification	231
Running an Interactive Backup Using the Data Protector GUI	233
Running an Interactive Backup Using the Data Protector CLI	234
Restoring a DB2 Database	236
Restoring a DB2 Object Using the Data Protector GUI	236

Contents

Restore Options	241
Restoring a DB2 Object Using the Data Protector CLI	243
Monitoring a DB2 Backup and Restore	248
Troubleshooting	250
General Troubleshooting	250
Backup Problems	251
Restore Problems	253

4. Integrating Sybase and Data Protector

In This Chapter	256
Organization of This Chapter	256
Overview	257
Prerequisites	259
Limitations	260
Integration Concepts	261
Installing and Upgrading the Integration	263
Configuring the Integration	265
Before You Begin Configuring	265
Configuring a Sybase User in Data Protector	268
Configuring a Sybase Server	270
Configuring a Sybase Backup	279
Testing the Integration	291
Using the Data Protector GUI	291
Using the Data Protector CLI	291
What Happens?	292
Backing Up a Sybase Database	294
Scheduling an Existing Backup Specification	296
Running an Interactive Backup	299
Backing Up Using Sybase Commands	301
Restoring a Sybase Database	303
The Data Protector omnidb Command	303
Using the load Command	305
The Data Protector syb_tool Command	310
Restoring Using Another Device	316
Disaster Recovery	317
Monitoring a Sybase Backup and Restore Session	319
Sybase Character Sets	320
Configuring the Integration as Cluster-Aware	321

Installation and Configuration	321
Backup and Restore	322
Troubleshooting	323
Before You Begin	323
Configuration Problems	323
Backup Problems	326
Restore Problems	331

5. Integrating Informix and Data Protector

In This Chapter	338
Organization of This Chapter	338
Overview	339
Prerequisites	342
Limitations	343
Integration Concepts	344
Installing and Upgrading the Integration	346
Configuring the Integration	348
Before You Begin Configuring	348
Configuring an Informix User in Data Protector	350
Configuring an OnLine Server	353
Configuring an Informix Backup	360
Testing the Integration	374
Using the Data Protector GUI	374
Using the Data Protector CLI	375
What Happens?	376
Backing Up an Informix Database	378
Scheduling an Existing Backup Specification	381
Running an Interactive Backup	383
Using Informix Commands	386
Using the Informix log_full.sh Script	388
On-Demand and Continuous Backups	388
Restoring an Informix Database	390
The Data Protector omnldb Command	390
Finding Information for Restore	392
Using the Data Protector GUI	394
Using Informix Commands	398
To Another OnLine Server	400
Using Another Device	401

Contents

Disaster Recovery	402
Monitoring an Informix Backup and Restore	403
Configuring the Integration as Cluster-Aware	404
Installation and Configuration	404
Backup and Restore	404
Troubleshooting	406
Before You Begin	406
Configuration Problems	406
Backup Problems	408
Restore Problems	417

6. Integrating the NDMP Server and Data Protector

In This Chapter	422
Overview	423
Prerequisites and Limitations	425
Integration Concept	427
Network Data Management Protocol (NDMP)	429
Installing the NDMP Server Integration	434
Configuring the Integration	435
Supported Configurations	436
Configuration Procedure	438
Importing the NDMP Server Host	439
Creating a Media Pool	441
Configuring an NDMP Backup Device	441
Network Appliance Configuration	446
EMC Celerra Configuration	448
Backing Up the NDMP Server Data	449
NDMP Environment Variables	453
Restoring the NDMP Server Data	454
Direct Access Restore	455
Restore Using Another Device	457
Media Management	458
The NDMP Related omnirc File Variables	459
Troubleshooting	460
Error Messages	460
Catalog Data Does Not Fit	460
Importing NDMP Media	460
Use of Media on Different Types of NDMP Servers	461

Use of NDMP Dedicated Media Pools with Standard Non-NDMP Devices	461
A Tape Remains in the Drive After the Scan Operation	461

7. Integrating Network Node Manager and Data Protector

In This Chapter	464
Overview	465
Prerequisites and Limitations	466
Prerequisites	466
Limitations	466
Integration Concept	467
Installing the NNM Integration	469
Configuring an NNM Backup	470
Tasks for the NNM Administrator	470
Creating a New Template	471
Creating a Backup Specification	471
NNM Backup Options	473
Testing the Integration	473
Backing Up an NNM Database	474
Scheduling a Backup	475
Starting an Interactive Backup	476
Restoring NNM	477
Disaster Recovery	477
Monitoring an NNM Backup and Restore	479
Troubleshooting	480
Error and Warning Messages	480
Backup/Restore Problems	481

8. Integrating Lotus Domino R5 Server and Data Protector

In This Chapter	484
Overview	485
Prerequisites and Limitations	488
Integration Concept	489
Installing the Lotus Notes Integration	492
Configuring the Integration	494
Configuring the Lotus Domino R5 Server	494
Configuring Data Protector Lotus Notes Integration	497
Testing the Integration	502
Backing Up Lotus Domino R5 Server	506

Contents

Configuring a Lotus Domino R5 Server Backup	508
Running an Online Backup	514
Restoring Lotus Domino R5 Server Data	517
Restore Procedure	517
Restore Options	520
Monitoring a Lotus Domino Server Backup and Restore	523
Monitoring Current Sessions	523
Viewing Previous Sessions	524
Troubleshooting	526
General Troubleshooting	526
Checking Prerequisites Related to the Lotus Domino R5 Server Side of the Integration	527
Configuration Problems	530
Backup Problems	530
Restore Problems	534
Recovery Problems	535
Before You Call Support	537

Glossary

Index

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
B6960-90061	August 2002	Data Protector Release A.05.00
B6960-90082	May 2003	Data Protector Release A.05.10

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>

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 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. There are two versions of this manual:

- ***HP OpenView Storage Data Protector Windows Integration Guide***

This manual describes integrations running the Windows operating systems, such as Microsoft Exchange, Microsoft SQL, Oracle, SAP R/3, Informix, Sybase, NetApp Filer, HP OpenView Network Node Manager, and Lotus Domino R5 Server.

- *HP OpenView Storage Data Protector UNIX Integration Guide*

This manual describes integrations running on the UNIX operating system, such as Oracle, SAP R/3, Informix, Sybase, NetApp Filer, IBM DB2 UDB, HP OpenView Network Node Manager, and Lotus Domino R5 Server.

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 EMC Symmetrix Integration Guide

This manual describes how to install, configure, and use the EMC Symmetrix integration. It is intended for backup administrators or operators.

It describes the integration of Data Protector with the EMC Symmetrix Remote Data Facility and TimeFinder features for Symmetrix Integrated Cached Disk Arrays. It covers the backup and restore of file systems and disk images, as well as online databases, such as Oracle and SAP R/3.

HP OpenView Storage Data Protector HP StorageWorks Disk Array XP Integration Guide

This manual describes how to install, configure, and use the integration of Data Protector with HP StorageWorks Disk Array XP. It is intended for backup administrators or operators. It covers the backup and restore of Oracle, SAP R/3, Microsoft Exchange, and Microsoft SQL.

HP OpenView Storage Data Protector EVA/VA/MSA Integration Guide

This manual describes how to install, configure, and use the integration of Data Protector with HP StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array or HP StorageWorks Modular SAN Array 1000. It is intended for backup administrators or operators. It covers the backup and restore of Oracle, SAP R/3, Microsoft Exchange, and Microsoft SQL.

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 MPE/iX System User Guide

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations

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).

HP OpenView Storage Data Protector Software Release Notes

This manual gives a description of new features of HP OpenView Storage Data Protector A.05.10. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.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 UNIX Integration Guide* describes how to install, configure, and use integrations of Data Protector with other software products on the UNIX platform.

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 in order to fully understand the fundamentals and the model of Data Protector.

Organization

The manual is organized as follows:

- Chapter 1** “Integrating Oracle8/9 and Data Protector” on page 1.
- Chapter 2** “Integrating SAP R/3 and Data Protector” on page 117.
- Chapter 3** “Integrating IBM DB2 UDB and Data Protector” on page 207.
- Chapter 4** “Integrating Sybase and Data Protector” on page 255.
- Chapter 5** “Integrating Informix and Data Protector” on page 337.
- Chapter 6** “Integrating the NDMP Server and Data Protector” on page 421.
- Chapter 7** “Integrating Network Node Manager and Data Protector” on page 463.
- Chapter 8** “Integrating Lotus Domino R5 Server and Data Protector” on page 483.
- 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 Windows Integration Guide*:

- Microsoft SQL Server 7.0/2000
- Microsoft Exchange
- Microsoft Exchange 2000

The integrations of Data Protector with the following applications are described in the *HP OpenView Storage Data Protector Administrator’s Guide*:

- Microsoft Cluster Server
- MC/ServiceGuard
- Data Source Integration
- Application Response Measurement
- ManageX

In This Chapter

This chapter explains how to install, configure, and use the Data Protector Oracle8/9 integration.

It is organized into the following sections:

“Overview” on page 3

“Prerequisites and Limitations” on page 5

“Integration Concept” on page 8

“Data Protector Oracle8/9 Configuration Files” on page 13

“Installing and Upgrading the Oracle8/9 Integration” on page 18

“Configuring the Integration” on page 20

“Configuring an Oracle8/9 Backup” on page 37

“Backing Up an Oracle8/9 Database” on page 52

“Restoring Oracle8/9 Databases” on page 66

“Monitoring an Oracle8/9 Backup and Restore” on page 83

“Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration” on page 84

“Oracle8i and Oracle9i RMAN Metadata and Data Protector Media Management Database Synchronization” on page 87

“Configuring the Integration as Cluster-Aware” on page 89

“Troubleshooting” on page 92

“Examples of an Oracle8/9 Database Restore” on page 107

Overview

Data Protector offers offline as well as online backup of the Oracle8/9 Server instances. In order to enable recovery from an online backup, the respective Oracle8/9 Server instance must operate in the ARCHIVELOG mode.

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

Backup Types

You can perform the following types of backups using the Data Protector integration:

- Online backup of a whole database or parts of it
- Online incremental backup (Oracle8/9 differential incremental backup 1 to 4)
- Offline backup of a whole database
- Backup of Archived Redo Logs only
- Backup of the Oracle8/9 recovery catalog
- Backup of the Oracle8/9 control file

Restore Types

Using the Data Protector Oracle8/9 integration, you can restore the following:

- The whole database or parts of it
- The database to a specific point in time
- From incremental backup
- To a host other than the one where the database originally resided
- A datafile to a location other than its original one
- A catalog before restoring the database
- From a chain of incremental backups

Why use the Data Protector User Interface?

Using the Data Protector Oracle8/9 integration with the Data Protector GUI or Recovery Manager (RMAN) offers several advantages over using RMAN alone:

- Central Management

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

- Media Management

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

- 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. See 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 built-in Data Protector database and in the Oracle8/9 recovery catalog database.

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

Prerequisites and Limitations

Prerequisites

- A special license is needed to use the Data Protector Oracle8/9 integration. See “HP OpenView Storage Data Protector On-Line Extension”, Appendix-A in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- Before you begin, ensure that you have correctly installed and configured the Oracle8/9 Server and 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 on various architectures.
 - ✓ *HP OpenView Storage Data Protector Administrator’s Guide* for instructions on how to configure and run backups.
 - ✓ *Oracle8i/9i Recovery Manager User’s Guide and References* for Oracle8/9 concepts and backup/recovery strategies.
 - ✓ *Oracle8i Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle8/9 backup terminology and concepts.
 - ✓ *Oracle8 Enterprise Manager User’s Guide* for information about backup and recovery with the Oracle8/9 Enterprise Manager, as well as information about Oracle8 Server Manager and about Oracle9i SQL Plus.
- The database you are about to recover must first have been backed up.
- An instance of Oracle 8/9 must be created and configured on the host to which you want to restore the database.
- If you want to restore tablespaces or datafiles, the Oracle database to which they belong must exist.

Prerequisites and Limitations

- The database must be in `Mount` state if the whole database is being restored, or in `NoMount` state if just the control file is being restored.
- It is assumed that you are familiar with Oracle8/9 database administration and basic Data Protector functionality.
- The Oracle8/9 *database* user used by this integration to connect to the target Oracle8/9 database during the backup must have the `SYSDBA` privilege granted. Refer to Oracle8/9 documentation for more information on user privileges in Oracle8/9.

Limitations

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

- Data Protector does not check whether database objects to be restored were backed up and exist in the Data Protector internal database. The restore procedure simply starts.
- When restoring a database to a host other than the one on which the database originally resided, the instance name on the new host must be the same as the original database instance name.
- When using RMAN scripts in Oracle 8/9 backup specifications, double quotes (“”) must not be used - single quotes should be used instead.
- When editing RMAN scripts in the Data Protector Restore GUI for Oracle, a maximum of two backup commands can be used. The first backup command is used for backing up any of the database objects (or the whole database), while the second is reserved for backing up only the archive logs. SQL commands can be used.
- The user will not be able to edit RMAN scripts before they are executed when performing a restore operation.
- When editing RMAN scripts in Oracle 8/9 backup specifications, usage of “ (“ and “)” is not supported.
- When editing an RMAN script, the term “backup”, which is a reserved word, must be written in capital letters when it is used in the SQL command. In all other RMAN commands, do not use capital letters for the reserved word “backup”. The term “backup” should not be used in the name of a backup specification e.g. “LAST Backup”.

- If the users perform the restore of tablespaces until point-in-time, the recovery of tablespaces until point in time should be done from RMAN. Alternatively users can perform the recovery of the database from the Data Protector Oracle Restore GUI.

Integration Concept

The Data Protector Oracle8/9 integration links the Oracle8/9 database management software with Data Protector. From the Oracle8/9 point of view, Data Protector represents a media management utility. On the other hand, the Oracle8/9 database management system can be seen as a data source for backup, using media controlled by Data Protector.

Components

The software components involved in backup and restore processes are:

- The Oracle8/9 Recovery Manager (RMAN)
- The Data Protector Oracle8/9 Integration software

Integration Functionality Overview

The Data Protector Oracle8/9 Integration agent works with Oracle to manage all aspects of backup, restore, and recovery operations on the Oracle8/9 target database. The following functionality is offered:

- Database startup and shutdown
- Backups (backup and copy)
- Recovery (restore and recovery)
- Catalog maintenance, catalog analysis and display
- Stored script maintenance, miscellaneous operations

How Does the Integration Work?

The Data Protector Oracle8/9 Integration agent for Oracle uses RMAN functionality to direct the Oracle8/9 server processes on the target database to perform backup, restore, and recovery operations. Furthermore, it maintains the required information about the Target Databases in the recovery catalog, the Oracle8/9 central repository of information, or in the control file of a particular Target Database.

The Data Protector Oracle8/9 Integration agent for Oracle uses the information in the recovery catalog to determine how to execute the requested backup and restore operations.

The Data Protector Oracle8/9 Integration agent for Oracle gets the following information about Oracle8/9 backup objects either from the control files in the Oracle8/9 Target Database or from the Oracle8/9 recovery catalog:

- The physical schema of the Oracle8/9 Target Databases

- Datafile and archive log backup sets and pieces
- Datafile copies, Archived Redo Logs
- Runtime information on backup and restore jobs.

You can back up Oracle8/9 control files, datafiles, and Archived Redo Logs using RMAN.

The RMAN and Oracle8/9 server integrate with Data Protector to provide complete backup and restore of Oracle8/9 Target Databases to tape storage.

The interface from the Oracle8/9 server processes to Data Protector is provided by the Data Protector Database Library (Media Management Layer or Media Management Library), which is a set of routines that allows the reading and writing of data to Media Agents.

Besides handling direct interaction with the media devices, Data Protector provides scheduling, media management, network backups, monitoring, and interactive backup.

Oracle Backup Types Handled by the Integration

Oracle8/9 full and incremental backup types (up to incremental level 4) can be performed using this integration. As opposed to a full backup, where all data blocks per data file are backed up, only the data blocks that have changed since the previous backup are included in an incremental backup. The full backup type is not related to the number of data files included in the backup, and can therefore be performed per single datafile. The data being backed up, regardless of the backup type (full or incremental), is selected and controlled by Oracle8/9.

NOTE

Regardless of the Oracle8/9 backup type specified, Data Protector always marks the Oracle8/9 backups as full in the Data Protector database, since the Data Protector incremental backup concept is different from the Oracle8/9 incremental backup concept.

A backup that includes all data files that belong to an Oracle8/9 Server instance is known as a whole database backup.

These features can be used for online or offline backup of the Oracle8/9 Target Database. However, you must ensure that the backup objects (such as tablespaces) are switched into the appropriate state before and after a backup session. For online backup, the database instance must

operate in the ARCHIVELOG mode; whereas for offline backup, objects need to be prepared for backup by using the Pre-exec and Post-exec options in the backup specification.

The Data Protector backup specification contains information about backup options, commands for RMAN, Pre- and Post-exec commands, media, and devices.

The Data Protector backup specification allows you to configure a backup and then use the same specification several times. Furthermore, scheduled backups can only be performed using a backup specification.

Backup and restore of an Oracle8/9 Target Database can be performed using the Data Protector User Interface, the RMAN CLI, or the Oracle8/9 Enterprise Manager utility.

The heart of the Data Protector Oracle8/9 integration is the **Data Protector Database Library**, which enables an Oracle8/9 server process to issue commands to Data Protector for backing up or restoring parts or all of the Oracle8/9 Target Database files. The main purpose is to control direct interaction with media and devices.

Backup Flow

Before a backup session begins, the Oracle8/9 Target Database instance is switched into backup mode.

A Data Protector scheduled or interactive backup is triggered by the Data Protector Backup Session Manager, which reads the backup specification and starts the `ob2rman.exe` script on the Oracle8/9 Server under a specific user. This user must be defined as the owner of the Data Protector Oracle8 backup specification. Further on, the `ob2rman.exe` script prepares the environment to start the backup, and issues the Recovery Manager (RMAN) backup command. RMAN instructs the Oracle8/9 Server processes to perform the specified command.

The Oracle8/9 Server processes initialize the backup through the Database Library, which establishes a connection to the Data Protector Backup Session Manager. The Backup Session Manager starts the Media Agent, sets up a connection between the Database Library and the Media Agent, and then monitors the backup process.

The Oracle8/9 Server processes read the data from the disks and send it to the backup devices through the Database Library and the Media Agent.

RMAN writes information regarding the backup either to the recovery catalog (if one is used) or to the control file of the Oracle8/9 Target Database.

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

The Data Protector Media Agent writes data to the backup devices.

Once the backup has finished, the Oracle8/9 database is switched into normal mode.

Restore Flow

A restore session can be started from the Data Protector Restore GUI for Oracle, by issuing the RMAN restore command from the RMAN command line, or from the Oracle8/9 Enterprise Manager graphical user interface. You must specify which objects are to be restored.

A restore from the Data Protector user interface is triggered by the Data Protector Restore Session Manager, which starts the `ob2rman.exe` script. The `ob2rman.exe` script prepares the environment to start the restore, and issues the RMAN restore command. RMAN checks the recovery catalog (if one is used) or the control file to gather the information about the Oracle8/9 backup objects. It also contacts the Oracle8/9 Server processes, which initialize the restore through the Database Library. The Database Library establishes a connection with the Restore Session Manager and passes along the information about which objects and object versions are needed.

The Restore Session Manager checks the IDB to find the appropriate devices and media, starts the Media Agent, establishes a connection between the Database Library and the Media Agent, and then monitors the restore and writes messages and information regarding the restore to the IDB.

The Media Agent reads the data from the backup devices and sends it to the Oracle8/9 Server processes through the Database Library. The Oracle8/9 Server Processes write the data to the disks.

The concept of Oracle8/9 integration, data and the control flow are shown in Figure 1-1 on page 12, and the related terms are explained in the following table.

Figure 1-1 Data Protector Oracle8/9 Integration Concept

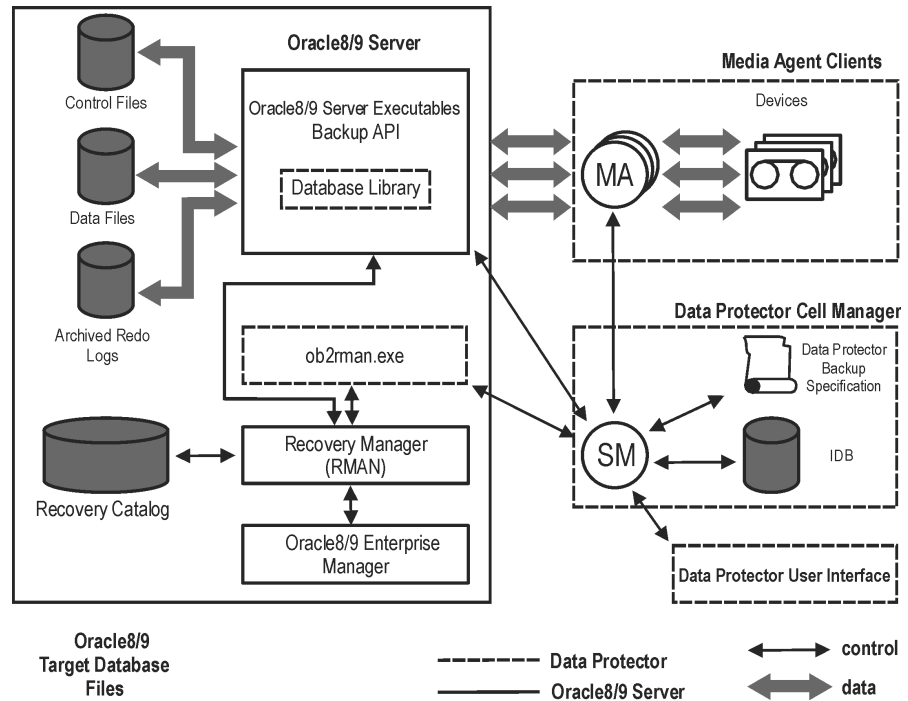


Table 1-1

Legend

SM	The Data Protector Session Manager, which can be the Data Protector Backup Session Manager during a backup session and the Data Protector Restore Session Manager during a restore session.
RMAN	The Oracle8/9 Recovery Manager.
Database Library	The Data Protector set of routines that enables data transfer between the Oracle8/9 Server and Data Protector.
Backup API	The Oracle-defined application programming interface.
IDB	The IDB where all the information about Data Protector sessions, including session messages, objects, data, used devices, and media is written.
MA	The Data Protector Media Agent, which reads and writes data from and to media devices.

Data Protector Oracle8/9 Configuration Files

Data Protector stores Oracle8/9 integration parameters in two files on the Cell Manager:

- For every configured Oracle8/9 *instance* in the:
/etc/opt/omni/integ/config/Oracle8/<client_name>%<ORACLE_SID> file (HP-UX and Solaris systems), or in the
<Data_Protector_home>\Config\integ\config\oracle8\<client_name>%<ORACLE_SID> file (Windows systems) on the Cell Manager.

The parameters stored in the **instance configuration file** are:

- ✓ Oracle home directory,
 - ✓ Oracle version,
 - ✓ encoded connection strings to the target database and recovery catalog and
 - ✓ the variables which need to be exported prior to starting a backup, and which affect the Oracle8/9 instance.
- Oracle8/9 *global* integration parameters in the:
/etc/opt/omni/integ/config/Oracle8/<client_name>%_OB2_GLOBAL file (HP-UX and Solaris systems), or in the
<Data_Protector_home>\Config\integ\config\oracle8\<client_name>%_OB2_GLOBAL file (Windows systems) on the Cell Manager.

The parameters stored in the **global configuration file** are:

- ✓ instance list (all Oracle8/9 instances on the Oracle8/9 server) and
- ✓ the variables that need to be exported prior to starting a backup, and which affect every Oracle8/9 instance on the Oracle8/9 server.

The configuration parameters are written to the Data Protector Oracle8/9 configuration files:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

**Global
Configuration File
Syntax**

The syntax of the file is as follows:

IMPORTANT

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

```
INSTANCE_LIST=( '<ORACLE_SID1>' [ , '<ORACLE_SID2>' , '<ORACLE_SID3>' .  
.. ] );  
Environment={  
    [<ENV var1>=<value1>'];  
    [<ENV var2>=<value2>'];  
    ...]  
}
```

**Example of Global
Configuration File**

This is an example of the Data Protector Oracle8/9 global configuration file:

```
INSTANCE_LIST=( 'ICE' , 'BUREK' );  
Environment={  
    OB2OPTS='-debug 1-200 GLOB.txt';  
    OB2FASTRAXDUMPDIR='/tmp';  
}
```

IMPORTANT

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

**Instance
Configuration File
Syntax**

The syntax of the file is as follows:

```
TGTLogin=<encoded connection string to target database>;  
RCVLogin=<encoded connection string to recover catalog  
database>;  
ORACLE_HOME=<Oracle's instance home directory>;  
[ORACLE_VERSION=<Version of Oracle software>;]  
Environment={
```

```
[<ENV var1>=<value>'];]  
[<ENV var2>=<value>'];  
...]  
}
```

Example of Instance Configuration File

This is an example of the Data Protector Oracle8/9 instance configuration file:

```
TGTLogin='EIBBKIB...BDGBBGFBBMFBB';  
RCVLogin='DIBBOH...BEFBBCFBBFGBB';  
ORACLE_HOME='/opt/oracle/product/8.1.6';  
ORACLE_VERSION='8.1.6 64bit';  
Environment={  
NLS_LANG='AMERICAN_AMERICA.US7ASCII';  
}
```

Setting, Retrieving and Listing Data Protector Oracle8/9 Configuration Files' Parameters Using the CLI

The Data Protector Oracle8/9 configuration files' parameters are normally written to the Data Protector Oracle8/9 configuration files after the completed configuration of the Oracle8/9 instance in Data Protector.

The variable definitions that are command or shell-based must be entered in the following file, since such variable definitions are not possible in the Data Protector Oracle8/9 configuration files:

/etc/opt/omni/oracle8/<ORACLE_SID>/ .profile (on HP-UX and Solaris systems) or

/usr/omni/config/oracle8/<ORACLE_SID>/ .profile (on other UNIX systems)

The util_cmd Command

You can set, retrieve or list the Data Protector Oracle8/9 configuration files' 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 Oracle8/9 client. The command resides in the `/opt/omni/lbin` (HP-UX and Solaris systems) directory or in the `/usr/omni/bin` (other UNIX systems) directory.

The util_cmd Synopsis

The syntax of the util_cmd command is as follows:

```
util_cmd -getconf[ig] Oracle8 <Oracle8_instance> [-local \  
<filename>]
```

```
util_cmd -getopt[ion] [Oracle8 <Oracle8_instance>] \  
<option_name> [-sub[list] <sublist_name>] [-local \  
<filename>]
```

```
util_cmd -putopt[ion] [Oracle8 <Oracle8_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, filename that the command output is written to the filename for the output of the command to be written to (if the -local option is not specified, the output is written to the standard output).
- When with the -getopt[ion], the filename of the file from which the parameter and its value are to be taken from and then written to the standard output (if the -local option is not specified, the parameter and its value are taken from one of the Data Protector Oracle8/9 configuration files and then written to the standard output).
- When with the -putopt[ion] option, the filename for the output of the command to be written to (if the -local option is not specified, the output is written to one of the Data Protector Oracle8/9 configuration files).

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 appears 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, one of the Data Protector Oracle8/9 configuration files is missing on the Cell Manager, etc.

Setting Parameters

To set the `OB2OPTS` and `NLS_LANG` parameters for the Oracle8/9 instance `ICE`, use the following commands:

```
/opt/omni/sbin/util_cmd -putopt Oracle8 ICE OB2OPTS \  
'-debug 1-200 INSTANCE.txt' -sublist Environment
```

```
/opt/omni/sbin/util_cmd -putopt Oracle8 ICE NLS_LANG \  
AMERICAN_AMERICA.US7ASCII -sublist Environment (HP-UX and  
Solaris systems)
```

```
/usr/omni/bin/util_cmd -putopt Oracle8 ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment
```

```
/usr/omni/bin/util_cmd -putopt Oracle8 ICE NLS_LANG \  
AMERICAN_AMERICA.US7ASCII -sublist Environment (other UNIX  
systems)
```

Retrieving Parameters

To retrieve the value of the `OB2OPTS` parameter for the instance `ICE`, use the following command:

```
/opt/omni/sbin/util_cmd -getopt Oracle8 ICE OB2OPTS \  
-sublist Environment (HP-UX and Solaris systems)
```

```
/usr/omni/bin/util_cmd -getopt Oracle8 ICE OB2OPTS \  
-sublist Environment (other UNIX systems)
```

Listing Parameters

To list the Data Protector configuration files' parameters for the instance `ICE`, use the following command:

```
/opt/omni/sbin/util_cmd -getconf Oracle8 ICE (HP-UX and  
Solaris systems)
```

```
/usr/omni/bin/util_cmd -getconf Oracle8 ICE (other UNIX  
systems)
```

Installing and Upgrading the Oracle8/9 Integration

- Prerequisite** You need to shut down any running Oracle8 databases on the Oracle8 servers to be integrated before upgrading or installing the Data Protector Oracle8 integration, and restart them after the upgrade or installation is complete.
- Upgrading** The Oracle 8/9 integration is upgraded automatically during the client upgrade. After the upgrade procedure has completed, some additional steps have to be performed manually to finish the upgrade. For the upgrade procedure, refer to “Upgrading to Data Protector A.05.10” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- Installation** Install the Data Protector Oracle8/9 integration software on your Oracle8/9 Server system either locally, from the CD-ROM, or remotely, using the Data Protector GUI.
- You must install the following Data Protector software components on the Oracle8/9 server system:
- Oracle8 Integration
 - Disk Agent
 - Media Agent (if you have devices connected to the system)
- It is recommended that you also install:
- User Interface
- Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details about the installation.
- Verifying the Installation** Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What's Next?

By now you should have installed the Data Protector Oracle8/9 integration software on the Oracle8/9 Server system so that the Oracle8/9 Server system has become a Data Protector client. At this stage you are ready to start the configuration procedure described in the next section.

Configuring the Integration

It is assumed that the installation of Data Protector software components on the Oracle8/9 Server system was successful.

Prerequisites

The following requirements must also be met before you start configuring the Oracle8/9 integration:

- The Oracle8/9 Server software must be installed and the Oracle8/9 Target Database must be online.
- The connection string for the Oracle8/9 target database and the recovery catalog must be properly configured in order to start the RMAN.

Refer to the *Oracle8i Recovery Manager User's Guide and References* for more information about different connection options.

- If the recovery catalog is used, make sure that it is properly installed and online. Refer to the *Oracle8i Recovery Manager User's Guide and References* for further information about how to create a recovery catalog.
- The TNS listener processes must be configured and running for the Oracle8/9 Target Database and the recovery catalog, if used.

Refer to the Oracle8/9 documentation for further information.

Refer to “Troubleshooting” on page 92 for details about how to check all the prerequisites listed above.

- Make sure to set any Oracle8/9-related environmental variables needed for the Oracle8/9 database to function properly (for example, the Oracle8/9 NLS_LANG environmental variable). Please refer to Oracle8/9 documentation for more information.

Configuration Overview

The following list gives an overview of the global tasks for configuring the Oracle8/9 integration.

1. “Linking Oracle8/9 with the Data Protector Database Library” on page 21.
2. “Configuring an Oracle8/9 User in Data Protector” on page 27.
3. “Configuring the Oracle8/9 Server” on page 29.

4. “Configuring an Oracle8/9 Backup” on page 37.

Linking Oracle8/9 with the Data Protector Database Library

In order to use the Data Protector Oracle8/9 integration, you need to manually link the Oracle8/9 server software and the Data Protector Database Library on the Data Protector Oracle8/9 Server system.

The Data Protector Database Library is invoked by the Oracle8/9 server when it needs to write to or read from devices using Data Protector.

IMPORTANT

After uninstalling the Data Protector Oracle8/9 integration on an Oracle8/9 server system, the Oracle8/9 server software is still linked to the Data Protector Database Library. You must rebuild (Oracle8) or relink (Oracle8i or Oracle9i) the Oracle8/9 binary to remove this link. If this is not done, the Oracle8/9 server cannot be started after the integration has been removed. Please refer to “Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration” on page 84 for more information on removing the integration link.

Linking Oracle8/9 with the Data Protector Database Library on HP-UX systems

On Oracle8/9 Server systems running on HP-UX, the Data Protector Database Library `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) is located in the `/opt/omni/lib` directory.

Proceed as follows:

1. On the Oracle8/9 Server system, connect to the Oracle8/9 database as an Oracle8/9 operating system user and shut down all Oracle8/9 instances.

Oracle8.x

2. If you have Oracle8.x installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```

Integrating Oracle8/9 and Data Protector

Configuring the Integration

- b. If you have a 32-bit integration, execute the following command:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8.sl" ioracle,
```

while for a 64-bit integration, the respective command is:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.sl" ioracle
```

Oracle8i and Oracle9i

3. If you have Oracle8i or Oracle9i installed, perform the following:

- a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),
```

```
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```

- b. `mv libobk.sl libobk.sl.orig`

IMPORTANT

If you intend to uninstall the Data Protector Oracle8/9 integration and to continue using Oracle8/9 on the same system after the integration is removed, do not delete the `libobk.sl.orig` file.

- c. `ln -s /opt/omni/lib/libob2oracle8.sl libobk.sl (32-bit Oracle) or`

```
ln -s /opt/omni/lib/libob2oracle8_64bit.sl libobk.sl  
(64-bit Oracle).
```

4. Start all Oracle8/9 instances.

Oracle8.x

5. If you have Oracle 8.x installed, check if the `libob2oracle8.sl (libob2oracle8_64bit.sl)` file is linked with the Oracle8 executable using the following command:

```
/usr/bin/chrtr <ORACLE_HOME>/bin/oracle (32-bit Oracle8)
```

```
/usr/ccs/bin/ldd <ORACLE_HOME>/bin/oracle (64-bit Oracle8)
```

The `/opt/omni/lib/libob2oracle8.sl (libob2oracle8_64bit.sl)` library must be listed in the shared library list.

The `chatr` command also verifies whether the shared library dynamic path `SHLIB_PATH` is enabled.

Example

The following is an example extract of the command output:

```
bin/oracle:
  shared executable
  shared library dynamic path search:
    SHLIB_PATH  enabled second
    embedded path disabled first Not Defined
  shared library list:
    static /opt/omni/lib/libob2oracle8.sl
    dynamic /usr/lib/librt.2
    dynamic /usr/lib/libnss_dns.1
    dynamic /usr/lib/libdld.2
  shared library binding:
    deferred
  static branch prediction disabled
  kernel assisted branch prediction enabled
  lazy swap allocation disabled
  text segment locking disabled
  data segment locking disabled
  data page size: D (default)
  instruction page size: D (default)
```

The line starting with the `SHLIB_PATH` entry should be as given in the example above. If this line is different, then enable the shared library dynamic path as follows:

- a. Stop all Oracle8 instances.
- b. `$ chatr +s enable $ORACLE_HOME/bin/oracle`
- c. Start all Oracle8 instances.

Linking Oracle8/9 with the Data Protector Database Library on Solaris Systems

On Oracle8/9 Server systems running on Solaris systems, the Data Protector Database Library `libob2oracle8.so` (`libob2oracle8_64bit.so`) is located in the `/opt/omni/lib` directory.

Proceed as follows:

1. On the Oracle8/9 Server system, connect to the database as an Oracle8/9 operating system user and shut down all Oracle8/9 instances.

Oracle8.x

2. If you have Oracle8.x installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```

- b. Execute the following command:

```
make -e -f ins_rdbms.mk LLIBMM=  
LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or  
make -e -f ins_rdbms.mk LLIBMM=  
LLIBOBK="-L/opt/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Oracle8i and Oracle9i

3. If you have Oracle8i or Oracle9i installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),
```

```
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```

- b. Execute the following commands:

```
mv libobk.so libobk.so.orig
```

IMPORTANT

If you intend to uninstall the Data Protector Oracle8/9 integration and to continue using Oracle8/9 on the same system after the integration is removed, do not delete the `libobk.so.orig` file.

```
ln -s /opt/omni/lib/libob2oracle8.so libobk.so (32-bit Oracle) or
```

```
ln -s /opt/omni/lib/libob2oracle8_64.so libobk.so (64-bit Oracle).
```

4. Start all Oracle8/9 instances.

Oracle8.x

5. If you have Oracle 8.x installed, check if the `libob2oracle8.so` file is linked with the Oracle8 executable using the following command:

```
/opt/bin/ldd <ORACLE_HOME>/bin/oracle
```

The `/opt/omni/lib/libob2oracle8.so` library has to be listed as required by the Oracle8 executable, and the `LD_LIBRARY_PATH` must be enabled. If the `/opt/omni/lib/libob2oracle8.so` library is not listed, then you need to re-build the Oracle binary as follows:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib /opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or
```

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/opt/omni/lib /opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit Oracle).
```

Linking Oracle8/9 with the Data Protector Database Library on Other UNIX Systems

On Oracle8/9 Server systems running on other UNIX systems, the Data Protector Database Library `libob2oracle8.so` (`libob2oracle8_64bit.so`) is located in the `/usr/omni/lib` directory.

NOTE

On AIX systems, the extensions for shared objects is `.a`. The Data Protector Database Library for the 32-bit version of Oracle8/9 is called `libob2oracle8.a`, and that for the 64-bit version `libob2oracle8_64bit.a`.

Integrating Oracle8/9 and Data Protector

Configuring the Integration

Proceed as follows:

1. On the Oracle8/9 Server system, connect to the database as an Oracle8/9 operating system user and shut down all Oracle8/9 instances.

Oracle8.x

2. If you have Oracle8.x installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```

- b. Execute the following command:

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/opt/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or
```

```
make -f ins_rdbms.mk LLIBMM= LLIBOBK="-L/usr/omni/lib  
/opt/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Oracle8i and Oracle9i

3. If you have Oracle8i or Oracle9i installed, proceed as follows:

- a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32bit Oracle),
```

```
cd <ORACLE_HOME>/lib64 (64bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64bit Oracle9i).
```

- b. Execute the following commands:

```
mv libobk.so libobk.so.orig
```

IMPORTANT

If you intend to uninstall the Data Protector Oracle8/9 integration and to continue using Oracle8/9 on the same system after the integration is removed, do not delete the `libobk.so.orig` file.

```
ln -s /usr/omni/lib/libob2oracle8.so libobk.so (32-bit  
Oracle) or
```

```
ln -s /usr/omni/lib /libob2oracle8_64.so libobk.so  
(64-bit Oracle).
```

4. Start all Oracle8/9 instances.

Oracle8.x

5. If you have Oracle 8.x installed, check if the `libob2oracle8.so` file is linked with the Oracle8 executable using the following command:

```
/usr/bin/dump -H <ORACLE_HOME>/bin/oracle (on AIX systems) or  
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle (on other UNIX  
systems).
```

The `/usr/omni/lib/libob2oracle8.so` library has to be listed as required by the Oracle8 executable and the `LD_LIBRARY_PATH` must be enabled. If the `/usr/omni/lib/libob2oracle8.so` library is not listed, then you need to re-build the Oracle binary as follows:

```
make -e -f ins_rdbms.mk LLIBMM= LLIBOBK="" -L/usr/omni/lib  
/usr/omni/lib/libob2oracle8.so" ioracle (32-bit Oracle) or  
make -e -f ins_rdbms.mk LLIBMM= LLIBOBK="" -L/usr/omni/lib  
/usr/omni/lib/libob2oracle8_64bit.so" ioracle (64-bit  
Oracle).
```

Configuring an Oracle8/9 User in Data Protector

In order to start an Oracle8/9 backup session, a user needs to perform an operating system logon to the system where an Oracle8/9 Server is running.

In addition, this operating system user must be registered in the Oracle8/9 database and identified by Oracle8/9 through the operating system identification.

This means that the Oracle8/9 Server does not request connection information from an application started under such a user account, but only checks whether the operating system user is registered in the database.

Refer to the Oracle8/9 documentation for further information about the different types of connections, the roles and privileges of Oracle8/9 database administrators, and security issues that should be considered.

If properly configured, this user is allowed to back up or restore an Oracle8/9 database. In order to start a backup of an Oracle8/9 database using Data Protector, the user must also become the owner of the Data Protector backup specification.

As the owner of the backup specification, the user must be added to either the Data Protector admin or operator user group.

You can identify this user by running the following command on the Oracle8/9 Server system:

```
ps -ef |grep ora_pmon_<ORACLE_SID>  
or  
ps -ef |grep ora_lgwr_<ORACLE_SID>
```

Figure 1-2

Finding the Oracle8/9 User



```
# ps -ef | grep ora_pmon  
ora 2675 1 4 Sep 24 ? 0:13 ora_pmon  
#
```

The example above states that the user ora has sufficient privileges within the Oracle8/9 database to back up and restore the database. Therefore, this user must be added to the corresponding Data Protector user group (admin or operator) and must also become the owner of the backup specification in order to be able to back up the Oracle8/9 database using Data Protector.

IMPORTANT

Additionally, the operating system root user on the Oracle8/9 Server also has to be added to either the Data Protector admin or operator user group.

After the two users are added to either the Data Protector admin or operator user group, Data Protector sessions can be started under the user account with all the necessary privileges required to perform an Oracle8/9 database backup with Data Protector.

If two or more Oracle users have the same user ID, all of them must be added to either the Data Protector admin or operator user group.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for detailed information on Data Protector user rights and how to add a user to a user group.

Configuring the Oracle8/9 Server

Before You Begin It is recommended that you configure and run a Data Protector filesystem backup of the Oracle8/9 Server system. A filesystem backup can be performed only if you have installed the Disk Agent on the Oracle8/9 Server system.

In case of problems, this type of backup is much easier to troubleshoot than an integration of the Oracle8/9 Server system with Data Protector.

Any device can be used for this test. Configure a standard filesystem backup, which can include one directory only. The test should include a partial restore to the Oracle8/9 Server system as well.

Configuring an Oracle8/9 Server system involves preparing the environment for starting a backup. The environment parameters such as the Oracle8/9 home directory and the connection string to the database are saved in the Data Protector Oracle8/9 configuration files on the Cell Manager. The database must be online during the configuration procedure. The configuration must be done for each Oracle8/9 Server instance.

If a recovery catalog has been created and the Oracle8/9 Target Database has not yet been registered in the recovery catalog database, this will occur during the configuration procedure.

The configuration is performed using the Data Protector User Interface.

Configuring the Oracle8/9 Server Using the Data Protector GUI

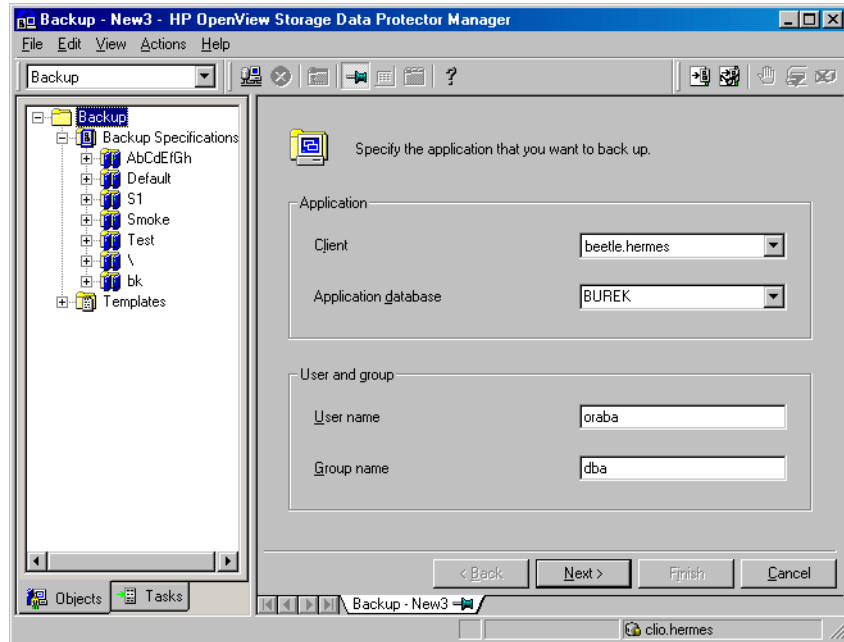
The procedure below describes the configuration of the Oracle8/9 Server system using the Data Protector GUI.

The configuration is performed at the same time that a new backup specification is created:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click Oracle8 Server.
3. Click Add Backup. In the Create New Backup dialog box, double-click any of the predefined backup templates.
4. In the Results Pane, enter the name of the Oracle8/9 Server system you want to configure and the name of the Oracle8/9 Server instance. Then type the UNIX user name and user group of the Oracle8/9 user

as seen in Figure 1-3. Refer to “Configuring an Oracle8/9 User in Data Protector” on page 27 for detailed information on how to identify that user.


Figure 1-3 Specifying an Oracle8/9 Server System



Click Next. If the Oracle8/9 instance has not yet been configured, the Configure Oracle8 dialog window is displayed.

In the General property page, enter the Oracle8/9 Server home directory.

Figure 1-4 **Configuring Oracle8/9 - General**



In the Connection property page, enter the login information for the Oracle8/9 Target Database. Note that the username entered must have the SYSDBA privilege granted.

To grant the sysdba privilege, start the `svrmgrl>` prompt and enter the following:

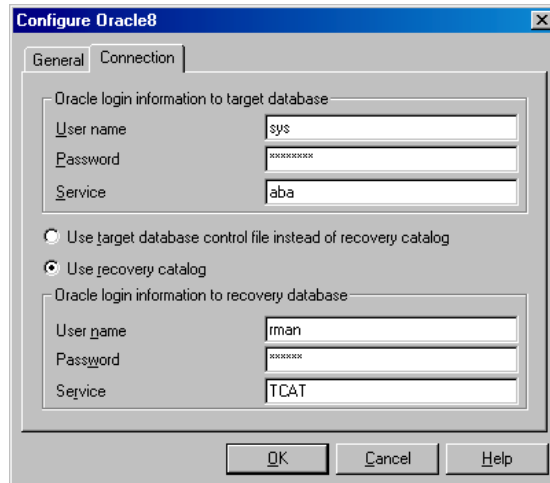
```
svrmgrl> connect internal;  
svrmgrl> alter user system identified by manager;  
svrmgrl> grant sysdba to system;
```

Refer to Oracle8/9 documentation for more information on user privileges in Oracle8/9.

To use the Oracle8/9 Target Database control files, check the Use target database controlfile option.

To use a recovery catalog, check the Use recovery catalog option and enter the login information for the recovery catalog as well.

Figure 1-5 **Configure Oracle8/9 - Connection**



What Happens? The following takes place after clicking OK in the Configure Oracle8 window.

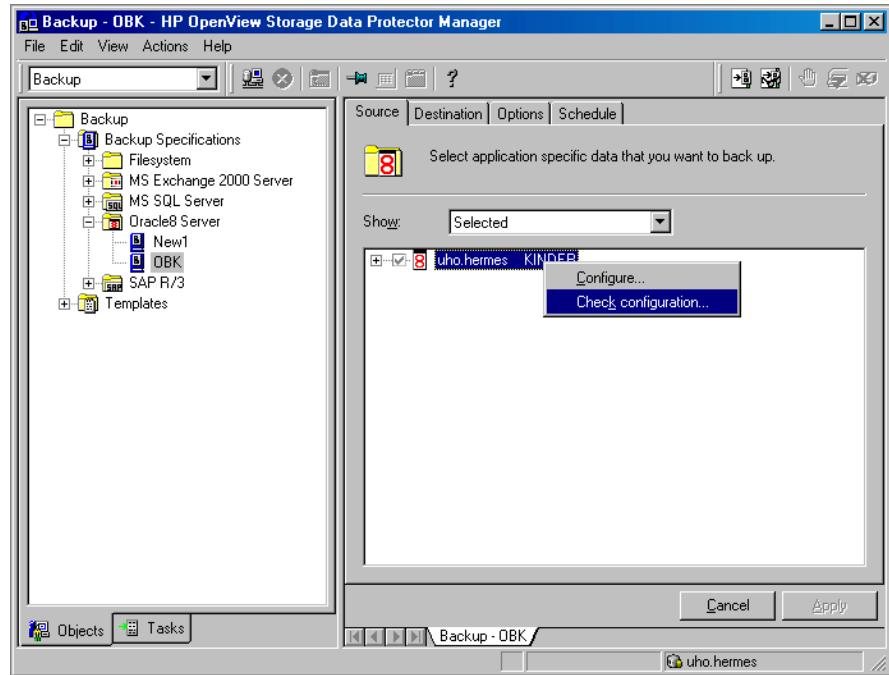
1. The `util_oracle8.exe` executable is started on the Oracle8/9 Server system, and saves the configuration parameters in the Data Protector Oracle8/9 configuration files. Refer to “Data Protector Oracle8/9 Configuration Files” on page 13.
2. If the Use recovery catalog button was selected, the `util_oracle8.exe` executable starts the Oracle8/9 RMAN command, which registers the Oracle8/9 Target Database in the recovery catalog.

Information about the Oracle8/9 Target Database’s structure is transferred to the recovery catalog from the Oracle8/9 Target Database’s control files.

Checking the Configuration After clicking OK in the Configure Oracle8 window, the next page in the wizard allows you to check the configuration.

Data Protector will verify the configuration by trying to connect to the Oracle8/9 Server system using the information specified and saved during the configuration procedure.

Figure 1-6 **Checking the Configuration**



Proceed as follows to check the configuration:

1. Right-click the Oracle8/9 Server system.
2. Click Check Configuration.

If the configuration is successful, you will receive a message confirming that the integration was properly configured. If it was not, you will receive a message explaining the reasons for the unsuccessful configuration.

IMPORTANT

When checking the Data Protector Oracle8/9 integration configuration using the Data Protector GUI, it is possible that although the GUI check returns a successful result, you may still receive the 12:8300 error when trying to start a backup session. Such a backup session will not start. For more information, please see “Troubleshooting” on page 92.

The configuration can also be checked if you have already created and saved a backup specification for backing up a particular Oracle8/9 Server:

1. In the Data Protector Manager, switch to the Backup context. In the Scoping Pane, expand Backup, Backup Specification, then Oracle8.
2. In the Results Area, right-click the backup specification.
3. In the Source property page, right-click the name of the client system, then click Check Configuration.

Configuring the Oracle8/9 Server Using the CLI

Execute the following command to configure an Oracle8/9 Server system using the Data Protector CLI:

On HP-UX and Solaris systems:

```
/opt/omni/sbin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> <Target_Database_Login>[<Recovery_Catalog \  
_Login>]
```

On other UNIX systems:

```
/usr/omni/bin/util_oracle8.exe -CONFIG <ORACLE_SID> \  
<ORACLE_HOME> <Target_Database_Login> \  
[<Recovery_Catalog_Login>]
```

The last connection string is required only if you are using the recovery catalog, otherwise the database control files are used instead.

The variables are defined as follows:

<ORACLE_SID>.	The name of the Oracle8/9 Server instance.
<ORACLE_HOME>.	The home directory of the Oracle8/9 Server instance.
<Target_Database_Login>.	The format of the login information is <user_name>/<password>@<service> . It is used to connect to the Target Database via the Recovery Manager.
<Recovery_Catalog_Login>.	The format of the login information is <user_name>/<password>@<service> . It is used to connect to the recovery

catalog database via the Recovery Manager. This parameter is optional and is to be used only if you are using the recovery catalog database.

In the example below, the Oracle8/9 service name and the ORACLE_SID of the Target database is aba, and the user name sys is identified by the password manager.

The recovery catalog owner is rman, the corresponding Oracle8/9 service name and the ORACLE_SID is TCAT, and the user name rman is identified by the password rman.

On HP-UX and Solaris systems:

```
/opt/omni/sbin/util_oracle8.exe -CONFIG aba \  
/applications/ORA8/app/oracle/product/8.0.4 \  
sys/manager@aba rman/rman@TCAT
```

On other UNIX systems:

```
/usr/omni/bin/util_oracle8.exe -CONFIG aba \  
/applications/ORA8/app/oracle/product/8.0.4 \  
sys/manager@aba rman/rman@TCAT
```

What Happens?

The following happens after typing the command at the command line:

1. The `util_oracle8.exe` executable is started on the Oracle8/9 server system, which saves the configuration parameters in the Data Protector Oracle8/9 configuration files. Refer to “Data Protector Oracle8/9 Configuration Files” on page 13.
2. If the `<Recovery_Catalog_Login>` option has been specified in the command line, the `util_oracle8.exe` executable starts the Oracle8/9 RMAN command line interface, which registers the Oracle8/9 Target Database to the recovery catalog.

Information about the Oracle8/9 Target Database’s structure is transferred to the recovery catalog from the Oracle8/9 Target Database’s control files.

If you need to export some variables before starting the Oracle8 Server Manager or Oracle9i SQL Plus, TNS listener, or Recovery Manager, these variables must be defined in the Environment section of the Data Protector Oracle8/9 global configuration file. Refer to “Data Protector Oracle8/9 Configuration Files” on page 13 for information on how to define such variables.

Checking the Configuration

To check the configuration, log in to the Oracle8/9 server system as an Oracle8/9 group dba user to the Oracle8/9 server system and run the following command:

```
/opt/omni/lbin/util_oracle8.exe -CHKCONF <ORACLE_SID> (HP-UX and Solaris systems) or
```

```
/usr/omni/bin/util_oracle8.exe -CHKCONF <ORACLE_SID> (other UNIX systems).
```

Data Protector attempts to connect to the Oracle8/9 server system using the information that was specified and saved during the configuration procedure.

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

To obtain an error description, start the following command on the Oracle8/9 Server system:

On HP-UX and Solaris systems:

```
/opt/omni/lbin/omnigetmsg 12 <Error_number>
```

On other UNIX systems:

```
/usr/omni/bin/omnigetmsg 12 <Error_number>
```

Configuring an Oracle8/9 Backup

To configure an Oracle8/9 backup, perform the following steps:

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

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

2. Create a backup specification for what and how to back up.

The Data Protector backup specification is stored on the Cell Manager system and contains instructions on how to perform a backup with Data Protector.

Once the backup specification is created and saved, it can be scheduled so that unattended backups may be performed. You may use one of the predefined backup templates for Oracle8/9 Target Database objects, or you can create a new, customized template.

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 than having to specify all the options over and over again. This task is optional, as you can use one of the default templates as well.

If you prefer using predefined templates, refer to “Creating a Backup Specification” on page 38 for a detailed explanation.

To create a new backup template, proceed as follows:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup and then Templates, and then right-click Oracle8 Server.
3. Click Add Template. Follow the wizard to define the appropriate backup options in your template.

Creating a Backup Specification

To create a new backup specification for the Oracle8/9 Target Database, proceed as follows:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications.
3. Right-click Oracle8 Server and then click Add Backup. The Create New Backup dialog box is displayed.
4. Double-click Blank Oracle8 Backup to create a backup specification without pre-defined options, or use one of the pre-defined templates given below:

NOTE

The first four templates in the list below are present because of the backward compatibility.

Archive.	Backs up the Archived Redo Logs.
Archive_Delete.	Backs up the Archived Redo Logs, then deletes them after the backup.
Whole_Online.	Backs up the database instance and the Archived Redo Logs.
Whole_Online_Delete.	Backs up the database instance and the Archived Redo Logs, and then deletes the Archived Redo Logs.
Database_Archive.	Backs up the database instance and the Archived Redo Logs.
Database_Switch_Archive.	Backs up the database instance, switches the Online Redo Logs and backs up the Archived Redo Logs.

- | | |
|------------------------------------|---|
| Database_Switch_ArchiveDel. | Backs up the database instance, switches the Online Redo Logs, backs up the Archived Redo Logs and then deletes the Archived Redo Logs. |
| Direct_Database. | Backs up the database instance and controlfile. |
| SMB_Database. | Backs up the database instance and controlfile in the ZDB (split mirror or snapshot) backup mode. To be used with Oracle8i and Oracle9i only. |
| SMB_Database_Oracle8. | Backs up the database instance and controlfile in the ZDB (split mirror or snapshot) backup mode. To be used with Oracle8 only. |
5. After starting the creation of a backup specification, enter the following information in the Results Area.
- Name of the Oracle8/9 Server system that you want to backup.
 - Name of the Oracle8/9 application instance, and the UNIX user name and group of the Oracle8/9 user as described in the section “Configuring an Oracle8/9 User in Data Protector” on page 27.
- Once you have entered the required information, the Backup Wizard is started, provided that the respective Oracle8/9 Server has already been configured. If not, you must configure the client at this stage by entering the appropriate connection strings.
- See “Configuring the Oracle8/9 Server” on page 29 for details.
6. In the next step of the wizard, select the database objects you want to back up.

TIP

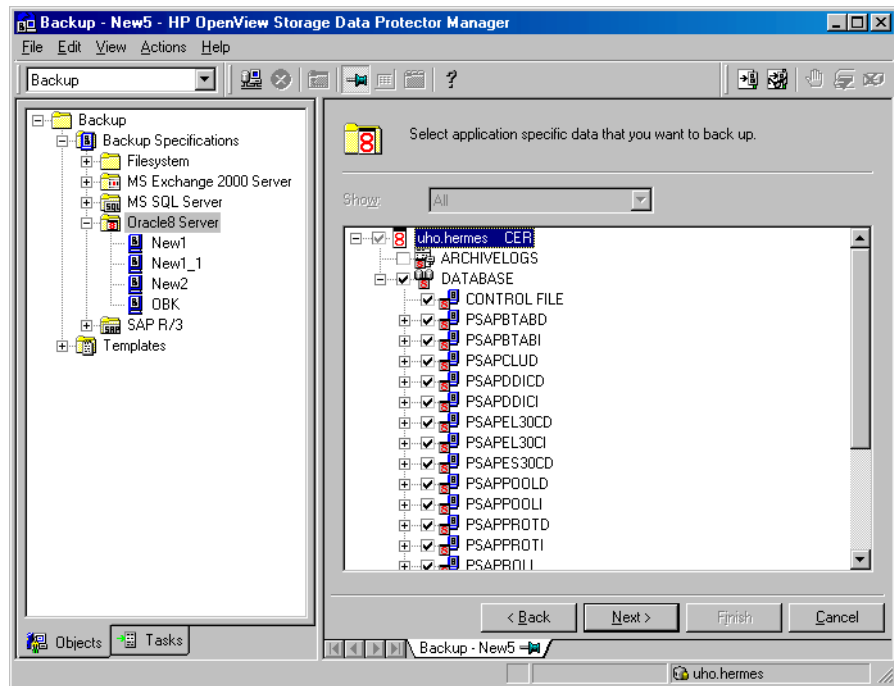
If no database objects are displayed in the Results Area, select All in the Show drop-down list.

For example, a single tablespace can be separately selected for the backup, but for a complete and consistent online backup of the Oracle8/9 Target Database, the ARCHIVELOGS must be selected.

NOTE

If your Oracle8/9 Target Database uses a recovery catalog, this is backed up by default after each Oracle8/9 Target Database backup. This is not the case if control files are used instead of the recovery catalog. The control files must be backed up separately.

Figure 1-7 Selecting Backup Objects:



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

Refer to Data Protector Online Help and the *HP OpenView Storage Data Protector Administrator's Guide* for backup options common to all objects.


See "Oracle8/9 Backup Options" on page 42 for details about Oracle8/9 specific options.

IMPORTANT

The word `DEFAULT` is a reserved word and therefore must not be used for backup specification names or labels of any kind. Oracle8/9 does not allow full stops in channel names. Therefore, do not use punctuation in names of backup specifications, since the Oracle8/9 channel format is created from the backup specification name.

Once you have defined all backup options, you must save and name your Oracle8/9 backup specification under a name of your choice. It is recommended that you save all Oracle8/9 backup specifications in the Oracle8 group.

Figure 1-8 Saving the Backup Specification



You should by now have completed the creation of a backup specification. After the backup specification is saved, verify that the owner of the backup specification is the specified Oracle8/9 user.

See “Configuring an Oracle8/9 User in Data Protector” on page 27 for details about this user.

8. You can examine the newly-created and saved backup specification in the Backup context, under the specified group of backup specifications. The backup specification itself is stored in the `/etc/opt/omni/barlists/oracle8/<Backup_Spec_Name>` file on the Cell Manager system.

It is recommended that you test the backup specification by clicking the `Start Preview` button. The test backup is performed to a null device (`/dev/null`).

You can start a real interactive backup that includes data transfer by clicking the `Start Backup` button.

See “Testing the Integration” on page 49 for more details.

Oracle8/9 Backup Options

Setting the Oracle8/9 options allows you to:

- Edit the Oracle8/9 RMAN script section of the Data Protector Oracle8/9 backup specification. The RMAN script section is created by Data Protector during the creation of a backup specification and reflects the backup specification’s selections and settings.

The RMAN script is used when the Data Protector backup specification is started to perform a backup of the Oracle8/9 objects. For more information on how to use the RMAN script in Data Protector, refer to “Editing the Oracle8/9 RMAN Script” on page 44.

IMPORTANT

The RMAN script section is not written to the backup specification file but is kept in memory until the backup specification is either saved or manually edited by clicking the `Edit` button.

IMPORTANT

The RMAN script can be edited only after the Data Protector Oracle 8/9 backup specification has been saved.

- Specify the Pre- and Post-exec commands/scripts.

The Oracle8/9 backup options are accessible in the Data Protector GUI in the `Application Specific Options` window.

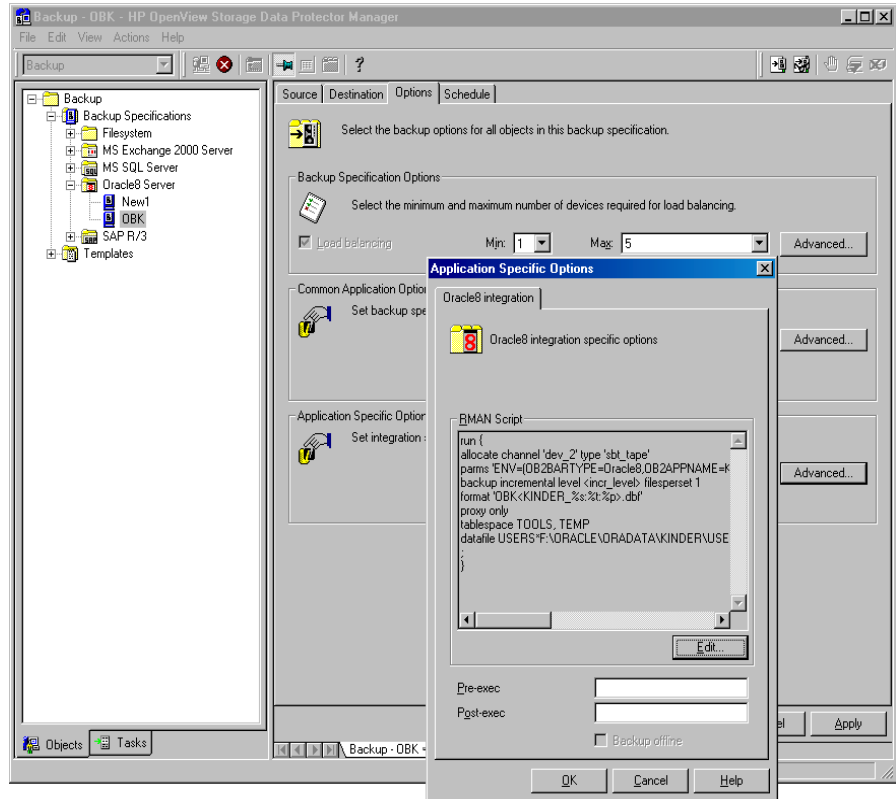
This window can be accessed from the `Options` property page of an Oracle8/9 backup specification by clicking the `Advanced` button.

TIP

If the Application Specific Options window is not displayed when clicking the Advanced button, this means that you have deselected all backup objects that were present in the backup specification, either by using the Data Protector GUI or by editing the RMAN script section.

Using the Data Protector GUI, select backup objects to display the Application Specific Options window when clicking the Advanced button.

Figure 1-9 Accessing the Application Specific Options Window



Specifying the Pre- and Post-exec commands

To specify Pre- or Post-exec commands, enter the command in the box next to the following two options:

- **Pre-exec**

Specifies a command/script that is started on the Oracle8/9 server system before backup. The command is started by the `ob2rman.exescrip` and runs under the account of the backup specification owner. See “Configuring an Oracle8/9 User in Data Protector” on page 27 for details about this user. The full pathname of the command must be provided in the backup specification.

- **Post-exec**

Specifies a command/script that is started on the Oracle8/9 server system after backup. The command is started by the `ob2rman.exe` script and runs under the account backup specification owner. See “Configuring an Oracle8/9 User in Data Protector” on page 27 for details about this user. The full pathname of the command must be provided in the backup specification.

Editing the Oracle8/9 RMAN Script

IMPORTANT

The RMAN script section of the Data Protector Oracle8/9 backup specification can be edited only after the Data Protector Oracle8/9 backup specification has been saved.

IMPORTANT

When editing the RMAN script section of the Data Protector Oracle8/9 backup specification, the Oracle8/9 manual configuration convention must be used and not the Oracle8/9 automatic configuration convention (introduced by Oracle9i).

When using RMAN scripts in Oracle 8/9 backup specifications, double quotes (“) must not be used - single quotes should be used instead.

Additionally, the format of RMAN commands must be as described on the following pages.

IMPORTANT

When editing the RMAN script section of the Data Protector Oracle8/9 backup specification, make sure that all manually-entered RMAN commands are *backup-related*. The RMAN script section of the Data Protector Oracle8/9 backup specification is not meant for any other Oracle8/9-related tasks, such as maintenance, configuration, registration, etc.

To edit an Oracle8/9 RMAN script, click the **Edit** button, edit the script and click the **Save** button to save changes to the script.

Refer to the *Oracle8i Recovery Manager User's Guide and References* for more information on Oracle8/9 RMAN commands.

The RMAN script created by Data Protector consists of the following parts (refer also to “Example of the RMAN Script” on page 47):

- The Oracle8/9 channel allocation together with the Oracle8/9 environment parameters' definition for every allocated channel:
 - ✓ If the `Blank Oracle8 Backup` template was used as a basis to create the backup specification, then the number of allocated channels is the same as the Data Protector concurrency number.

NOTE

Once the backup specification has been saved, changing the concurrency number does not change the number of allocated channels in the RMAN script. This has to be done manually by editing the RMAN script.

- ✓ If any other template was used as a basis to create the backup specification, Data Protector allocates one channel.

NOTE

When an Oracle8/9 channel is manually defined by editing the RMAN script, the environment parameters must be defined in the following format:

```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=<ORACLE_SID>,
OB2BARLIST=<Backup_Specification_Name>';
```

- Depending on the backup objects selection, an RMAN backup statement for the backup of the whole database instance, and/or for any combination of RMAN commands to back up tablespaces, the control file or the datafile. The backup statement consists of the following:
 - ✓ The Oracle8/9 type of backup (`full | incr1-4`) together with the RMAN `filesperiset=1` command.

Limitation

Only one file per set is possible.

- ✓ The Oracle8/9 format of the backup file in the following format:

```
format
'<Backup_Specification_Name><<ORACLE_SID>_%s:%t:%p>.db
f'
```

NOTE

When an Oracle8/9 format of the backup file is manually defined or changed by editing the RMAN script, any user-defined combination of the Oracle8/9 substitution variables can be *added* to the `%s:%t:%p` substitution variables, which are obligatory.

- ✓ The RMAN backup objects definitions (depending on the backup objects selection) as follows:
 - ❑ The RMAN database command, and/or
 - ❑ Any combination of the RMAN commands to back up tablespaces, the control file or the datafile (the RMAN `tablespace <tablespc1>[, <tablespc2>...]` command, the RMAN `include current controlfile` command or the RMAN `datafile <tablespace_name>*<datafile_name>` command).
- If the Archived Redo Logs were selected for a backup, an RMAN backup statement for the backup of Oracle8/9 archive logs. The backup statement consists of the following:
 - ✓ The RMAN `filesperiset=1` command.

Limitation

Only one file per set is possible.

- ✓ The Oracle8/9 format of the backup file in the following format:

```
format  
'<Backup_Specification_Name><<ORACLE_SID>_%s:%t:%p>.db  
f'
```

NOTE

When an Oracle8/9 format of the backup file is manually defined or changed by editing the RMAN script, any user-defined combination of the Oracle8/9 substitution variables can be *added* to the %s:%t:%p substitution variables, which are obligatory.

-
- ✓ The RMAN `archive log all` command.

IMPORTANT

It is not permitted to manually add additional RMAN backup commands, it is only possible to edit them.

-
- If an appropriate template was selected, or if the statement was manually added, the RMAN sql statement to switch the Online Redo Logs before backing up the Archived Redo Logs:

```
sql 'alter system archive log current';
```

- If an appropriate template was selected, or if the statement was manually added, the RMAN statement to delete the Archived Redo Logs after they are backed up:

```
archive log all delete input;
```

Example of the RMAN Script

The following is an example of the RMAN script section as created by Data Protector based on the Blank Oracle8 Backup template, after the whole database instance selection:

```
run {  
  
allocate channel 'dev_0' type 'sbt_tape' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';  
  
allocate channel 'dev_1' type 'sbt_tape' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';  
  
allocate channel 'dev_2' type 'sbt_tape' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1)';  
  
backup incremental level <incr_level> filesper set 1  
  
format 'New1<DIPSI_%s:%t:%p>.dbf'
```

Integrating Oracle8/9 and Data Protector

Configuring an Oracle8/9 Backup

```
database
include current controlfile
;
backup filesperset 1
format 'New1<DIPSI_%s:%t:%p>.dbf'
archivelog all
;
}
```

Creating Copies of Backed Up Objects

Oracle 8i/9i Duplex Mode

Oracle 8i and 9i support the duplex mode, which allows you to create copies of every backed up object to a separate backup device. To enable the duplex feature, perform the following steps:

1. Add the following command to the RMAN script before any allocate channel command:

```
set duplex=<on | 2 | ... >
```

IMPORTANT

If more than one allocated channel is used, it may happen that some original and copied objects are backed up to the same medium. To prevent this, you should use only one allocated channel when backing up using the duplex mode.

2. Add the following parameter to every format string used for backup:

```
%c
```

3. Set the concurrency of each device used for backup to 1.
4. Set the MIN and MAX load balancing parameters according to the following formula:

```
<number of duplex copies>*<number of allocated channels>
```

Example

If the duplex is set to 2 and the backup runs with 2 allocated channels, then the MIN and MAX parameters should be set to 4.

IMPORTANT

If the MIN and MAX load balancing parameters are set to lower values, the backup will hang.

If the MIN and MAX load balancing parameters are set to higher values, it may happen that the original and copied objects are backed up to the same medium.

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a backup. The test verifies both parts of the integration, the Oracle8/9 side and the Data Protector side. In addition, the configuration is tested as well.

The procedure consists of checking both the Oracle8/9 and the Data Protector parts of the integration to ensure that communication between Oracle8/9 and Data Protector is established, that the data transfer works properly, and that the transactions are recorded either in the recovery catalog (if used) or in the control file.

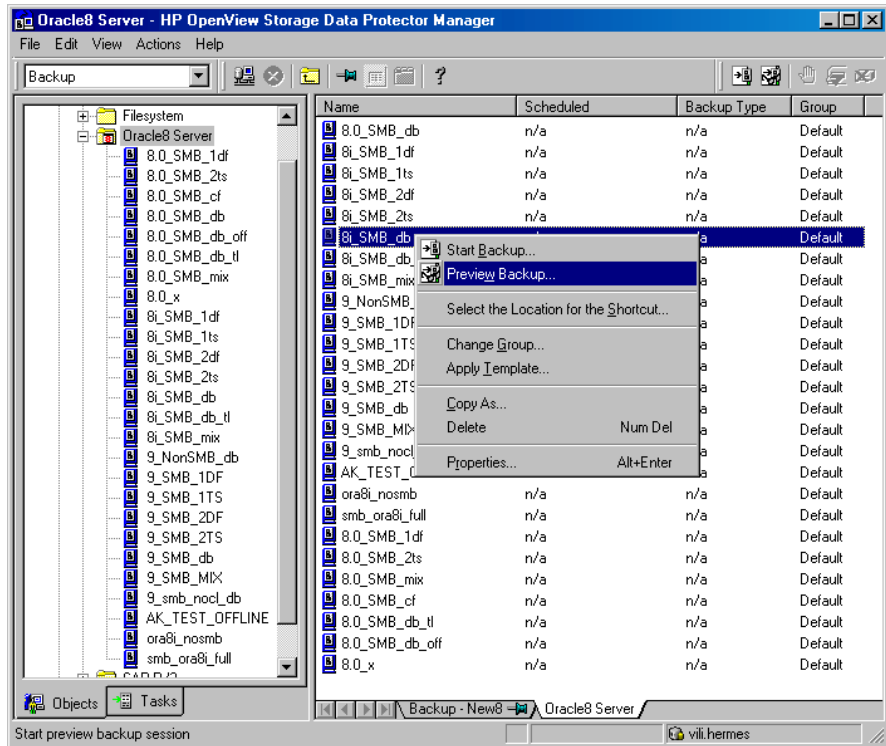
The test backup of the Oracle8/9 database is written to `/dev/null/` on the Oracle8/9 server system. Details of the test backup, such as media protection, backup user and backup status are registered in the Data Protector database and in the Oracle8/9 control files. Set the `Protection` option of your test backup specification to `None`.

Testing Using the Data Protector GUI

Follow the procedure below to test the backup of an Oracle8/9 backup specification:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Expand Oracle8 Server and right-click the backup specification you want to preview.
3. Click Preview Backup.

Figure 1-10 Previewing a Backup



Testing Using the CLI

A test can be executed from the command line on the Oracle8/9 Server system or on any Data Protector client system within the same Data Protector cell, provided that the system has the Data Protector User Interface installed.

You have to run the omnib command with the `-test_bar` option. Execute the following command:

On HP-UX and Solaris systems:

```
/opt/omni/bin/omnib -oracle8_list \  
<backup_specification_name> -test_bar
```

On other UNIX systems:

```
/usr/omni/bin/omnib -oracle8_list \  
<backup_specification_name> -test_bar
```

What Happens?

The session messages are displayed on the screen during the command execution, while the following happens:

1. The `ob2rman.exe` script is started, which then starts the Oracle8/9 RMAN backup command.

The Oracle8/9 Target Database is backed up to `/dev/null/` on the Oracle8/9 server system without using Data Protector for data transfer. This is how the Oracle8/9 side of the integration is checked.

2. The Data Protector `testbar2` command is started by the `ob2rman.exe` script, which checks:
 - The communication within Data Protector
 - The syntax of the Oracle8/9 backup specification
 - If the devices are correctly specified
 - If the required media reside in the devices

Backing Up an Oracle8/9 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. Special attention is required to reach a consistent state with an online backup.

A decision about your database backup strategy depends on a number of factors. If the database must be opened and available all the time, then online backup is your only choice. If you can afford to have the database offline at a certain time, then you are more likely to make periodic offline backups of the entire database, supplementing them with online backups of the dynamically changing tablespaces.

Oracle8/9 Offline

An offline backup of a database is a backup of the datafiles and control files which 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 either closed or mounted.

The offline backup of an Oracle8/9 Target Database can be performed using a Data Protector filesystem backup specification or a Data Protector Oracle8/9 backup specification, based on which Data Protector automatically generates and executes the RMAN script. The Data Protector Disk Agent is used in the first case, and the Data Protector Oracle8/9 Integration software component in the second case.

Typically, you would perform an offline backup of the entire database, which must include all datafiles and control files, while the parameter files may be included optionally.

The whole offline database backup is performed as follows:

1. Shut down the database cleanly.
A clean shutdown means that the database is not shut down using the ABORT option.
2. Mount the database if you are backing it up using RMAN.
3. Back up all datafiles, control files and, optionally, parameter files.
4. Restart the database in the normal online mode

Oracle8/9 Online

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

The backup of an opened database is inconsistent, because portions of the database are being modified and written to disk while the backup is progressing. Such changes to the database are entered into the online redo logs as well. A database running in ARCHIVELOG mode enables the archiving of online redo logs. In the case of a restore, this feature is essential to bring a database to a consistent state as part of the entire restore process.

When using an online backup, the following must be done in order to bring the database to a consistent state:

1. Restore the database files (which are inconsistent) to disk
2. Restore the Archived Redo Logs to disk
3. Perform a database recovery, which requires applying the Archived Redo Logs. This is an Oracle8/9 operation.

An Oracle8/9 online database backup can be performed using the Oracle8/9 RMAN utility or Data Protector GUI. In the latter case, Data Protector creates and executes the RMAN script automatically based on data entered in the Data Protector GUI. During an Oracle8/9 online backup, the Oracle8/9 Target Database is open, while tablespaces, datafiles, control files, and archived redo logs are being backed up.

The database must operate in the ARCHIVELOG mode so that the current Online Redo Logs are archived to the Archived Redo Logs.

In order to be able to restore to a consistent state, a complete online database backup must be performed as follows:

1. Put the tablespaces in backup mode.
2. Back up tablespaces, datafiles, and control files.
3. Put the tablespaces back in normal mode.
4. Back up the Archived Redo Logs.

IMPORTANT

Before you run an Oracle8/9 online backup, make sure that the database is really operating in ARCHIVELOG mode. This can be done on the Oracle8/9 server system by starting the Oracle8 Server Manager or Oracle9i SQL Plus and issuing the following command:

```
archive log list;
```

If the Oracle8/9 Target Database is not operating in the ARCHIVELOG mode, shut it down first and mount it again. Issue the following command at the Oracle8 Server Manager or Oracle9i SQL Plus prompt:

```
alter database archive log;  
archive log start;  
alter database open;
```

Now you are ready to run an online backup of the Oracle8/9 database, using any of the following methods:

- Schedule the backup of a saved Oracle8/9 backup specification using the Data Protector Scheduler. See “Scheduling a Backup” on page 55.
- Start an interactive backup of the Oracle8/9 backup specification. See “Starting an Interactive Backup” on page 57.
- Start a backup on the Oracle8/9 server using either Oracle8/9 Recovery Manager or Oracle8/9 Enterprise Manager. See “Using the Oracle8/9 Recovery Manager (RMAN)” on page 60.

Backup Procedure The following happens when you start a backup using the Data Protector user interface:

1. Data Protector executes the `ob2rman.exe` script on the client. This command starts the Recovery Manager (RMAN) and sends the Oracle8/9 RMAN Backup Command Script to the standard input of the RMAN command.
2. The Oracle8/9 RMAN contacts the Oracle8/9 Server, which contacts Data Protector via the Database Library interface and initiates a backup.
3. During the backup session, the Oracle8/9 Server reads data from the disk and sends it to Data Protector for writing to the backup device.

Messages from the Data Protector backup session and messages generated by Oracle8/9 are logged to the Data Protector database.

**Automatic
Recovery Catalog
Database Backup**

A backup of the Oracle8/9 recovery catalog is performed automatically following each Oracle8/9 Target Database backup. Using the standard Oracle8/9 export utility, the Data Protector `ob2rman.exe` script starts an export of the Oracle8/9 recovery catalog to a file which is then backed up by Data Protector.

IMPORTANT

When backing up or restoring the catalog database, the value of the NLS_LANG variable, which is set for the target database, is used as the default NLS_LANG value.

If you want to use a different NLS_LANG value for importing or exporting the catalog database, you can specify this in the Oracle instance configuration file in the section Environment. Refer to the section “Setting, Retrieving and Listing Data Protector Oracle8/9 Configuration Files’ Parameters Using the CLI” on page 15.

Deleting Data from the Recovery Catalog

When backing up an Oracle8/9 database using the recovery catalog database, all information about the backup, restore and recovery of the database is stored in the recovery catalog. This information is used by RMAN during the restore. If you overwrite or format the media on which this data is backed up, Data Protector exports the object from the Data Protector database. You must manually delete the data from the recovery catalog while logged on to RMAN. Refer to the *Oracle8i Recovery Manager User’s Guide and References* for detailed information about deleting data from the recovery catalog.

NOTE

You can obtain the primary keys of the records to be deleted by issuing the `list` command, as shown in the following example:

If the (data) file is stored on a disk, issue the following command:

```
list backup of tablespace temp;
```

Primary keys in the `temp` tablespace are listed. You need these keys to delete the records, as shown in the following example:

```
allocate channel for delete type 'sbt_tape';  
change backupset <primary_key> delete;  
release channel;
```

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, including the backup of the Archived Redo Logs, which is required in case you need a recovery to a particular point in time.

For example, you may decide to perform daily backups and make multiple copies of the online redo logs and the Archived Redo Logs to several different locations.

An example of scheduling backups of production databases:

- Weekly full backup
- Daily incremental backup
- Archived Log backups as needed

To schedule an Oracle8/9 backup specification, proceed as follows:

1. In the Data Protector Manager window, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click Oracle8 Server.

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 Incr 4. Refer to RMAN documentation for details on incremental backup levels. The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to Full.

In the case of zero downtime backup, but only for ZDB disk or ZDB disk/tape backups (instant recovery enabled), specify the Split mirror/snapshot backup option.

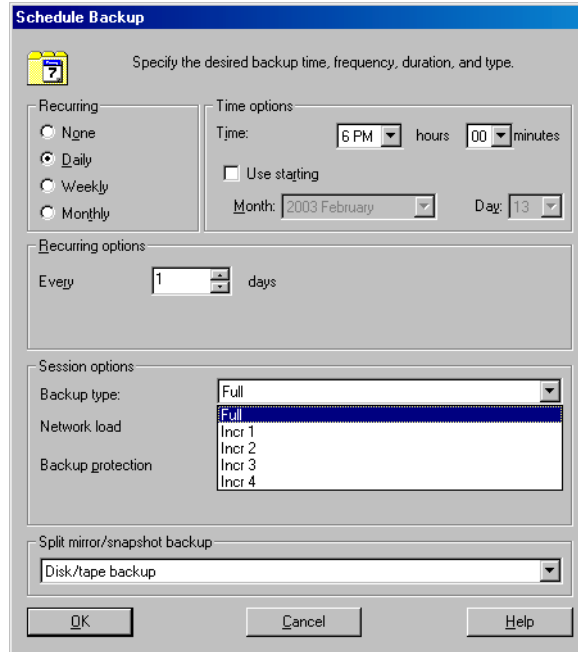
See Figure 1-11 on page 57.

6. Click OK to return to the Schedule property page.

7. Click Apply to save the changes.

Figure 1-11

Scheduling Backups



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, while the corresponding backup specification is scheduled at a later time.

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

When you start a backup, Data Protector invokes the `ob2rman.exe` script on the Oracle8/9 Server 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 Oracle8/9 backup specification:

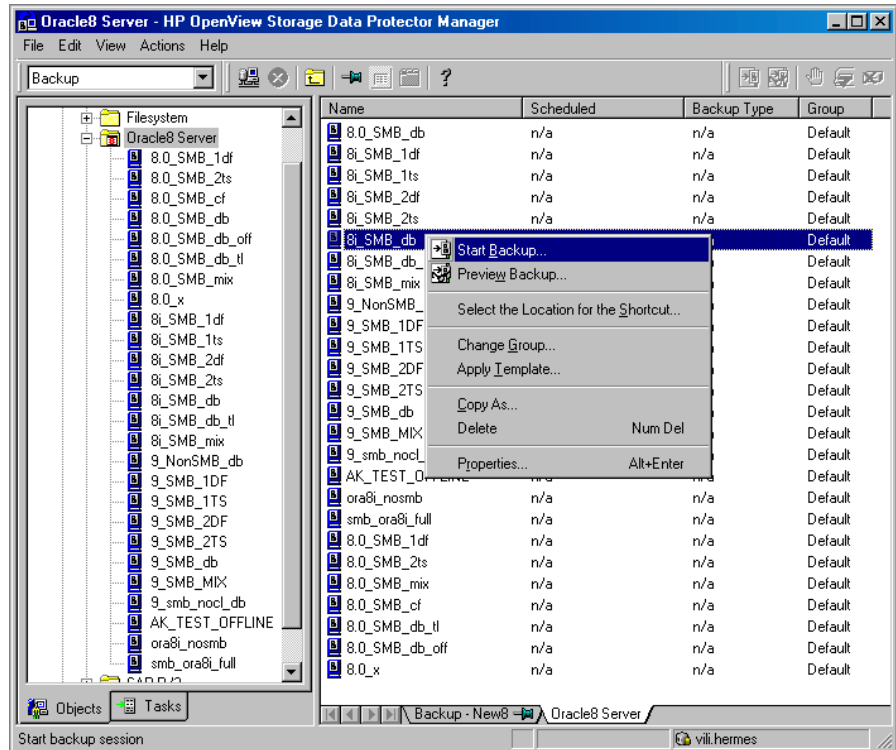
1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then expand Backup Specifications and then Oracle8 Server.
3. Right-click the backup specification, then click Start Backup.

In the Start Backup dialog box, select the backup type and network load. The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to Full.

In the case of zero downtime backup, but only for ZDB disk or ZDB disk/tape backups (instant recovery enabled), specify the Split mirror/snapshot backup option.

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

Figure 1-12 Starting an Interactive Backup



Running Backup Interactively Using the CLI

An interactive backup can also be started from the CLI. Change to the /opt/omni/bin directory (HP-UX and Solaris systems), or the /usr/omni/bin directory (other UNIX systems)

on an Oracle8/9 Server system and run the following command:

```
omnib -oracle8_list <backup_specification_name> [-barmode <Oracle8Mode>] [list_options]
```

You can select among the following *list_options*

- protect {none | weeks *n* | days *n* | until *date* | permanent}
- load {low | medium | high}

```
-crc  
-no_monitor  
Oracle8Mode = { -full | -incr1 | -incr2 | -incr3 |  
-incr4 }
```

Refer to the man page for details.

Example

To start a backup using an Oracle8/9 backup specification called RONA, execute the following command:

```
omnib -oracle8_list RONA
```

Using the Oracle8/9 Recovery Manager (RMAN)

The Oracle8/9 Recovery Manager (RMAN) utility is an Oracle8/9 CLI that allows you to perform a backup, restore, or recovery of Oracle8/9 database objects. To start an Oracle8/9 backup using RMAN, an Oracle8/9 backup specification must be created.

See “Configuring an Oracle8/9 Backup” on page 37 for information on how to create an Oracle8/9 backup specification. To start an Oracle8/9 backup using RMAN, follow these steps:

1. Connect to the Oracle8/9 Target Database specified in the backup specification.

To connect to an Oracle8/9 Target Database that uses the recovery catalog, enter the following command:

```
<ORACLE_HOME>/bin/rman target <Target_Database_Login>  
rcvcat <Recovery_Catalog_Login>
```

To connect to an Oracle8/9 Target Database without using the recovery catalog, enter the following command:

```
<ORACLE_HOME>/bin/rman target <Target_Database_Login>  
nocatalog
```

See the “Glossary” on page G- 1 for details on the login information syntax.

2. Allocate the Oracle8/9 channels.

Allocating a channel tells RMAN to initiate an Oracle8/9 Server process for backup, restore, or recovery on the Oracle8/9 Target Database. For example:

```
allocate channel 'dev_0' type 'disk';  
or  
allocate channel 'dev_1' type 'sbt_tape';
```

where you specify the backup directly to disk in the first case and directly to tape in the second case. Note that if Oracle8/9 is linked to Data Protector, Data Protector will perform the backup to the tape in the second case.

If you specify more than a single `allocate channel` command, RMAN can establish multiple logon sessions and conduct multiple backup sets in parallel. This "parallelization" of backup and restore commands is handled internally by RMAN.

To use Data Protector backup media, specify the channel type `SBT_TAPE`.

3. You must specify the `parms` operand in the following form:

```
parms 'ENV(OB2BARTYPE=Oracle8,  
OB2APPNAME=<ORACLE_SID>,OB2BARLIST=<backup_  
specification_name>)' ;
```

Note that the RMAN script will not work without the above parameters being specified in this form.

4. You must specify the backup command operand `filesperset`.

`filesperset` represents the number of files per backup set. Data Protector limits this number to 1. Each channel can operate on only one file, and there are as many channels as files used.

5. Optionally, you may specify `format` as follows:

If you have created and saved a backup specification named `credo` for backing up an Oracle8/9 database identified by the Oracle8/9 instance called `alma`, you would enter the following string:

```
format 'credo<alma_%s:%t>.dbf'
```

See the *Oracle8i Recovery Manager User's Guide and References* for information on substitution variables. The Oracle8/9 channel format specifies which Oracle8/9 backup specification to use for the backup.

6. Optionally, you may specify `backup incremental level`.

Note that a Data Protector full backup performs the same operation as an incremental level 0 backup type in the Oracle8/9 RMAN scripts. They both back up all the blocks that have ever been used.

In order to run a backup using RMAN, start the Recovery Manager by typing the following from the `<ORACLE_HOME>` directory (if you are using recovery catalog):

```
bin/rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>
```

Some examples of RMAN scripts that must be executed from the RMAN> prompt are listed below:

Backing Up a Single Channel

To back up the Oracle8/9 instance `credo`, using a backup specification named `alma`, enter the following command sequence:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  backup  
  incremental level 0  
  filesperset 1  
  format 'alma<credo_%s:%t>.dbf' database;  
}
```

Backing Up Three Channels in Parallel

The RMAN backup script for backing up the database by using three parallel channels for the same backup specification would look like this:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_1' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  allocate channel 'dev_2' type 'sbt_tape'  
  parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  backup  
  incremental level 0  
  filesperset 1  
  format 'alma<credo_%s:%t>.dbf' database;  
}
```

Backing Up All Archived Logs and Tablespaces

If you want to back up the Archived Redo Logs and the tablespace SYSTEM and RONA of the previous database using three parallel channels and a backup specification named alma, the RMAN script should look like this:

```
run {
allocate channel 'dev_0' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_1' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_2' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
backup
incremental level 0
filesperset 1
format 'alma<credo_%s:%t>.dbf'
tablespace SYSTEM, RONA
archivelog all;
}
```

Backing Up Particular Archived Logs

To back up all Archived Redo Logs from sequence #5 to sequence #105 and delete the Archived Redo Logs after backup of the instance named alma is complete, run the following script:

```
run {
allocate channel 'dev_0' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_1' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
allocate channel 'dev_2' type 'sbt_tape'
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
backup
incremental level 0
filesperset 1
(archivelog low logseq 5 high logseq 105 thread 1 all delete input
format 'alma<credo_%s:%t>.dbf');
}
```

If the backup fails, the logs are not deleted.

Including Control File in a Backup Specification

The current control file is automatically backed up when the first datafile of the system tablespace is backed up. The current control file can also be explicitly included in a backup, or backed up individually. To include the current control file after backing up a tablespace named COSTS, run the following script:

```
run {  
  
allocate channel 'dev_0' type 'sbt_tape'  
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  
  
allocate channel 'dev_1' type 'sbt_tape'  
'parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  
  
allocate channel 'dev_2' type 'sbt_tape'  
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  
backup  
  
incremental level 0  
  
filesperset 1  
  
(tablespace COSTS current controlfile  
format 'alma<credo_%s:%t>.dbf');  
  
};
```

Backing Up While Allowing for Some Corrupted Blocks

The set `maxcorrupt` command determines the number of corrupted blocks per datafile that can be tolerated by RMAN before a particular backup will fail.

If a backup specification named `alma` backs up the database and allows for up to 10 corrupted blocks per datafile `/oracle/data1.dbs`, then the appropriate RMAN script would be:

```
run {  
  
allocate channel 'dev_0' type 'sbt_tape'  
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  
  
allocate channel 'dev_1' type 'sbt_tape'  
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';  
  
  
allocate channel 'dev_2' type 'sbt_tape'  
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=credo,OB2BARLIST=alma)';
```



```
backup  
incremental level 0  
filesperset 1  
format 'alma<credo_%s:%t>.dbf'  
database;  
set maxcorrupt for datafile  
'/oracle/data1.dbs' to 10;}
```

Restoring Oracle8/9 Databases

You can restore the database using one of the following tools within Data Protector:

Restore Methods

- Data Protector Restore GUI for Oracle
- Oracle Recovery Manager (RMAN). For information on how to use RMAN to restore the database see “Restoring Oracle8/9 Using RMAN” on page 80.

Restorable Items

It is possible to restore the following items using both the Data Protector Restore GUI for Oracle and RMAN.

- Control files
- Datafiles
- Tablespaces
- Databases

Restoring Oracle8/9 Using the Data Protector Restore GUI for Oracle

For the restore, RMAN scripts are generated with necessary commands, depending on selections made in the GUI. If you want to perform additional actions, you cannot edit the RMAN restore script, but you can perform them manually from RMAN itself.

Restoring Database Items in a Disaster Recovery

In a disaster recovery situation, database objects must be restored in a certain order. The following table shows you in which order database items must be restored. Under normal conditions it is possible to restore database items in any order.

Table 1-2 Oracle8/9 Data Restore Order in Disaster Recovery

Oracle version	Data restore order
Oracle8/8i	<ol style="list-style-type: none"> 1. Restore the recovery catalog database 2. Restore the control file 3. Restore the entire database or data items
Oracle9	<ol style="list-style-type: none"> 1. Restore the control file from automatic backup 2. Restore the database or data items <p>OR:</p> <ol style="list-style-type: none"> 1. Restore the recovery catalog database. 2. Restore the control file 3. Restore the database or data items

Changing The Database State

Before you restore any database item you need to ensure that the database is in the correct state. The following table tells you which state the database needs to be in to restore a particular type of database item:

Table 1-3 Required Database States

Item to restore	Database state
Control file	NoMount
All other items	Mount

To put the database into the correct state carry out the following procedure:

Open a command window and enter the following:

```
sqlplus/nolog
```

In the SQL> prompt, enter:

```
SQL>connect/as sysdba
```

```
SQL>shutdown immediate
```

If you are restoring a control file put the database into NoMount state.

```
SQL>startup nomount
```

If you are restoring any other database item put the database in Mount state.

```
SQL>startup mount
```

Restoring the Recovery Catalog Database

The Oracle recovery catalog database is exported using the Oracle export utility to a binary file and backed up by Data Protector. This file has to be restored back to the disk and then imported into the Oracle database using the Oracle import utility. Data Protector provides a facility to do this automatically using the Oracle8/9 Integration. Carry out the following procedure to restore the recovery catalog database:

1. Ensure the recovery catalog database exists and is empty. To check if the recovery catalog database was used as a repository during backup execute the following from the command line on the client system:

```
# ./util_cmd -getconf Oracle8 <SID>
```

If the recovery catalog database was selected as one of the Data Protector backup options when the original database backup was configured, this command gives you output to the following:

```
ORACLE_HOME='/app/oracle9i/product/9.2.0.1.0';  
TGTLogin='EIBBKIBBEIBBQDBBOHBBCHBBPHBBBIBBCHBBEIBBB  
FBBFGBBFFBBDFFBB';  
RCVLogin='DIBBOHBBCHBBPHBBQDBBDIBBOHBBCHBBPH  
FBBFGBBFFBBDFFBB';  
ORACLE_VERSION='9.2.0';  
Configuration read/write operation successful.
```

If the RCVLogin entry is present in the output, the recovery catalog database was used.

2. Identify the recovery catalog database owner and the instance name of the recovery catalog database using the Data Protector Restore GUI for Oracle.

3. Ensure that the recovery catalog database is in the Open state. In the command line enter:

```
sqlplus /nolog
```

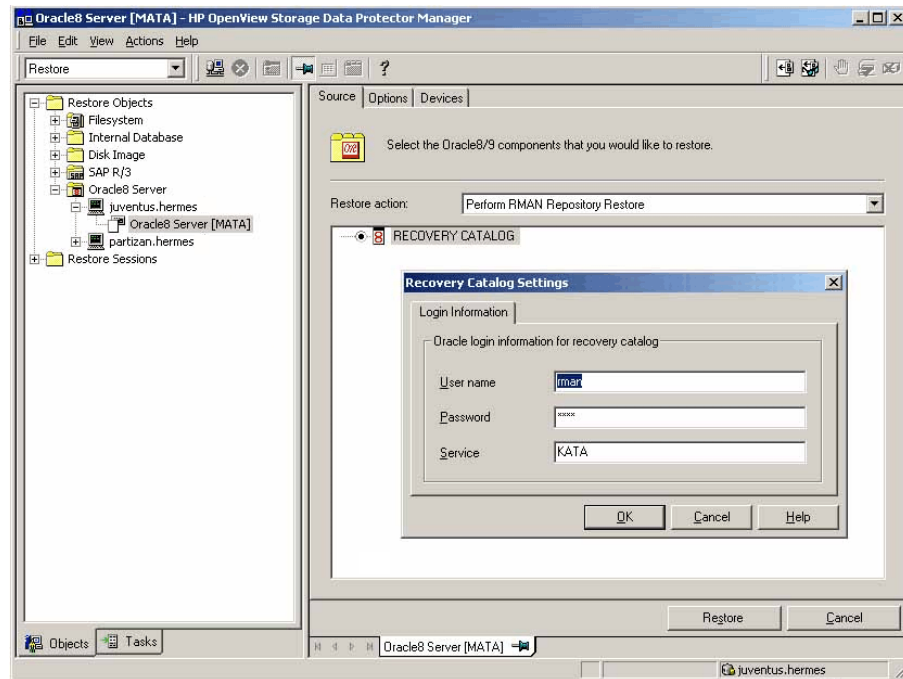
In the SQL> prompt, enter

```
SQL>connect sys/password@CATALOG_NAME as sysdba;
```

```
SQL>select status from v$instance;
```
4. In the Data Protector Manager, switch to the Restore context and select the Oracle server and Oracle instance for which you need to restore the recovery catalog database.
5. In the Results Area, select Recovery Catalog. If you want to change or enter the recovery catalog login strings, right-click on the Recovery Catalog and select Change Settings.

Figure 1-13

Recovery Catalog and Change Recovery Catalog Settings Dialogs



6. Enter the recovery catalog database login strings in the Recovery Catalog Settings dialog.

7. Select the `Options` pane on the Data Protector Restore GUI for Oracle. Select the `Session ID` from the `Session ID` drop-down list. For further information see “Restore and Recovery Options” on page 76.
8. Enter the user name and password to the recovery catalog database in the `User name` and `User group` fields on the `Options` pane.
9. Click the `Restore` button. At this point the recovery catalog database is restored. You can now proceed to restore your control file.

Restoring the Control File

The control file contains all the information about the database structure. If the control file has been lost, you must restore it before you can restore any other part of the database.

IMPORTANT

If you are using Oracle8i it is not possible to restore the control file immediately. You need to restore the recovery catalog database first, for information on how to do this see “Restoring the Recovery Catalog Database” on page 68. If you did not use the recovery catalog database when backing up objects in the database, call Oracle Support for help recovering the control file.

The procedure for restoring the control file is as follows:

1. Open a SQLPlus window and put the database in the nomount state.
2. In the Data Protector GUI switch to the `Restore` context. Expand `Restore` and select the Oracle Server and the Oracle instance for which you need to restore the control file.
3. In the `Results` area select the control file.
4. Click `Restore`.

If you want to restore the control file and the entire database to a different client refer to “Restoring the Database to Another Client” on page 71 for more information about this dialog.

You can now proceed to recover the Oracle database or items within the database.

Restoring the Database to Another Client

It is possible to restore the database and the control file to a machine other than the one where they originally resided. To do this, carry out the following procedure:

1. In the `Source` pane select the database.
2. In the `Options` pane select `Perform Restore` from the `Restore Action` drop-down list.
3. Select the client to which to restore the database by selecting the name of the client from the `Restore to client` drop-down list.
4. Click `Restore`.

At this point both the control file and the database will be restored to the chosen client.

Restoring Oracle Database Objects

Before you restore Oracle database objects you need to ensure that you have an up-to-date version of the recovery catalog database and the control file. These contain the database structure information. If you do not already have up-to-date versions of these files you need to restore them, refer to “Restoring the Recovery Catalog Database” on page `HIDDEN` and “Restoring the Control File” on page 70 for more information.

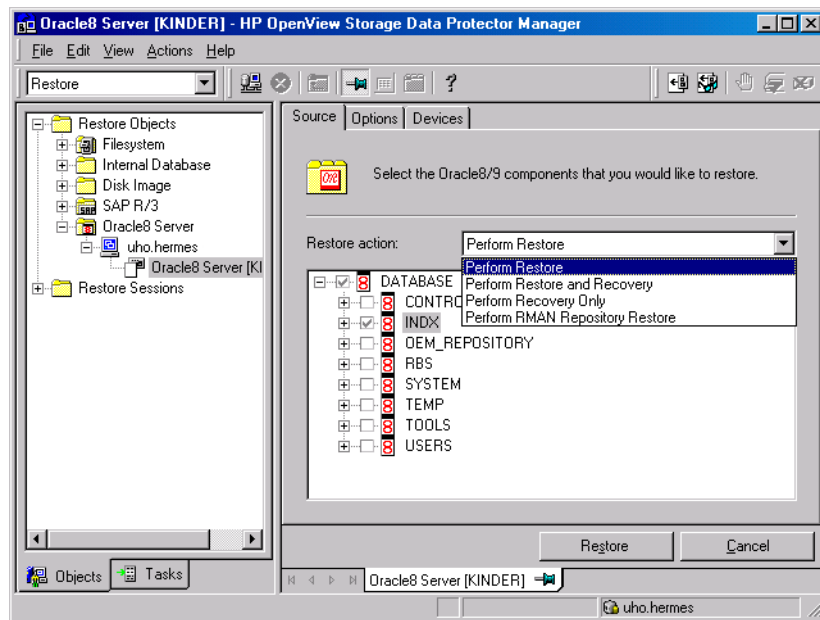
To restore Oracle database items proceed as follows:

1. First ensure that the database is in the correct state before you begin to restore database items. Details of this can be found in “Changing The Database State” on page 67.
2. In the `Context List` select `Restore`.
3. In the `Scoping Pane` select the Oracle server instance you would like to restore. Data Protector now displays all the Oracle objects that were backed for that instance of Oracle.

IMPORTANT

Data Protector does not check if the Oracle objects were all backed up and that they can all be restored. You need to be aware of what was in the backup. You can check this by going to the Backup context and selecting the backup specification for the Oracle database you are restoring.

Figure 1-14 Source Pane



4. Select the type of restore action you wish to perform from the Restore action drop-down list. For details on these options refer to “Restore and Recovery Options” on page 76.

IMPORTANT

If you do not select Perform Restore & Recovery or Perform Recovery Only in the Restore actions drop-down list you will have to recover database items manually from the command line. For more information about using RMAN refer to “Restoring Oracle8/9 Using RMAN” on page 80.

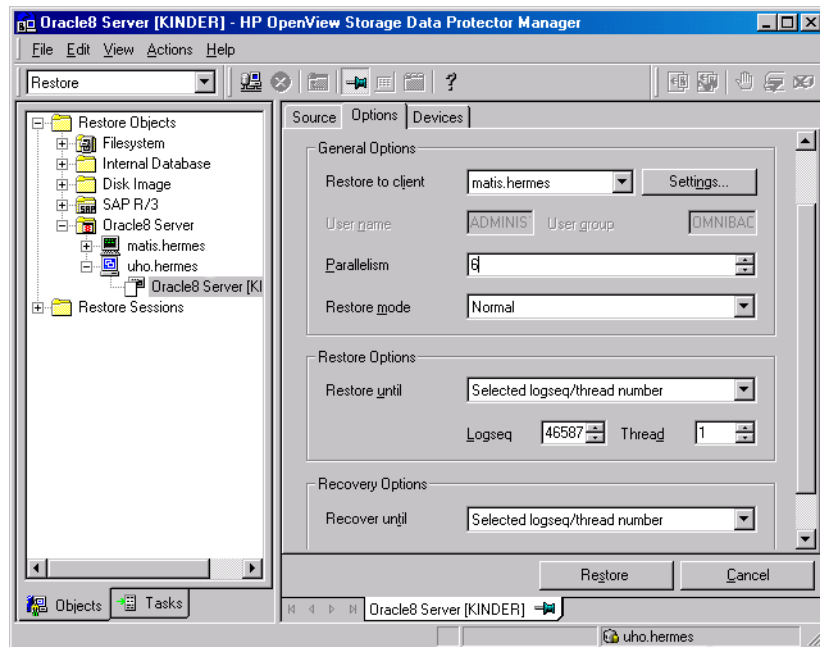
5. In the Results Area click on the items you would like to restore. If you are restoring datafiles it is possible to restore the files to a new location by right-clicking on the filename. A dialog will appear into which you can enter the new datafile location.

If you select to restore the datafile to another location or with a different name, the datafile will be restored to the selected location. If you want Oracle to use the datafile under the new name, you must issue a switch statement afterwards using Oracle tools. For more details refer to “*Recovery Manager User’s Guide and Reference*”.

6. In the Options pane enter the restore or recovery details. See the “Restore and Recovery Options” on page 76 for more information about the fields in this pane.

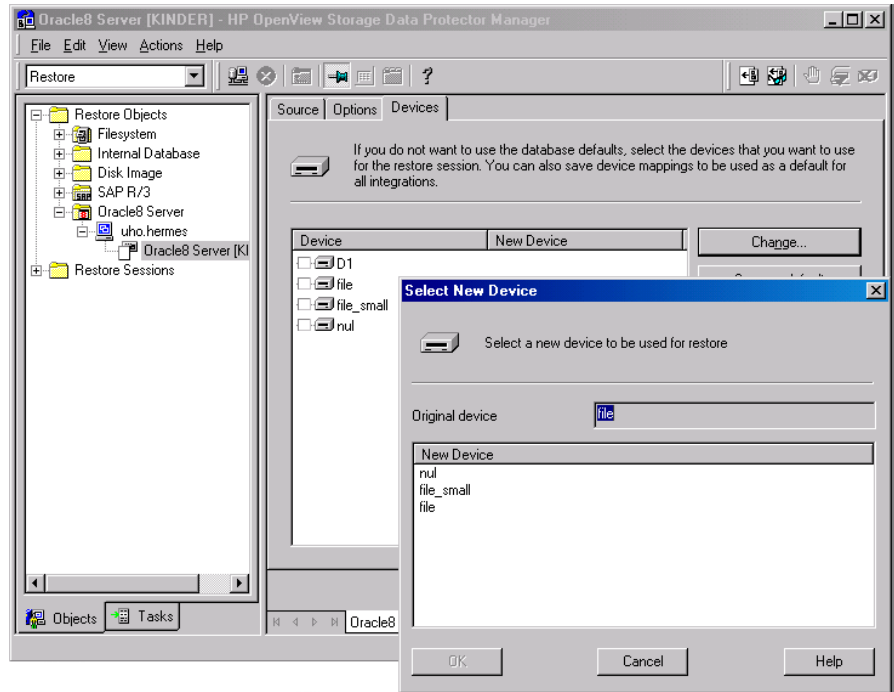
Figure 1-15

Options Pane



7. Go to the Devices pane and select the devices to be used for the restore. You can select devices from which to restore the database other than those which were used for the original backup, although Data Protector defaults to the original device on which the backup was made. For more information on the Devices pane, press **F1**.

Figure 1-16 Devices Pane



8. If you would like to change the device from which an item is restored, select your desired device and click *Change*.
9. After selecting all the devices click *Restore*. The restore procedure starts.

When the restore session starts, messages are displayed in the *Computer Output* pane. When the session is finished a message is issued in the *Session Information* dialog.

After Recovering the Database

Once you have restored the lost database objects you need to ensure that the database is in the correct state.

If you used one of the options on the *Source* pane containing the word "Recovery" then the database will automatically be put into *Open* state by Data Protector.

If you are performing an Oracle database restore and recovery until point in time, and the session has finished successfully, reset the database in order to register the new incarnation of database in the recovery catalog. Connect to the target and recovery catalog database using RMAN and reset the database:

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>RMAN> RESET DATABASE;  
  
RMAN> exit
```

If you did not choose to use Data Protector to Recover database items, you need to carry out the following procedure after the database has been restored:

Open a command line window and enter the following commands:

```
sqlplus/nolog  
SQL>recover database;  
SQL>connect/as sysdba  
SQL>alter database open;
```

Restoring Tablespaces and Datafiles

You can restore tablespaces and datafiles as follows:

1. Open a command line window and enter the following commands if you have the database in the Open state:

```
sqlplus/nolog  
SQL>connect/as sysdba  
SQL>alter database datafile '<datafile name>' offline;
```

If you are restoring a tablespace enter:

```
SQL>alter tablespace '<tablespace name>' offline;
```

2. When the restore has been completed put the datafiles and tablespaces back online with the following procedures:

Open a command line window and enter the following commands:

```
sqlplus/nolog  
SQL>connect/as sysdba
```

If you are restoring a datafile enter:

```
SQL>alter database datafile '<datafile name>' online;
```

If you are restoring a tablespace enter:

```
SQL>alter tablespace '<tablespace name>' online;
```

Restore and Recovery Options

Source Pane

The following describe each of the options on the `Source` pane. This pane is used to define the combination of `Restore` and `Recovery` you would like to perform with the GUI.

In the context of Data Protector “Restore” means to restore the datafiles. Users can select which database, tablespace or datafiles they would like to restore and up to which point in time they would like them to be restored. “Recover” means applying all the redo logs. The user can select which redo logs to apply according to SCN number, logseq, or can apply all the redo logs to the time of the last backup.

The options below are available on the `Restore` action drop-down list.

Perform Restore

This option specifies that the user will restore the database objects using Data Protector and will then perform the recovery manually using RMAN. For information on how to recover database items using RMAN see “Restoring Oracle8/9 Using RMAN” on page 80.

Perform Restore and Recovery

This option means that the user will perform both the `Restore` and the `Recovery` from the GUI.

Perform Recovery Only

The option specifies that the user will perform the `recovery` only using the GUI.

This option is meant to be used after `Instant Recovery` has been successfully completed.

Perform RMAN Repository Restore

This option is used to restore the recovery catalog or the control file if the database objects are not available in the `Source` property page.

The `Source` window allows you to specify another location to which to restore an Oracle database file. This is done using the `Restore As` dialog which appears when you right-click the mouse on a database item listed in the `Results Area`.

Options Pane

The following describe each of the fields on the `Options` pane.

Restore to client

This option specifies the name of the Oracle server to which the user wants to restore the database item. This defaults to the original backup server.

NOTE

When restoring to another host, you should select the database SID on the `Source` pane and then select the host to which to restore the database on the `Options` pane. At this point Data Protector automatically restores the control file and then all the objects in the database.

User name

Use this field to enter the Oracle User Name. When the user account has been specified the Oracle restore process can start. The user needs to be a member of the DBA group.

User group

The User group the user in the `User name` field belongs to. This has to be the Oracle DBA group.

The User name and the User group must be the same as defined in the backup ownership.

Parallelism

This field is used to specify the number of concurrent data streams that can read from the backup device. If you do not enter a value, the number of parallel streams defaults to one.

To optimize restore performance, specify the same number of data streams as were used during the backup. If, for example you set the backup concurrency to 3 then set the number of parallel data streams to 3

as well. Note that if a very high number of parallel data streams are specified this may result in a resource problem because too much memory is being used.

Restore mode

This drop-down list allows you to specify which type of restore you would like perform. The options are:

- Normal

This option should be used when either a Standalone or Zero Downtime Backup was made of the database using a version of Data Protector which is older than 5.0.

- Proxy copy

This option should be used when the original Oracle backup was made using the Oracle Proxy copy feature, such as ZDB/Snapshot backups of Oracle8i/9 using Data Protector version 5.0.

This option is disabled when you are performing a restore after Instant Recovery.

Restore until

The options in this drop-down list allow you to specify to which point in time you would like the restore to be performed.

This option is disabled when you are performing a restore after Instant Recovery.

- Now

This is the default option. Data Protector finds the most recent backups to use when performing the restore.

- Selected time

Using this option you can specify an exact time to which Data Protector restores database objects.

IMPORTANT

Check the internal database to find out what time the backup you are restoring was completed. Specify the time immediately after the backup was finished.

- Selected logseq/thread number

A logseq number is a redo log sequence number. You can enter a particular redo log sequence number which will act as an upper limit of redo logs to restore.

- Selected SCN number

This option allows you to enter the SCN number to which to perform the restore.

Recover until

The options in this drop-down list allow you to specify to which point in time you would like the recovery to be performed.

- Now

This is the default option. Data Protector finds the most recent backups to use when performing the recovery.

- Selected time

Using this option you can specify an exact time to which Data Protector restores database objects.

- Selected logseq/thread number

A logseq number is a redo log sequence number. You can enter a particular redo log sequence number which will act as an upper limit of redo logs to recovery.

This option is used for Point in Time Restore.

- Selected SCN number

This option allows you to enter the SCN number to which to perform the recovery.

Target DB login

This option lets you change the target database login information, i.e. the username and password of the user who has SYSDBA privileges and the service name to which Data Protector should connect.

Restoring Oracle8/9 Using RMAN

Data Protector acts as a media management utility for the Oracle8/9 system, therefore the Oracle Recovery Manager (RMAN) can be used for a restore.

Refer to the *Oracle8i Recovery Manager User's Guide and References* for detailed information on how to perform database, tablespace, control file, and datafile restore and recovery.

Restoring the Recovery Catalog

Data Protector can restore the binary file which contains the logical backups of the Oracle8/9 recovery catalog. This file is made using the Oracle8/9 Export utility, which creates it by reading the Oracle8/9 database and writing the output to the binary file, which is then backed up by Data Protector.

This file can be restored back to the disk and then imported to the Oracle8/9 database by the Oracle8/9 Import utility.

To restore the Oracle8/9 recovery catalog, proceed as follows:

1. Login to the Oracle8/9 Target Database. Ensure that the recovery catalog is empty and that you have identified the Oracle8/9 recovery catalog owner.

If you have lost (deleted) the recovery catalog then you need to create it again. Data Protector determines the Oracle8/9 login information for the recovery catalog from the Data Protector Oracle8/9 configuration files. Refer to "Data Protector Oracle8/9 Configuration Files" on page 13.

2. Export the OB2APPNAME environment variable. Its value must be set to the Oracle8/9 SID of the target database, not of the Oracle8/9 recovery catalog:

- if you are using an sh - like shell, enter the following commands:

```
OB2APPNAME=" <ORACLE_SID>"  
export OB2APPNAME
```


- if you are using a csh - like shell, enter the following commands:

```
setenv OB2APPNAME "<ORACLE_SID>"
```

3. Switch to the

```
<Data_Protector_home>/lbin
```

directory and issue the following Data Protector command:

```
ob2rman.exe -restore_catalog -session <session_ID>
```

After issuing the above command, the `/var/opt/omni/tmp/rcvcat.exp` file is restored by the Data Protector `obkrestore` utility, which is then read by the Oracle8/9 import utility. The Oracle8/9 Import utility then restores the file back to the Oracle8/9 database.

Restoring Using Another Device

Data Protector supports the restore of Oracle8/9 database objects from devices other than those on which the database objects were backed up.

Specify these devices in the: `/etc/opt/omni/cell/restoredev` 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.

Example

Suppose you have Oracle8/9 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 intended to be used as a guideline.

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

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.
2. Install, configure, and initialize the database/application so that data on the Data Protector media can be loaded back to the system. Consult the documentation from the database/application vendor for a detailed procedure and the steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and in the section. See also the section of this manual about the Data Protector Restore GUI for Oracle for information about using this to restore database items, "Restoring Oracle8/9 Using the Data Protector Restore GUI for Oracle" on page 66.
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 Oracle8/9 Backup and Restore

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

Details about Oracle8/9 backup and restore sessions are also written in the following logs on the Oracle8/9 Server system:

- Data Protector writes the logs in the `/var/opt/omni/log/oracle8.log` (HP-UX and Solaris systems) or `/usr/omni/log/oracle8.log` (other UNIX systems) file
- Oracle8/9 usually writes the logs in the `<Oracle8 user dump directory>/sbtio.log` file.

Monitoring Tools

The progress of backups and restores can be monitored by querying the Oracle8/9 Target Database using the following SQL statement:

```
select * from v$SESSION_LONGOPS where  
compnam='dbms_backup_restore';
```

You can also monitor the progress of backups and restores from the Data Protector GUI by selecting the `Monitor Context` and then selecting the session which interests you from the list in the `Results Area`.

Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration

After uninstalling the Data Protector Oracle8/9 integration on an Oracle8/9 server system, the Oracle8/9 server software is still linked to the Data Protector Database Library. You must rebuild (Oracle8.x) or relink (Oracle8i or Oracle9i) the Oracle8/9 binary to remove this link. If this is not done, the Oracle8/9 server cannot be started after the integration has been removed.

After you have uninstalled the Data Protector Oracle 8 integration on the Oracle8/9 server system, proceed as described in the sections “Removing the Data Protector Oracle8/9 Integration Link on HP-UX Systems” on page 84 or “Removing the Data Protector Oracle8/9 Integration Link on Solaris and other UNIX Systems” on page 85.

Removing the Data Protector Oracle8/9 Integration Link on HP-UX Systems

1. On the Oracle8/9 Server system, connect to the Oracle8/9 database as an Oracle8/9 operating system user and shut down all Oracle8/9 instances.

Oracle8.x

2. If you have Oracle8.x installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```
 - b. Execute the following command:

```
make -f ins_rdbms.mk ioracle
```

IMPORTANT

The `make -f ins_rdbms.mk ioracle` command will work only if the `env_rdbms.mk` file was not changed.

**Oracle8i and
Oracle9i**

3. If you have Oracle8i or Oracle9i installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or  
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```
 - b. Execute the following command:

```
mv libobk.sl.orig libobk.sl
```

where `libobk.sl.orig` is the Oracle soft link as it existed before configuring the integration.
4. Start all Oracle8/9 instances.

Removing the Data Protector Oracle8/9 Integration Link on Solaris and other UNIX Systems

1. On the Oracle8/9 Server system, connect to the Oracle8/9 database as an Oracle8/9 operating system user and shut down all Oracle8/9 instances.

Oracle8.x

2. If you have Oracle8.x installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/rdbms/lib` directory:

```
cd <ORACLE_HOME>/rdbms/lib
```
 - b. Execute the following command:

```
make -f ins_rdbms.mk ioracle
```

IMPORTANT

The `make -f ins_rdbms.mk ioracle` command will work only if the `env_rdbms.mk` file was not changed.

**Oracle8i and
Oracle9i**

3. If you have Oracle8i or Oracle9i installed, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle8i) or
```

```
cd <ORACLE_HOME>/lib (64-bit Oracle9i).
```

- b. Execute the following command:

```
mv libobk.so.orig libobk.so
```

where `libobk.so.orig` is the Oracle soft link as it existed before configuring the integration.

4. Start all Oracle8/9 instances.

Oracle8i and Oracle9i RMAN Metadata and Data Protector Media Management Database Synchronization

This section describes how to synchronize Oracle8i and Oracle9i Recovery Manager metadata with the Data Protector Media Management Database.

Recovery Manager metadata contains information about the target database. RMAN uses this information for all backup, restore and maintenance operations. The metadata can be stored either in the recovery catalog database or in the control files.

Data Protector is the media manager that Oracle needs to perform tape storage backups and restores.

Data Protector has its own data protection policy that is not automatically synchronized with Oracle RMAN metadata. In order to have both catalogs synchronized, run the following command using RMAN:

```
allocate channel for maintenance type 'sbt_tape';  
crosscheck backup;  
release channel;
```

RMAN checks every backup piece in the repository and queries the MMDB for the availability of that backup piece. RMAN then mark the backup piece as expired or available, depending on media availability. Note that in the above example, RMAN does not delete backup pieces that are reported as expired by the MMDB, but instead marks them as expired.

Refer to the *Oracle8i Recovery Manager User's Guide and References* for more details on recovery catalog maintenance.

TIP

It is recommended that synchronization be performed in the following cases:

- after a Data Protector import or export of media with Oracle objects and
 - whenever protection for media with Oracle objects has expired.
-

Configuring the Integration as Cluster-Aware

Installation and Configuration

The Data Protector Oracle8/9 integration can be configured in the MC/ServiceGuard cluster. This means that either the Data Protector Cell Manager can be configured in a cluster, or the Data Protector client can be configured in the cluster. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for more information on supported configurations.

Refer to “Configuring the Integration” on page 20 and to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on how install and configure the Data Protector Oracle8/9 MC Service Guard integration.

The following steps must be performed when installing Oracle8/9 in an MC/ServiceGuard cluster:

1. Data Protector must be installed on all physical nodes. All physical nodes must also be imported in the corresponding cell.
2. All packages must be imported as virtual hosts. The packages must be running at the time of the import.

When configuring the Data Protector Oracle8/9 integration, configure it only on one of the cluster nodes for each Oracle8/9 server, since the Data Protector Oracle8/9 configuration files reside on the Cell Manager. Use the virtual hostname when configuring the integration. However, when linking Oracle8/9 with the Data Protector database library, link it on all nodes.

When configuring a backup specification for a Data Protector Oracle8/9 MC Service Guard integration, the virtual hostname should be used as the client hostname.

NOTE

The environment variable `OB2BARHOSTNAME` must be set to the virtual hostname before running the configuration or backup from the command line (on the client). When the GUI is used, this is not required. In the case of ZDB backup, you must use the command line.

For example, if the configuration is run on the cluster physical node `physical_1.domain.com` and Oracle8 is running on the virtual host `virtual.domain.com`, the variable setting is:

```
export OB2BARHOSTNAME=virtual.domain.com
```

Add the Oracle8/9 group `dba` user to Data Protector for the virtual server and for every node in the cluster. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on how to add a user to Data Protector.

For information on the Data Protector Cell Manager package configuration (if you want to install and configure a Data Protector Cell Manager in the MC/SG cluster), refer also to the *HP OpenView Storage Data Protector Administrator's Guide*.

Backup and Restore

When creating a Data Protector Oracle8/9 MC/SG cluster backup specification, always select the virtual hostname in the cluster rather than a particular node.

The following are extra requirements that must be fulfilled:

- Before performing an offline backup, be sure to take the Oracle8/9 Database resource offline and bring it back online after the backup.

This can be done using the `fscmd` command in the Pre-exec and Post-exec commands for the client system in a particular backup specification, or by using the Cluster Administrator.

- When restoring the database from a backup performed on a virtual host, you should set either the `OB2BARLIST` or `OB2BARHOSTNAME` environment variable in the RMAN script. For example,

```
run {
```

```
allocate channel dev1 type 'sbt_tape'  
parms 'ENV=(OB2BARHOSTNAME=virtual.domain.com)';  
  
restore datafile '/opt/ora9i/oradata/MAKI/example02.dbf';  
  
release channel dev1;  
  
}
```

Refer to “Backing Up an Oracle8/9 Database” on page 52 for information on how to create a Data Protector Oracle8/9 backup specification and to the *HP OpenView Storage Data Protector Administrator’s Guide* for information on MC/SG cluster backup specifics.

Refer to “Restoring Oracle8/9 Databases” on page 66 for information on how to restore an Oracle8/9 database.

Troubleshooting

Before you start troubleshooting the Data Protector Oracle8/9 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, problems and workarounds, and a list of related Data Protector patches.

The following sections provide some testing procedures you should perform before calling 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, actions are proposed to help you work around the problem.

Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration

After uninstalling the Data Protector Oracle8/9 integration on an Oracle8/9 server system, the Oracle8/9 server software is still linked to the Data Protector Database Library. You must rebuild (Oracle8.x) or relink (Oracle8i or Oracle9i) the Oracle8/9 binary to remove this link. If this is not done, the Oracle8/9 server cannot be started after the integration has been removed.

Please refer to “Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration” on page 84 for more information on how to make the Oracle8/9 server functional again.

Setting Up the Environment Variables

If you need to export some variables before starting the Oracle8 Server Manager or Oracle9i SQL Plus, TNS listener, or Recovery Manager, these variables must be exported as described in “Setting, Retrieving and Listing Data Protector Oracle8/9 Configuration Files’ Parameters Using the CLI” on page 15.

Checking Prerequisites Related to the Oracle8/9 Side of the Integration

For more detailed information about how to perform any of the following procedures, refer to the Oracle8/9 documentation.

1. Verify that you can access the Oracle8/9 Target Database and that it is opened as follows:

Export `<ORACLE_HOME>` and `<ORACLE_SID>` as follows:

- if you are using an sh - like shell, enter the following commands:

```
ORACLE_HOME="<ORACLE_HOME>"
export ORACLE_HOME
ORACLE_SID="<ORACLE_SID>"
export ORACLE_SID
```

- if you are using a csh - like shell, enter the following commands:

```
setenv ORACLE_HOME "<ORACLE_HOME>"
setenv ORACLE_SID "<ORACLE_SID>"
```

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl (Oracle8) or
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

```
connect internal
select * from dba_tablespaces;
exit
```

If this fails, open the Oracle8/9 Target Database.

2. Verify that you can access the recovery catalog (if used) as follows:

Export `<ORACLE_HOME>` and `<ORACLE_SID>` as described on page 93.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl (Oracle8) or
```

```
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

```
connect <Recovery_Catalog_Login>
```

```
select * from rcver;
```

```
exit
```

If this fails, open the recovery catalog.

3. Verify that the TNS listener is correctly configured for the Oracle8/9 Target Database and for the recovery catalog database. This is required for properly establishing network connections:

Export `<ORACLE_HOME>` as described on page 93.

Start the listener from the `<ORACLE_HOME>` directory:

```
bin/lsnrctl80 status <service>
```

```
exit
```

Note that with Oracle8i and Oracle9i, the `lsnrctl` and not the `lsnrctl80` command is used.

If it fails, start up the TNS listener process and refer to the Oracle documentation for instructions on how to create a TNS configuration file (`LISTENER.ORA`).

Export `<ORACLE_HOME>` as described on page 93.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl (Oracle8) or
```

```
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

```
connect <Target_Database_Login>
```

```
exit
```

and then

```
connect <Recovery_Catalog_Login>
```

```
exit
```

If this fails, refer to the Oracle8/9 documentation for instructions on how to create a TNS configuration file (TNSNAMES.ORA).

4. **Verify that the Oracle8/9 Target Database and the recovery catalog database are configured to allow remote connections with system privileges:**

Export <ORACLE_HOME> as described on page 93.

Start the Server Manager (Oracle8) or SQL Plus (Oracle9i) from the <ORACLE_HOME> directory:

```
bin/svrmgrl (Oracle8) or
```

```
bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt type:

```
connect <Target_Database_Login> as SYSDBA
```

```
exit
```

and

```
bin/svrmgrl connect <Recovery_Catalog_Login> as SYSDBA  
(Oracle8) or
```

```
bin/sqlplus connect <Recovery_Catalog_Login> as SYSDBA  
(Oracle9i)
```

```
exit
```

Repeat the procedure using SYSOPER instead of SYSDBA.

If this fails, refer to the Oracle8/9 documentation for instructions about how to set up the password file and any relevant parameters in the init<ORACLE_SID>.ora file.

5. **If you are using the recovery catalog database, verify that the Target Database is registered in the recovery catalog:**

Export `<ORACLE_HOME>` as described on page 93 and start the Oracle8 Server Manager or Oracle9i SQL Plus:

```
bin/svrmgrl (Oracle8) or bin/sqlplus /nolog (Oracle9i).
```

At the SVRMGR (Oracle8) or SQL (Oracle9i) prompt, type:

```
connect <Recovery_Catalog_Login>;
select * from db;
exit
```

If this fails, start the configuration using Data Protector or refer to the Oracle8/9 documentation for details about how to register an Oracle8/9 Target Database in the recovery catalog database.

6. **Verify backup and restore directly to disk using a Recovery Manager channel type disk.**

If you are using the recovery catalog:

Export `<ORACLE_HOME>` as described on page 93 and start Recovery Manager:

```
bin/rman target <Target_Database_Login> rcvcat
<Recovery_Catalog_Login> cmd_file=rman_script
```

If you are not using the recovery catalog:

Export `<ORACLE_HOME>` as described on page 93 and start Recovery Manager: `bin/rman target <Target_Database_Login> nocatalog cmd_file=rman_script`

An example of the RMAN script is presented below:

```
run {allocate channel 'dev0' type disk;
backup tablespace <tablespace_name>
format '<ORACLE_HOME>/tmp/<datafile_name>';}
```

After a successful backup, try to restore the backed up tablespace by running the following restore script:

```
run {
allocate channel 'dev0' type disk;
sql 'alter tablespace <tablespace_name> offline immediate';
```



```
restore tablespace <tablespace_name>;  
recover tablespace <tablespace_name>;  
sql 'alter tablespace <tablespace_name> online'  
release channel 'dev0';  
}
```

If this fails, refer to the Oracle8/9 documentation for details about how to execute a backup and restore directly to disk using the Recovery Manager.

Configuration Problems

IMPORTANT

If you have encountered any errors up to this point when performing the procedures described in the previous section, please contact Oracle8/9 support. The respective tests must be done before you even start checking the Data Protector Oracle8/9 configuration.

- 1. Verify that the Data Protector software has been installed properly.**

See “Installing and Upgrading the Oracle8/9 Integration” on page 18 for details.

- 2. Verify that the Data Protector Database Library is linked with the Oracle8/9 executable:**

Use the following command to check if the `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) file is linked with the Oracle8 executable.

Export `<ORACLE_HOME>` and `<ORACLE_SID>` as described on page 93.

On HP-UX platforms:

```
/usr/bin/chatr <ORACLE_HOME>/bin/oracle (32-bit Oracle8/9)
```

```
/usr/ccs/bin/ldd <ORACLE_HOME>/bin/oracle (64-bit Oracle8/9)
```

On Solaris systems:

```
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle
```

On other UNIX systems:

Integrating Oracle8/9 and Data Protector

Troubleshooting

```
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle
```

On IBM AIX systems:

```
/usr/bin/dump -H <ORACLE_HOME>/bin/oracle (32-bit Oracle8/9)
```

```
/usr/bin/dump -H -X64 <ORACLE_HOME>/bin/oracle (64-bit Oracle8/9)
```

On Linux systems:

```
/usr/bin/ldd <ORACLE_HOME>/bin/oracle
```

The output must state that the respective Data Protector library is required by the Oracle8/9 executable.

The following is an extract from the command output on HP-UX:

```
bin/oracle:
    shared executable
    shared library dynamic path search:
        SHLIB_PATH  enabled second
        embedded path disabled first Not Defined
    shared library list:
        static
/opt/omni/lib/libob2oracle8.sl(libob2oracle8_64bit.sl)
    dynamic /usr/lib/librt.2
    dynamic /usr/lib/libnss_dns.1
    dynamic /usr/lib/libdld.2
```

The line starting with SHLIB_PATH should be as presented in the example above. If this line is different, then enable the Data Protector Database Library dynamic path as follows:

```
/usr/bin/chatr +s enable <ORACLE_HOME>/bin/oracle
```

On Solaris, HP-UX (64-bit) and other UNIX platforms, LD_LIBRARY_PATH is used instead of SHLIB_PATH as on HP-UX (32-bit).

The following is an extract from the command output on other UNIX systems:

Figure 1-17

Output of the ldd command on other UNIX systems:

```
find library=/usr/omni/lib/libob2oracle8.so; required by /app/oracle8/product/8.0.4/bin/oracle  
/usr/omni/lib/libob2oracle8.so
```

3. Perform a filesystem backup of the Oracle8/9 Server system:

Perform a filesystem backup of the Oracle8/9 Server system so that you can eliminate any potential communication problems between the Oracle8/9 Server and the Data Protector Cell Manager system.

Do not start troubleshooting an online database backup unless you have successfully completed a filesystem backup of the Oracle8/9 Server system.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for details about how to do a filesystem backup.

4. Verify the permissions of the current user account:

Your user account should enable you to perform an Oracle8/9 backup or restore with Data Protector. Use the testbar2 utility to check the permissions:

```
/opt/omni/bin/testbar2 -perform:checkuser (HP-UX and  
Solaris systems) or
```

```
/usr/omni/bin/testbar2 -perform:checkuser (other UNIX  
systems).
```

If the user account holds all required permissions, you will receive only NORMAL messages displayed on the screen. See also “Configuring an Oracle8/9 User in Data Protector” on page 27.

5. Examine the system errors

The system errors are reported in the

```
/var/opt/omni/log/debug.log (HP-UX and Solaris systems) or  
/usr/omni/log/debug.log (other UNIX systems) file on the  
Oracle8/9 Server system.
```

Backup Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Check your Oracle8/9 Server configuration:

To check the Oracle8/9 Server configuration, login as an Oracle8/9 group dba user to the Oracle8/9 server system and start the following command on the Oracle8/9 Server system:

```
/opt/omni/lbin/util_oracle8.exe -CHKCONF <ORACLE_SID>  
(HP-UX and Solaris systems) or
```

```
/usr/omni/bin/util_oracle8.exe -CHKCONF <ORACLE_SID> (other  
UNIX systems).
```

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 <set_number> <Error_number>  
(HP-UX and Solaris systems) or
```

```
/usr/omni/bin/omnigetmsg <set_number> <Error_number>  
(other UNIX systems).
```

The *RETVAL*0 indicates successful configuration.

IMPORTANT

It is possible to receive a *RETVAL*0 even though the backup still fails. This can happen due to the fact that the backup owner is not the Oracle dba user.

2. Verify Data Protector internal data transfer using the testbar2 utility.

Before you run the testbar2 utility, verify that the Cell Manager name is correctly defined on the Oracle8/9 Server system. Check the /etc/opt/omni/cell/cell_server (HP-UX and Solaris systems) or /usr/omni/config/cell/cell_server (other UNIX systems) file, which contains the name of the Cell Manager system. Then run the following command:

On HP-UX and Solaris systems:

```
/opt/omni/bin/testbar2 -type:Oracle8 -appname:<ORACLE_SID>  
-bar:<backup_specification_name> -perform:backup
```

On other UNIX systems:

```
/usr/omni/bin/testbar2 -type:Oracle8 -appname:<ORACLE_SID>  
-bar:<backup_specification_name> -perform:backup
```

Switch to the Data Protector Manager and examine the errors reported by the testbar2 utility by clicking the Details button in the Data Protector Monitor context.

If the messages indicate problems on the Data Protector side of the integration, proceed as follows:

Create an Oracle8/9 backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

Data Protector reports "Export of the Recovery Catalog Database Failed" when backing up Oracle 9i

Problem

The following errors are listed in the Data Protector monitor:

```
EXP-00008: ORACLE error 6550 encountered
```

```
ORA-06550: line 1, column 13:
```

```
PLS-00201: identifier 'SYS.LT_EXPORT_PKG' must be declared
```

```
ORA-06550: line 1, column 7:
```

```
PL/SQL: Statement ignored
```

```
EXP-00083: The previous problem occurred when calling  
SYS.LT_EXPORT_PKG.schema_info_exp
```

```
. exporting statistics
```

```
Export terminated successfully with warnings.
```

```
[Major] From: ob2rman.exe@machine "MAKI" Time: 10/01/01 16:07:53
```

```
Export of the Recovery Catalog Database failed.
```

Action Log in to Oracle9i SQL Plus and grant the execute permission to the Oracle9i LT_EXPORT_PKG as follows (make sure that the user sys has the SYSDBA permission granted beforehand):

```
sqlplus `sys/<password>@CDB as sysdba`
```

```
SQL> grant execute on sys.lt_export_pkg to public;
```

Restart the failed backup session.

Data Protector reports “Cannot allocate/attach shared memory”

Problem Backup fails and the following error message is displayed:

```
Cannot allocate/attach shared memory (IPC Cannot Allocate Shared  
Memory Segment)
```

```
System error: [13] Permission denied) => aborting
```

Action Set the OB2SHMEM_IPCGLOBAL omnirc option in the /opt/omni/.omnirc file to 1 in order to use the memory windowing properly, and restart the failed backup session. Refer to the *HP OpenView Storage Data Protector Administrator’s Guide* for details on using omnirc options.

Backup Fails After a Point in Time Restore and Recovery

Problem Backup fails after a Point in time restore and recovery was performed and the following error is displayed:

```
RMAN-06004: ORACLE error from recovery catalog database:  
RMAN-20003: target database incarnation not found in  
recovery catalog
```

Action Connect to the target and recovery catalog database using RMAN and reset the database:

```
rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login>RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Restore Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Verify that an object exists on the backup media:

This can be done by executing:

```
/opt/omni/bin/omnidb -oracle8 "<object_name>" -session  
"<Session_ID>" -media (HP-UX and Solaris systems) or  
/usr/omni/bin/omnidb -oracle8 "<object_name>" -session  
"<Session_ID>" -media (other UNIX systems)
```

on the Oracle8/9 Server system.

The output of the command lists detailed information about the specified Oracle8/9 object, as well as the session IDs of the backup sessions containing this object and a list of the media used. For detailed syntax of the `omnidb` command, run:

```
man omnidb
```

2. Simulate a restore session

Once you know the information about the object to be restored, you can simulate a restore using the Data Protector `testbar2` utility.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the Oracle8/9 Server system. Check the `/etc/opt/omni/cell/cell_server` (HP-UX and Solaris systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system.

Test Data Protector internal data transfer using the `testbar2` utility:

On HP-UX and Solaris systems:

```
/opt/omni/bin/testbar2  
-type:Oracle8  
-appname:<ORACLE_SID>  
-perform:restore  
-object:<object_name>  
-version:<object_version>  
-bar:<backup_specification_name>
```

On other UNIX systems:

```
/usr/omni/bin/testbar2  
-type:Oracle8  
-appname:<ORACLE_SID>  
-perform:restore  
-object:<object_name>  
-version:<object_version>  
-bar:<backup_specification_name>
```

IMPORTANT

The hostname should not be specified in the `object` option. It is automatically provided by `testbar2`.

You should see only `NORMAL` messages displayed on your screen, otherwise examine the errors reported by the `testbar2` utility by clicking the `Details` button in the `Data Protector Monitor` context.

If the messages indicate problems on the `Data Protector` side of the integration, proceed as follows:

Run the `omnidb` command to view the objects in the database.

3. Ensure that the database is in the correct state.

If you are trying to restore a database item using the `Data Protector Restore GUI for Oracle` and the GUI hangs try one of the following:

- If you are restoring the control file the database should be in the `NoMount` state.

Open a command window and enter the following:

```
sqlplus/nolog  
SQL>connect/as sysdba  
SQL>shutdown immediate  
SQL>startup nomount
```


- If you are restoring datafiles the database should be in the Mount state.

Open a command window and enter the following:

```
sqlplus/nolog
SQL>connect/as sysdba
SQL>shutdown immediate
SQL>startup mount
```

4. Check your environment variables.

The message below sometimes appears when you are restoring database items to a new host:

```
"Binary util_orarest is missing. Cannot get information
from the remote host."
```

To resolve this problem do as follows:

- a. Close Data Protector.
- b. Set the environment variable on the system where the Cell Manager resides:

```
OB2_ORARESTHOSTNAME = <target Oracle host>
```
- c. Restart Data Protector and try to restore the database items again.
- d. When the restore is complete, close Data Protector and re-set the following environment variable:

```
OB2_ORARESTHOSTNAME = <empty>
```
- e. Restart Data Protector.

5. Try using the RMAN CLI to restore the database items.

If there is a problem you cannot resolve while you are trying to restore a database item using the Data Protector Restore GUI for Oracle try using the RMAN CLI to restore the database items.

For information about using the CLI see "Restoring Oracle8/9 Using RMAN" on page 80.

6. Try putting the database into the Open state manually after using the Data Protector Restore GUI for Oracle to recover and restore a Zero Downtime Backup or Standalone Backup session.

If you have used the Data Protector Restore GUI for Oracle to recover and restore a Zero Downtime Backup (ZDB) session or a standalone backup, and you see the following error message:

```
Oracle Error: ORA-1589: must use RESETLOGS or NORESETLOGS
option for database open.
```

Open a SQLplus window and use the following command:

```
SQL>sqlplus/nolog
SQL>connect as sysdba
SQL>alter database open no resetlogs/resetlogs;
```

If this does not work try using the following command:

```
SQL>alter database open resetlogs;
```

“Binary util_orarest failed” error message is displayed when browsing Oracle9 database for restore on Linux

Problem

The following error message is displayed when browsing *Oracle9* database for restore on Linux:

```
Binary util_orarest failed. Cannot get information from the
remote host.
```

Action

Replace the `util_orarest.exe` utility with the new `util_orarest9.exe` (both located in the `/usr/omni/bin` directory on Linux):

1. Rename the `util_orarest.exe` to `util_orarest.exe.orig`
2. Rename the `util_orarest9.exe` to `util_orarest.exe`

Examples of an Oracle8/9 Database Restore

This section describes some examples of how you can restore an Oracle8/9 database. The following examples of restore are given:

- Full database restore
- Point-in-time restore
- Tablespace restore
- Datafile restore
- Archive log restore
- Control files restore

IMPORTANT

The restore of an Oracle8/9 database is performed using the RMAN utility, which is not a part of Data Protector. This section only describes *examples* of how you can perform a restore. The examples provided do not apply to all situations where a restore is needed. For additional information on how you can restore an Oracle8/9 database using the RMAN utility, refer to the Oracle documentation.

Preparing the Oracle8/9 Database for Restore

The restore of an Oracle8/9 database can be performed when the database is in a mount (quiescent) mode. However, when you are performing the restore of tablespaces or datafiles, only a part of the Oracle8/9 database can be put offline.

Prerequisites

The following requirements must be met before you start a restore of an Oracle8/9 database:

- If you are using a recovery catalog database, make sure that the database is running. If the recovery catalog database cannot be brought online, you will probably need to restore the database. Refer to “Restoring Oracle8/9 Databases” on page 66 for details on how to restore the recovery catalog database.

- If you are using control files, they must be accessible in order to perform the restore. If the control files are not available, you will probably need to restore them. Refer to “Control File Restore” on page 115 for details on how to restore the control files.

If you have to perform a restore of the recovery catalog database or control files, you must perform this restore first. Only then can you perform a restore of other parts of the Oracle8/9 database.

When you are sure that the recovery catalog database or control files are in place, start the recovery catalog database and the listener.

- Make sure that the following environment variables are set:

- ✓ ORACLE_BASE
- ✓ ORACLE_HOME
- ✓ ORACLE_TERM
- ✓ ORACLE_SID
- ✓ PATH
- ✓ NLS_LANG
- ✓ NLS_DATE_FORMAT

Example of Environment Variables

```
ORACLE_BASE=/opt/oracle
ORACLE_HOME=/opt/oracle/product/8.1.6
ORACLE_TERM=hp
ORACLE_SID=PROD
PATH=$PATH:/opt/oracle/product/8.1.6/bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

- Check that the `/etc/oratab` file has the following line:

```
PROD:/opt/oracle/product/8.1.6:N
```

The last letter determines whether the database will automatically start upon bootup (Y) or not (N).

Examples of Oracle8/9 Database Restore

In the examples below, the following connection strings are used:

- Target connection string for target database:

```
sys/manager@PROD
```

where `sys` is the username, `manager` is the password and `PROD` is the name of the Oracle8/9 database.

- Recovery catalog connection string for recovery catalog database:

```
rman/rman@CATAL
```

where `rman` is the username and password and `CATAL` is the name of the Oracle8/9 database.

Full Database Restore

To perform a full database recovery, you also need to restore and apply all the archive logs. To perform a full database restore, follow the steps below:

1. Log in to the Oracle 8/9 RMAN:
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog
```

2. Start the full database restore:

```
run{  
  
allocate channel <dev1> type '<sbt_tape>' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)';  
  
restore database;  
  
recover database;  
  
sql 'alter database open';
```

```
release channel <dev1>;  
}
```

You can also save the script into a file and perform a full database restore using the saved files. The procedure in such cases is as follows:

1. Create a file `restore_database` in the `/var/opt/omni/tmp` directory.
2. Start the full database restore.
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL  
cmdfile=/var/opt/omni/tmp/restore_datafile
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog cmdfile=/var/opt/omni/tmp/restore_datafile
```

Point-in-Time Restore

To perform a point-in-time restore, you also need to restore and apply the archive logs to the specified point in time. To perform a point-in-time database restore, follow the steps below:

1. Log in to the Oracle 8/9 RMAN:
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog
```

2. Start the point-in-time restore:

```
run{  
allocate channel <dev1> type 'sbt_tape' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;  
set until time 'Mar 14 2001 11:40:00';
```

```
restore database;  
recover database;  
sql 'alter database open';  
release channel <dev1>;  
}
```

3. After you have performed a point-in-time restore, reset the database in the Recovery Catalog.

You can also save the script into a file and perform a point-in-time restore using the saved files. Follow the steps below:

1. Create a file `restore_PIT` in the `/var/opt/omni/tmp` directory.
2. Start the point-in-time restore.
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_PIT
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog cmdfile=/var/opt/omni/tmp/restore_PIT
```

Tablespace Restore

If a table is missing or corrupted, you need to perform a restore of the entire tablespace. To restore a tablespace, you may take only a part of the database offline, so that the database does not have to be in the mount mode. You can use either a recovery catalog database or control files to perform a tablespace restore. Follow the steps below:

1. Log in to the Oracle 8/9 RMAN:
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog
```

2. Start the tablespace restore.

- a. If the database is in the open state, the script to restore the tablespace should have the following format:

```
run{  
allocate channel <dev1> type '<sbt_tape>' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;  
sql 'alter tablespace "TEMP" offline immediate';  
restore tablespace 'TEMP';  
recover tablespace 'TEMP';  
sql 'alter tablespace "TEMP" online';  
release channel <dev1>;  
}
```

- b. If the database is in the mount state, the script to restore the tablespace should have the following format:

```
run{  
allocate channel <dev1> type '<sbt_tape>' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;  
restore tablespace 'TEMP';  
recover tablespace 'TEMP';  
release channel <dev1>;  
}
```

You can also save the script into a file and perform a tablespace restore using the saved files. The procedure in such cases is as follows:

1. Create a file `restore_TAB` in the `/var/opt/omni/tmp` directory.
2. Start the tablespace restore.

- a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL cmdfile=/var/opt/omni/log/restore_TAB
```


- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog cmdfile=/var/opt/omni/log/restore_TAB
```

Datafile Restore

To restore a datafile, you may take only a part of the database offline. To perform a datafile restore, follow the steps below:

1. Log in to the Oracle 8/9 RMAN:

- a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog
```

2. Start the datafile restore:

- a. If the database is in an open state, the script to restore the datafile should have the following format:

```
run{  
  allocate channel <dev1> type '<sbt_tape>' parms  
  'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;  
  sql 'alter database datafile  
  '/opt/oracle/data/oradata/DATA/temp01.dbf' offline;  
  restore datafile  
  '/opt/oracle/data/oradata/DATA/temp01.dbf' ;  
  recover datafile  
  '/opt/oracle/data/oradata/DATA/temp01.dbf' ;  
  sql 'alter database datafile  
  '/opt/oracle/data/oradata/DATA/temp01.dbf' online;  
  release channel <dev1>;  
}
```

- b. If the database is in a mount state, the script to restore the datafile should have the following format:

```
run{
  allocate channel <dev1> type '<sbt_tape>' parms
  'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)' ;
  restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' ;
  recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' ;
  release channel <dev1>;
}
```

You can also save the script into a file and perform a datafile restore using the saved files. The procedure in such cases is as follows:

1. Create a file `restore_dbf` in the `/var/opt/omni/tmp` directory.
2. Start the datafile restore.
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat
rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_dbf
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD
nocatalog cmdfile=/var/opt/omni/tmp/restore_dbf
```

Archive Log Restore

To restore an archive log, follow the steps below:

1. Login to the Oracle 8/9 RMAN:
 - a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat
rman/rman@CATAL
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog
```

2. Start the archive log restore:

```
run{  
  allocate channel <dev1> type '<sbt_tape>' parms  
  'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>)';  
  restore archivelog all;  
  release channel <dev1>;  
}
```

You can also save the script into a file and perform an archive log restore using the saved files. The procedure in such cases is as follows:

1. Create a file `restore_arch` in the `/var/opt/omni/tmp` directory.
2. Start the archive log restore.

- a. If you are using the recovery catalog database, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD rcvcat  
rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_arch
```

- b. If you are using control files, run the following command:

```
<ORACLE_HOME>/bin/rman target sys/manager@PROD  
nocatalog cmdfile=/var/opt/omni/tmp/restore_arch
```

Control File Restore

The restore and recovery procedure of Oracle 8/9 control files is a very delicate operation, which depends on whether you are using the recovery catalog or control file as a central repository and on the version of the Oracle database you are using. For detailed steps to perform the restore of control files, please refer to the *Recovery Manager User's Guide and Reference*.

Integrating Oracle8/9 and Data Protector
Examples of an Oracle8/9 Database Restore

In This Chapter

This chapter explains how to install, configure, and use the HP OpenView Storage Data Protector SAP R/3 integration. It explains the concepts and methods that you need to understand in order to back up and restore SAP R/3 databases.

It is organized into the following sections:

“Overview” on page 119

“Prerequisites and Limitations” on page 121

“Integration Concept” on page 123

“Data Protector SAP R/3 Configuration File” on page 132

“Installing and Upgrading the Data Protector SAP R/3 Integration” on page 138

“Configuring the Integration” on page 140

“Backing Up an SAP R/3 Database” on page 167

“Restoring an SAP R/3 Database” on page 174

“Monitoring an SAP R/3 Backup and Restore” on page 180

“Configuring the Integration as Cluster-Aware” on page 181

“Troubleshooting” on page 183

“Example of SAP R/3 Database Restore” on page 197

Overview

Data Protector integrates with the SAP R/3 Database Server to offer online backup of your SAP R/3 databases.

If the SAP R/3 system uses an Oracle database, then the Data Protector SAP R/3 integration can be used for backup. If any other database is used by SAP, then the corresponding Data Protector integration of that database (for example, Informix) must be used instead.

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 integration.

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

The SAP R/3 part of the integration provides storage management utilities. These utilities communicate with Data Protector via the Data Protector `backint` executable, which complies with the SAP R/3 backup interface.

Using Data Protector with the SAP R/3 Database Server offers several advantages over using SAP R/3 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 that allows you to keep track of all media and the status of each medium, set the protection for stored data, fully automate operations as well as organize and manage devices and media.

- 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 configure run unattended at the periods you specify.

Overview

- Local versus Network Backups

When configuring an SAP R/3 backup, the location of devices is completely transparent to the user. They can be connected to the SAP R/3 Database Server or any other Data Protector clients on the network.

- 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.

- 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

This section provides you with a list of prerequisites and limitations you must be aware of before using the integration.

Prerequisites

- The database used by SAP R/3 must be an Oracle database. If any other database is used by SAP, then the corresponding Data Protector integration of that database (for example, Informix) must be used instead.
- You need a license to use the Data Protector SAP R/3 integration. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for licensing instructions.
- Before you begin, ensure that you have correctly installed and configured the SAP R/3 Database 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.
 - ✓ *SAP R/3 System Online Documentation* for instructions on how to install and configure the SAP R/3 database and the SAP R/3 backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE).

Limitations

See the *HP OpenView Storage Data Protector Software Release Notes* for a list of general Data Protector limitations. The following is a list of integration-specific limitations:

- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands. These commands are optionally entered as integration-specific options during the creation of backup specifications.

Prerequisites and Limitations

- Do not configure RMAN Offline SAP R/3 backups using the Internal user since the backups will not work.

Integration Concept

This integration links SAP R/3 backup utilities (BRTOOLS) with Data Protector. SAP R/3 backup utilities provide an interface between an SAP R/3 Database Server and media management applications, like Data Protector. They enable the backup or restore of the following SAP R/3 data objects:

- data files
- control files
- online redo logs
- offline (archived) redo logs
- SAP R/3 logs and parameter files

Because SAP R/3 Database Servers run on top of Oracle databases, the SAP R/3 backup objects are very similar to those of Oracle. The main difference is that SAP R/3 backup utilities hide the database from Data Protector, which sees those objects as plain files.

Version 4.5 and higher of the SAP R/3 backup utilities allows Oracle8 data files to be backed up directly using the Oracle8 Recovery Manager (hereinafter RMAN), as well as using the Data Protector Oracle8 Integration. This is very useful because RMAN supports incremental backups, and thus the backup time and the amount of backed up data can be significantly reduced. This mode of SAP R/3 backup and restore will be referred to as RMAN mode in the rest of this chapter. The previous mode (where data files were backed up as plain files) will be referred to as the backint mode.

SAP R/3 Backup Utilities

SAP R/3 backup utilities are the following:

- BRBACKUP

This utility performs online and offline backup of control files, data files, and online redo log files. Additionally, BRBACKUP saves the profiles and logs relevant for a particular backup session.

- BRARCHIVE

This utility performs backups of the offline (archived) redo logs, written by Oracle to the archiving directory.

- BRRESTORE

This utility restores the backed up data using the BRBACKUP and BRARCHIVE utilities.

These backup utilities can be started directly using Data Protector, or interactively using SAPDBA, which is an SAP R/3 administration utility.

NOTE

Data Protector supports all SAP BRTOOLS options that are also supported by `backint`, except for the `BRARCHIVE -b` and `BRBACKUP -a` options.

**Data Protector
Integration
Software**

The Data Protector integration software consists of the following components, as depicted in Figure 2-1 on page 125.

- The `backint` program is a backup interface between the Data Protector software and the SAP R/3 backup and restore tools.
It is started using `BRBACKUP` or `BRARCHIVE` during a backup session, and `BRRESTORE` during a restore session.
- The `sapback` program performs the actual backup of files.
- The `saprest` program performs the actual restore of files.
- The Data Protector Database Library links Data Protector and Oracle8 Server software. This is required only if SAP R/3 is backed up in the RMAN mode.
- The `omnisap.exe` program is used by Data Protector to start the SAP R/3 backup tools.
- The `testbar2` utility checks the Data Protector part of the integration.
- The `util_sap.exe` program is used by Data Protector to configure the integration.
- The configuration file on the Cell Manager system contains data needed by Data Protector to run backups and restores.

Figure 2-1 SAP R/3 Backup Concept

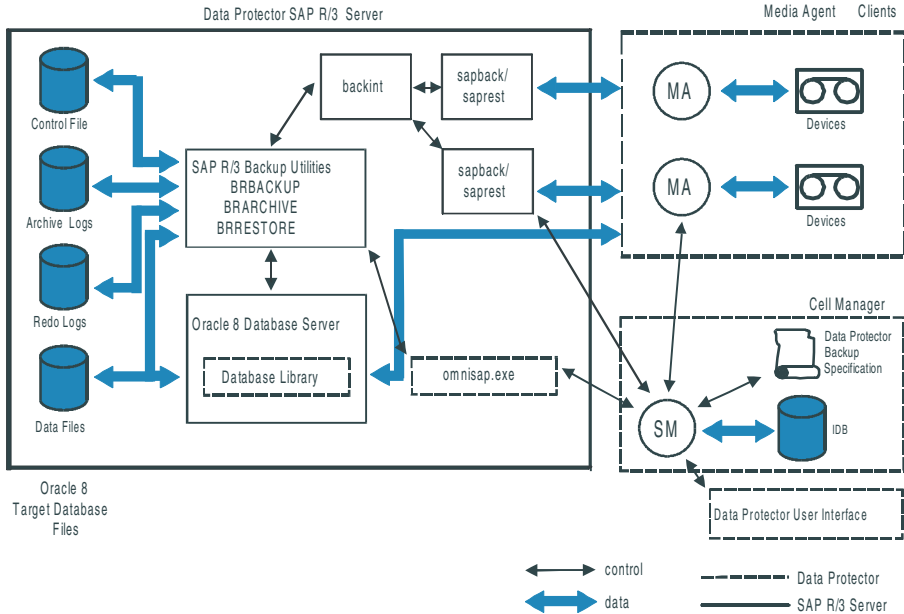


Table 2-1 Legend

SM	The Data Protector Session Manager, which is the Data Protector Backup Session Manager during backup or the Data Protector Restore Session Manager during restore.
Database Library	The interface between SAP R/3 Server processes and Data Protector
IDB	The IDB, which stores information about Data Protector sessions, such as session messages, and information about objects, data, used devices, and media.
MA	The Data Protector Media Agent

SAP R/3 Architecture

Depending on the backup mode, there are two possible backup scenarios (backint mode or RMAN mode) that can be used.

Backup Flow Using backint

The backup session undergoes the following stages if the backup is performed in backint mode. Refer to Figure 2-2 for details.

NOTE

It is not possible to perform an incremental backup in backint mode.

Figure 2-2 SAP R/3 Architecture: Backint Mode

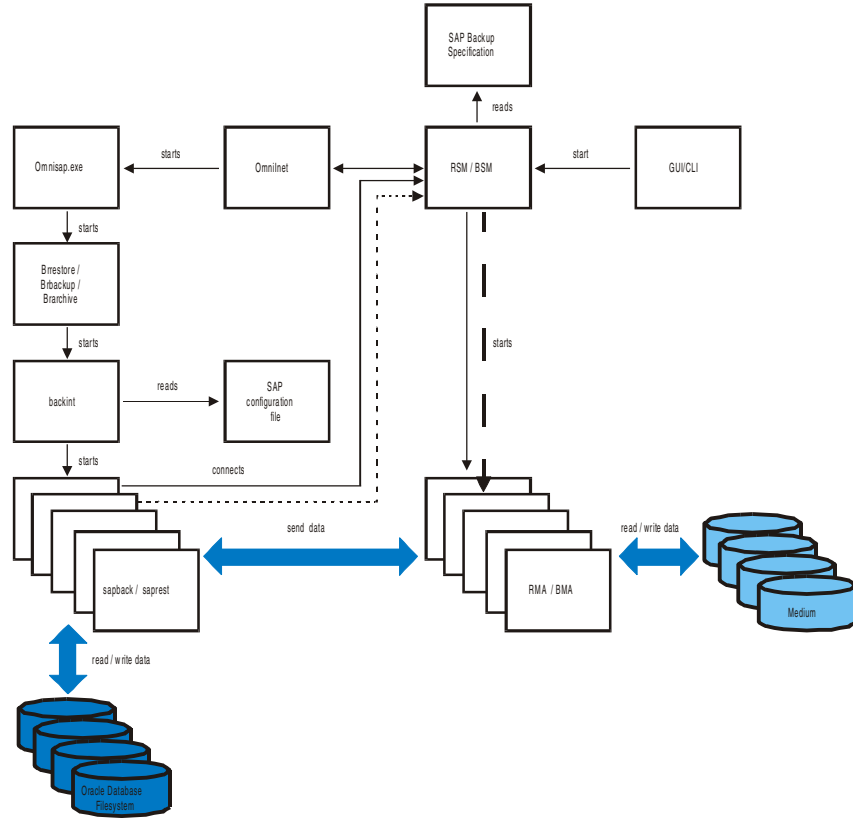


Table 2-2

Legend

BSM	Data Protector Backup Session Manager
RSM	Data Protector Restore Session Manager
BMA	Data Protector Backup Media Agent

Table 2-2

Legend

RMA	Data Protector Restore Media Agent
GUI/CLI	Data Protector User Interface

1. The backup session can be started using the Data Protector GUI, or interactively using the SAP R/3 utilities.

If the backup session is started using the Data Protector User Interface (or using the scheduler), then the Backup Session Manager (BSM) is started. The BSM then reads the appropriate Data Protector backup specification, checks if the devices are available, and starts the `omnisap.exe` program on the SAP R/3 Database Server.

The `omnisap.exe` program exports the appropriate environment variables and starts either the BRBACKUP or BRARCHIVE utilities. These utilities then initiate the first `backint` command to back up the Oracle Target Database's data files and the control files (BRBACKUP) or to back up archived redo log files (BRARCHIVE).

If the backup is started interactively using the SAPDBA program, then the BRBACKUP or BRARCHIVE utilities are started directly.

2. BRBACKUP does the following:

- Automatically changes the state of the Oracle Target Database (opened or closed), according to the backup type (online or offline).
- Switches the Oracle Target Database to the ARCHIVELOG mode before the backup.

The archived redo log files are written to the archiving directory by Oracle and are backed up later using BRARCHIVE.

- Writes the BRBACKUP log during the backup session, with information about the backup file and the backup ID. These logs must be available in order to determine the location of the database files and archived redo log files during restores.
- Sets the tablespace mode (BEGIN / END BACKUP) in the case of online backup using `backint`.

In this way, the SAP R/3 locks the tablespace just before it is backed up, and opens the tablespace immediately after the backup is completed. The tablespaces are therefore locked for a minimal amount of time.

3. The `backint` program obtains the SAP R/3 configuration from the Cell Manager, divides the files for backup into subsets (provided that the specified concurrency is greater than 1) and starts the `sapback` program for each subset. Each `sapback` process connects to the BSM, which then starts Media Agents on the corresponding client systems and establishes a connection between the `sapback` processes and Media Agents.

Data transfer can begin at this point. The `sapback` processes read data from disks and send it to Media Agents. The first `backint` program stops as soon as all `sapback` processes have finished and control is returned to the parent process, either the `BRBACKUP` or `BRARCHIVE` utility.

The second `backint` command is initiated by either the `BRBACKUP` or `BRARCHIVE` command. This command attempts to back up the SAP R/3 log files and parameter files (in the case of `BRBACKUP`), or the archived redo logs (in the case of `BRARCHIVE`) that have been created since the first `backint` command.

If new archived redo logs have been created, they are backed up and another `backint` command is started. Otherwise, the SAP R/3 log files and the parameter files are backed up, and the second `backint` program is started using `BRBACKUP`.

Therefore, more than two `backint` commands may be initiated by `BRARCHIVE`, while there are only two `backint` commands initiated by `BRBACKUP`.

NOTE

The total number of `sapback` processes started in one session using Data Protector is limited to 256.

4. Media Agents finish transferring data when all the `sapback` processes are complete. When all of the Media Agents have finished data transfer, the BSM waits for a timeout (`SmWaitForNewClient` global variable) and completes the backup session, as long as no `backint` is started within this time frame.

**Backup Flow
Using Recovery
Manager**

A backup session using `RMAN` mode differs from a backup session in `backint` mode in step 3. Refer to Figure 2-3 on page 130 for details.

BRBACKUP starts RMAN, which then connects to the Data Protector Database Library via the Oracle8 Server processes. The Database Library provides a connection to the Data Protector BSM, which starts Media Agents and establishes a connection between the Oracle8 Server and Media Agents.

The data transfer begins at this point. The Oracle8 Server sends data to Media Agents, which then write the data to the media.

Once the Oracle8 Target Database's data files have been written to the media, the respective Oracle8 Server processes are completed, and so, subsequently, is RMAN. The backup control is now returned to BRBACKUP, which starts the first backint command to back up the Oracle8 Target Database's control file and the SAP R/3 log files. Archive logs are backed up in the same manner as in backint mode.

Figure 2-3 SAP R/3 Architecture: RMAN Mode

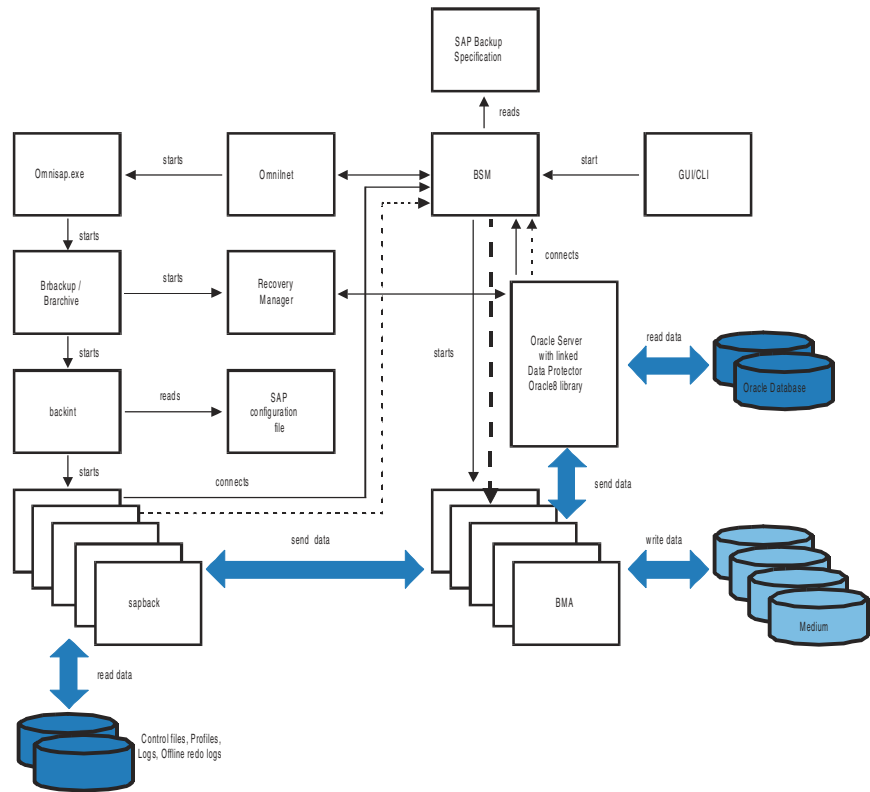


Table 2-3 Legend

BSM	Data Protector Backup Session Manager
BMA	Data Protector Backup Media Agent
GUI/CLI	Data Protector User Interface

Restore Flow Using Backint

SAP R/3 restore can be initiated using Data Protector, or interactively using the SAP R/3 utilities. However, only a standard filesystem restore is performed using Data Protector.

The restore session proceeds according to the following stages if the restore is performed in backint mode.

1. Using the SAPDBA utility, the objects to be restored are selected.
2. The BRRESTORE first checks whether the required free disk space is available to allow the files to be restored. It then starts the first backint command to restore the Oracle Target Database's data files. The backint command reads the SAP R/3 configuration file, divides the files for restore into subsets (provided that the specified concurrency is greater than 1) and starts the saprest process for each subset.

The first saprest process starts the Data Protector Restore Session Manager (RSM), while the subsequent saprest processes connect to the same RSM. In addition, the saprest process checks whether the specified objects have been backed up.

The RSM checks the availability of the restore devices, starts Media Agents and establishes a connection between the saprest processes and Media Agents. Data transfer begins at this stage. Data is sent from the media to the target disks. The Media Agent finishes as soon as all saprest processes connected to it are completed.

3. When all the Media Agents have finished, the RSM waits for a timeout (`SmWaitForNewClient` global variable) and completes the restore session, if no backint is started within this time frame.

Restore Flow Using Recovery Manager

A restore session using RMAN differs from a restore session using the backint mode in the step 2 as follows:

BRRESTORE starts RMAN in order to restore the Oracle8 Target Database data files. RMAN then connects to the Data Protector Database Library via the Oracle8 Server processes.

Data Protector SAP R/3 Configuration File

Data Protector stores the SAP R/3 integration parameters for every configured SAP R/3 instance in the

`/etc/opt/omni/integ/config/SAP/<client_name>%<ORACLE_SID>`
file (HP-UX and Solaris systems), or in the

`<Data_Protector_home>\Config\integ\config\sap\<client_name>%<ORACLE_SID>` file (Windows systems) on the Cell Manager. The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- concurrency number and balancing (for each backup specification), and number of channels for RMAN backup
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

IMPORTANT

To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

NOTE

You can set up the parameters in the Environment section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='<ORACLE_HOME>';
ConnStr='<ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE>';
BR_directory='<BRTOOLS_HOME>;
SAPDATA_HOME='<SAPDATA_HOME>';
Environment={
  [<ENV var1>='<value1>'];
  [<ENV var2>='<value2>';
  ...]
}
SAP_Parameters={<bckup_spec_name>=('-concurrency <# of
concurrency>' | '-time_balance' | '-load_balance' |
'-manual_balance' | '-channels <#_of_RMAN_channels>');
}
speed={
  AVERAGE=1;
  '<filename>'=<# of seconds needed to backup this file>;
}
compression={'<filename>'=<size of the file in bytes after the
compression>;
}
manual_balance={<backup_specification_name>={'<filename>'=<device_
number>;
}
}
```

Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIIBBFIBBGHBBQDBBOFBBCFBFBPFBBFCFBBIFBBGFBBBDGBBB
FBBCFBBDFFBBCFBB';

BR_directory='/usr/sap/ABA/SYS/exe/run'; SAPDATA_HOME='/sap';

Environment={ }

SAP_Parameters={
    sap_weekly_offline=('-concurrency 1','-no_balance');
    sap_daily_online=('-concurrency 3','-load_balance');
    sap_daily_manual=('-concurrency 3','-manual_balance');
}

speed={
    AVERAGE=1;
'/file1'=2345;
'/file2'=6789;
}

compression={
'/file1'=1234;
'/file2'=5678;
}

manual_balance={
    sap_daily_manual={
'/file1'=1; /* file 1 is backed up by the first sapback */
'/file2'=2; /* file 2 is backed up by the second sapback */
'/file3'=1; /* file 3 is backed up by the first sapback */
'/file4'=1;
    }
}
```

Setting, Retrieving, Listing and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI

The Data Protector SAP R/3 configuration file parameters are normally written to the Data Protector SAP R/3 configuration file after:

- the configuration of the SAP R/3 instance in Data Protector is completed.
- a new backup specification is created.
- a backup that uses balancing by time algorithm is completed.

NOTE

The variable definitions that are command or shell-based must be hard-coded in the Data Protector SAP R/3 configuration file, since such variables definitions are not possible in the Data Protector SAP R/3 configuration file.

The `util_cmd` Command

You can set, retrieve, list or delete the Data Protector SAP R/3 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 SAP R/3 client. The command resides in the `/opt/omni/lbin` (HP-UX and Solaris systems) or in the `/usr/omni/bin` (other UNIX systems) directory.

The `util_cmd` Synopsis

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

```
util_cmd -getconf[ig] SAP <sap_instance> [-local <filename>]
util_cmd -getopt[ion] [SAP <sap_instance>] <option_name>
[-sub[list] <sublist_name>] [-local <filename>]
util_cmd -putopt[ion] [SAP <sap_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 to which a parameter is written to or taken from.

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

- When it is used with the `-getconf [ig]` option, the filename for the output of the command to be written to (if the `-local` option is not specified, the output is written to the standard output).
- When it is used with the `-getopt [ion]`, the filename of the file from which the parameter and its value are to be taken and then written to the standard output (if the `-local` option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output).
- When it is used with the `-putopt [ion]` option, the 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 SAP R/3 configuration file).

NOTE

If you are setting the `option_value` parameter as a number, the number must be put in single quotes, surrounded by double quotes.

Return Values

The `util_cmd` command displays a short status message after each operation (writes it 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 appears 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, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, etc.

Setting Parameters

To set the Data Protector OB2OPTS and the Oracle8 NLS_LANG parameters for the SAP R/3 instance ICE, use the following commands on the Data Protector SAP R/3 client:

```
/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment  
  
/opt/omni/lbin/util_cmd -putopt SAP ICE NLS_LANG \  
'AMERICAN_AMERICA.US7ASCII' -sublist Environment  
  
/opt/omni/lbin/util_cmd -putopt SAP TOR BR_TRACE "'10'" \  
-sublist Environment (HP-UX and Solaris systems)  
  
/usr/omni/bin/util_cmd -putopt SAP ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment  
  
/usr/omni/bin/util_cmd -putopt SAP ICE NLS_LANG \  
'AMERICAN_AMERICA.US7ASCII' -sublist Environment  
  
/usr/omni/bin/util_cmd -putopt SAP TOR BR_TRACE "'10'" \  
-sublist Environment (other UNIX systems)
```

Retrieving Parameters

To retrieve the value of the OB2OPTS parameter for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

```
/opt/omni/lbin/util_cmd -getopt SAP ICE OB2OPTS -sublist \  
Environment (HP-UX and Solaris systems)  
  
/usr/omni/bin/util_cmd -getopt SAP ICE OB2OPTS -sublist \  
Environment (other UNIX systems)
```

Listing Parameters

To list all the Data Protector SAP R/3 configuration file parameters for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

```
/opt/omni/lbin/util_cmd -getconf SAP ICE (HP-UX and Solaris  
systems)  
  
/usr/omni/bin/util_cmd -getconf SAP ICE (other UNIX systems)
```

Deleting Parameters

To remove the value of the OB2OPTS parameter for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

```
/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS -sublist  
Environment (HP-UX and Solaris systems)  
  
/usr/omni/bin/util_cmd -putopt SAP ICE OB2OPTS -sublist  
Environment (other UNIX systems)
```

Installing and Upgrading the Data Protector SAP R/3 Integration

- Prerequisites** You have to shut down any running Oracle8 or SAP R/3 databases on the SAP R/3 servers to be integrated before upgrading or installing the Data Protector SAP R/3 integration, and start them again after the upgrade or installation.
- Upgrading** The SAP R/3 integration is upgraded automatically during the client upgrade. After the upgrade procedure has completed, some additional steps have to be performed manually to finish the upgrade. For the upgrade procedure, refer to “Upgrading to Data Protector A.05.10” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- Installation** Install the Data Protector software on your SAP R/3 Database Server following the instructions given in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- Which Components Should I Install?** Install the following components:
- SAP R/3 Integration
 - Oracle8 Integration
- Install this component if you intend to use the Oracle8 Recovery Manager to back up the SAP R/3 database files.
- User Interface
- Install this component to gain access to the Data Protector GUI and the Data Protector CLI on the system.
- Disk Agent
- Data Protector requires a Disk Agent to be installed on Backup Servers (clients with (filesystem) data to be backed up). Install the Disk Agent for two reasons:
- Why Install the Disk Agent?** To run a filesystem backup of the SAP R/3 Database Server. Run this backup (a test backup) *before* configuring your integration and resolve all problems related to the SAP R/3 Database Server and Data Protector.

To run a filesystem backup of data that *cannot* be backed up using the SAP R/3 Database Server.

- Media Agent

Install this component on Drive Servers (clients with connected devices).

Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

You have now successfully installed Data Protector software on the SAP R/3 Database Server system. The integration is not yet ready for use. The next section gives you instructions on the procedure for configuring the integration to make it ready for use.

Configuring the Integration

Configuring the Data Protector SAP R/3 integration consists of these steps:

1. If you intend to use the Oracle8 Recovery Manager to backup the SAP R/3 database files, install and configure the Data Protector Oracle8 integration. When the Data Protector Oracle8 integration is configured, it is recommended to run a test Data Protector Oracle8 backup using the Oracle8 Recovery Manager. Refer to the *HP OpenView Storage Data Protector Integration Guide* for information on how to do this. If the test backup fails, refer to the “Troubleshooting” section of the same guide to resolve the problems.
2. Configure the SAP R/3 user.
3. Configure the SAP R/3 Database Server.
4. Configure the SAP R/3 backup.

Configuring an SAP R/3 User in Data Protector

In order to start an SAP R/3 backup session, a user needs an operating system logon on the system where an SAP R/3 Database Server is running.

In addition, this user has to be registered in the Oracle database and identified by SAP R/3 through the operating system identification.

This means that Oracle Server does not request connection information from an application started under such user account, but only checks whether the user is registered in the database.

Refer to the SAP R/3 and Oracle documentation for further information about different types of connections, about roles and privileges of Oracle database administrators, and about security issues that should be considered.

Further on, this user is allowed to backup and restore an SAP R/3 database. In order to start a backup of an SAP R/3 database using Data Protector, this user has to become the owner of the Data Protector backup specification.

As the owner of the backup specification, the user has to be added to either the Data Protector admin or operator user group.

Such a user is the user ora<SID> from the group sapsys; or, you can identify such a user by running the following command on the SAP R/3 Database Server system:

```
ps -ef |grep ora_pmon_<ORACLE_SID>  
or  
ps -ef |grep ora_lgwr_<ORACLE_SID>
```

Figure 2-4

Finding the Oracle User



```
# ps -ef | grep ora_pmon  
ora 2675 1 4 Sep 24 ? 0:13 ora_pmon  
#
```

It can be seen from the example above that the user ora has sufficient privileges within the SAP R/3 database to backup and restore the SAP R/3 database. Therefore, this user has to be added to the corresponding Data Protector user group (admin or operator) and have to become the owner of the backup specification, so that the user is able to backup the SAP R/3 database using Data Protector.

IMPORTANT

Additionally, the operating system root user on the SAP R/3 Server also has to be added to either the Data Protector admin or operator user group.

After the two users are added, Data Protector sessions can be started under the user account with all the privileges required to perform an SAP R/3 database backup with Data Protector.

Sometimes SAP administrators want to enforce more security and allow restores to be performed only by using a specific user account (for example SAP administrator). In this case, this user should also be configured as a Data Protector user and have to be added in either the operator or administrator group.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for a detailed information on Data Protector user rights and how to add a user to a user group.

Configuring an SAP R/3 Database Server

Before You Begin It is recommended that you configure and run a Data Protector test filesystem backup of the SAP R/3 Database Server (a client system in the Data Protector cell).

In case of problems, this type of backup is much easier to troubleshoot than the integration itself.

A test filesystem backup includes installing a Disk Agent on the SAP R/3 Database Server. Any device can be used for the test purposes only. Configure a standard filesystem backup, which can include one directory only. The test should include a partial restore to the SAP R/3 Database Server as well.

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

Configuring the SAP R/3 Database Server involves preparing the environment for performing backups. The environment parameters such as the Oracle home directory and the connection string to the Oracle Target Database are saved on the Cell Manager. The database must be online during the configuration procedure.

NOTE

Each SAP R/3 instance must be configured separately.

NOTE

Make sure to set any Oracle8 and SAP R/3 related environment variables needed for the Oracle8 and SAP R/3 databases to function properly (for example, the Oracle8 NLS_LANG environment variable) on the SAP R/3 Database Server. Refer to Oracle8 and SAP R/3 documentation for more information.

Configuration of an SAP R/3 Database Server is performed using the `/opt/omni/sbin/util_sap.exe` command on HP-UX and Solaris, or the `/usr/omni/bin/util_sap.exe` command on other UNIX systems.

**The
util_sap.exe
Command**

Use the `util_sap.exe` command to get the information you may need to configure your SAP R/3 Database Server. This will:

- List all Oracle instances on a particular system.

```
util_sap.exe -APP
```

- List the tablespaces that belong to a particular Oracle8 instance:

```
util_sap.exe -OBJS0 <ORACLE_SID>
```

- List the database files that belong to a particular tablespace of the Oracle instance:

```
util_sap.exe -OBJS1 <ORACLE_SID> <TABLESPACE>
```

Using the CLI

To configure an SAP R/3 Database Server, execute the following command with root privileges on the SAP R/3 Database Server:

NOTE

Each instance must be configured separately.

```
util_sap.exe -CONFIG <ORACLE_SID> <ORACLE_HOME> \  
<targetdb_connection_string> <SAPTOOLS_DIR> \  
[<SAPDATA_HOME>], where:
```

- *<ORACLE_SID>*
is the name of the Oracle database instance to be configured
- *<ORACLE_HOME>*
is the directory in which Oracle binaries are installed
- *<targetdb_connection_string>*
is the login information to the target database of the
<user_name>/<password>@<service> format, described in
“Glossary” on page G-1.

The *<user_name>* is the name by which a user is known to Oracle Server and to other users. Every user is identified by a password, and both must be entered to connect to an Oracle database. This user is, by default, used by *brbackup* and *brarchive* during backup. To define a different user when backing up, use the *-u <user_name>* as a BR Backup SAP R/3 backup option. Refer to “SAP R/3 Backup Options” on page 155.

NOTE

The user `<user_name>` is visible during backup when the `ps -ef` command is run.

- `<SAPTOOLS_DIR>`
is the directory in which SAP R/3 backup utilities are stored.
- `<SAPDATA_HOME>`
Directory where SAP R/3 database files are installed. This is an optional parameter. By default, it is set to `<ORACLE_HOME>`.

Using the GUI

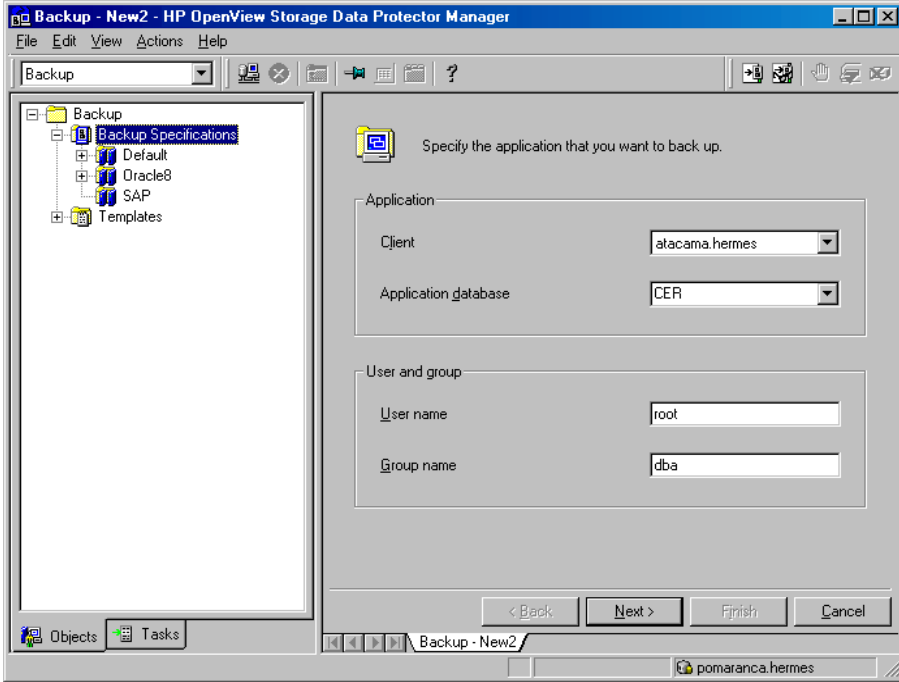
To configure an instance of the SAP R/3 Database Server, perform the following steps using the Data Protector GUI:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then expand Backup Specifications, and right-click SAP R/3.
3. Click Add Backup. In the Create New Backup dialog box, double-click the Blank SAP Backup template or any of the pre-defined templates.

The properties of a particular backup template can be seen in the corresponding pop-up window.

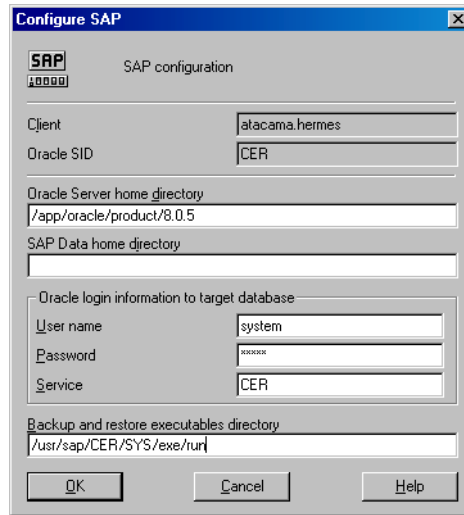
4. In the Results Area of the next page of the wizard, enter the following information:
 - Name of the SAP R/3 Database Server you want to configure.
 - Name of the Oracle Server instance (`ORACLE_SID`) on which the SAP R/3 Database Server is running.
 - UNIX user name and user group of the SAP R/3 user, as described in the “Configuring an SAP R/3 User in Data Protector” on page 140.

Figure 2-5 Specifying the SAP R/3 Database Server and the Oracle SID



Once you have provided the required information, click Next. If the selected system is configured for the first time, the configuration window is displayed.

Figure 2-6 Configuring an SAP R/3 Database Server



5. Enter the following information in the Configure SAP dialog box:

- The Oracle Server home directory. If not specified, this is set to the default Oracle8 home directory.
- SAP data home directory (if not entered, this is set to `<ORACLE_HOME>`)
- The connection string to the Oracle Target Database.

See “Glossary” on page G-1 for more information on login connection strings.

- The directory where the SAP R/3 backup utilities are stored.

By default, the utilities reside in the `/usr/sap/ <ORACLE_SID>/SYS/exe/run` directory.

What Happens?

The following happens after saving the configuration.

Data Protector starts the `util_sap.exe` file on the SAP R/3 Database Server, which performs the following:

1. Saves the configuration parameters in the Data Protector integration configuration on the Cell Manager in the

`/etc/opt/omni/integ/config/SAP/<client_name>%<ORACLE_SID>` file (HP-UX or Solaris Cell Manager), or in the `<Data_Protector_home>\Config\integ\config\sap\<client_name>%<ORACLE_SID>` file (Windows Cell Manager).

2. Creates a UNIX soft link for `backint` from the directory in which SAP R/3 utilities are stored to `/opt/omni/1bin` (on HP-UX and Solaris systems), or to `/usr/omni/bin` on other UNIX systems.

Checking the SAP R/3 Configuration - Data Protector GUI

To check the configuration of your SAP R/3 Database Server, proceed as follows:

1. Right-click the SAP R/3 Database Server system.
2. Click Check Configuration.

If the configuration is successful, you should receive a message confirming that the integration was properly configured.

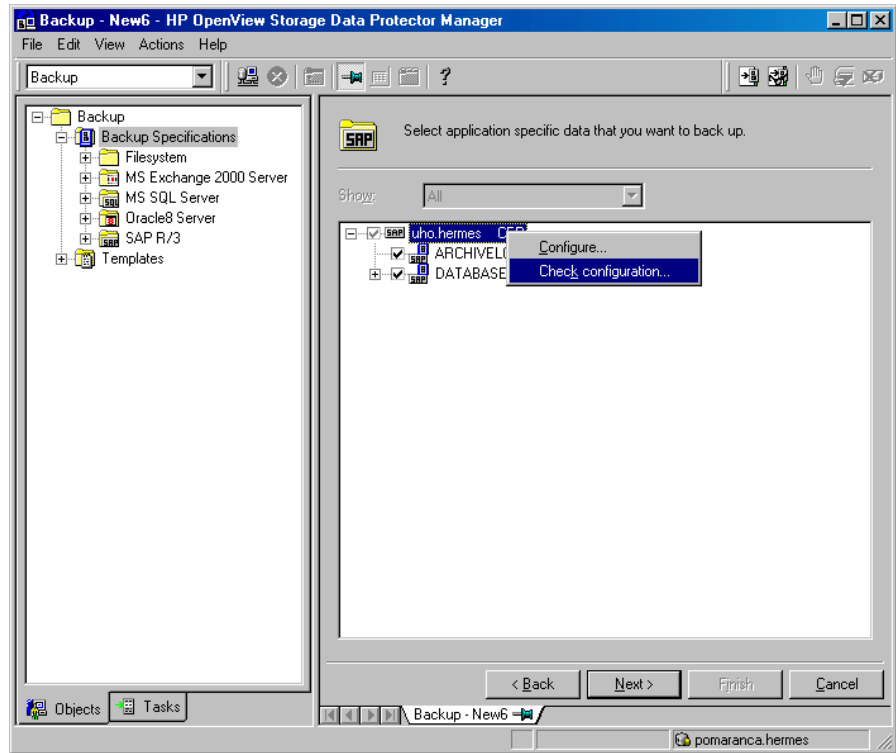
NOTE

The target database must be online during the check.

The configuration can also be checked if you have already created and saved a backup specification for a particular SAP R/3 Database Server. Proceed as follows:

1. In the Data Protector Manager, switch to the Backup context. In the Scoping Pane, expand Backup, Backup Specification, then SAP R/3.
2. In the Results Area, double-click the backup specification, then select Properties.
3. In the Source property page, right-click the name of the SAP R/3 Database Server, then click Check Configuration.

Figure 2-7 Checking the SAP R/3 Configuration



You can also (re)configure an SAP R/3 Database Server by right-clicking it and selecting Configure.

Checking the SAP R/3 Configuration - Data Protector CLI

To check the SAP R/3 configuration, start the following command on the client:

```
util_sap.exe -CHKCONF <ORACLE_SID>.
```

Data Protector verifies the configuration by attempting to connect to the SAP R/3 Database Server using the information that was specified and saved during the configuration.

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

To get the error description, start the `/opt/omni/lbin/omnigetmsg 12 <error number>` command for Solaris and HP-UX platforms or the `/usr/omni/bin/omnigetmsg 12 <error number>` command for other UNIX platforms.

Configuring an SAP R/3 Backup

To configure an SAP R/3 backup, perform the following steps:

1. Configure media and devices needed for backup. See the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instructions.
2. Create a backup specification, specifying what to back up and how to back it up.
3. Create or modify the parameter file on the SAP R/3 Database Server.

Creating a New Template

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

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

To create a new backup template, proceed as follows in the Data Protector Manager:

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Templates and right-click SAP R/3.
3. Click Add Template. Follow the wizard to define the appropriate backup options in your template.

You can also modify any of the existing pre-defined templates.

Creating a Data Protector SAP R/3 Backup Specification

SAP R/3 backup specifications are located in the following directory on the Cell Manager:

`/etc/opt/omni/barlists/sap/` (HP-UX and Solaris systems) or

`/usr/omni/config/sap/` (other UNIX systems)

An SAP R/3 backup specification is created using the Data Protector GUI.

Creating an SAP R/3 Backup Specification

To create an SAP R/3 backup specification using the Data Protector User Interface, proceed as follows:

1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications.
3. Right-click SAP R/3 and then click Add Backup. The Create New Backup dialog box is displayed.
4. Double-click Blank SAP Backup to create a backup specification without predefined options, or use one of the pre-defined templates given below:

Brarchive_CopyDeleteSave	Creates a second copy of the offline redo logs, saves them, deletes them after the backup, and then archives the newly-created redo logs.
Brarchive_Save	Archives the offline redo logs.
Brarchive_SaveDelete	Archives the offline redo logs, and then deletes them after the backup.
Brarchive_SecondCopyDelete	Creates a second copy of the offline redo logs that have been already archived, and then deletes them after the backup.
Brbackup_Offline	Backs up the shut-down database.
Brbackup_Online	Backs up the active database. The tablespace is locked for the time of the whole backup.
Brbackup_Online_Fast	Backs up the active database. The tablespace is locked for the time of its backup.
Brbackup_RMAN_Offline	Backs up the shut-down database.

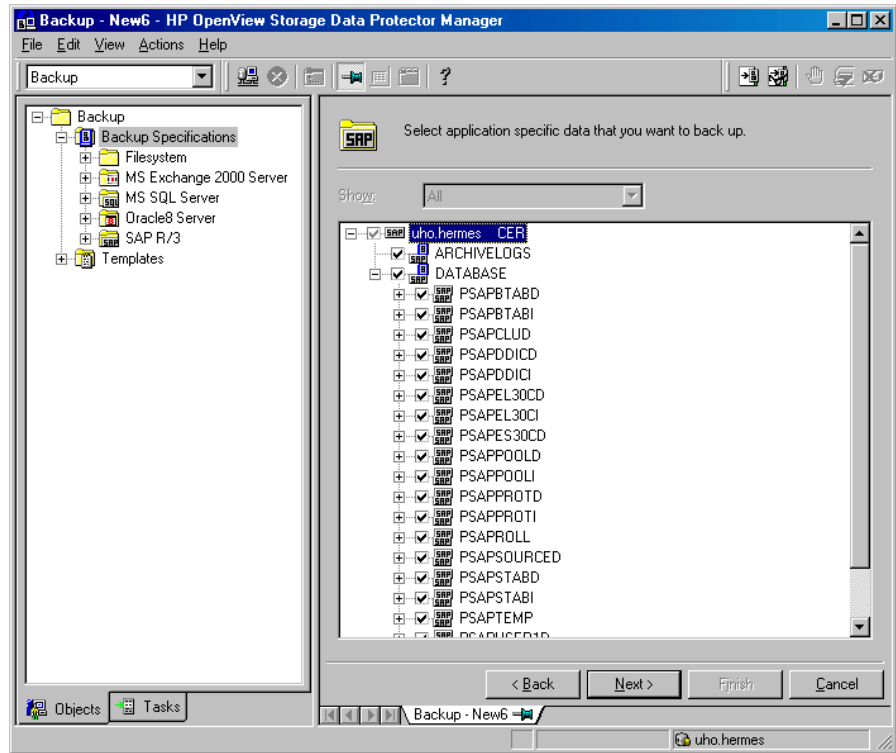
Brbackup_RMAN_Online	Backs up the active database. The tablespace is locked for the time of the whole backup.
Brbackup_SMB_Offline	Performs the database backup of the mirror disks. The database is stopped during the splitting of the mirror disks.
Brbackup_SMB_Online	Performs the database backup of the mirror disks. The database is active during the splitting of the mirror disks.

5. In the Results Area, enter the following information:
 - Name of the SAP R/3 Database Server you want to configure.
 - Name of the Oracle Server instance (ORACLE_SID) on which the SAP R/3 Database Server is running.
 - UNIX user name and user group of the SAP R/3 user as described in the “Configuring an SAP R/3 User in Data Protector” on page 140.

Click Next. If the SAP R/3 Database Server is configured successfully, the Source dialog box is displayed. Otherwise, you are prompted to configure the SAP R/3 Database Server. See “Configuring an SAP R/3 Database Server” on page 142 for details.

6. In the Source property page, select the database objects you want to back up. Database objects include archive logs, tablespaces, and data files.

Figure 2-8 Selecting Backup Objects



Refer to “Why Archive Redo Logs?” on page 155, for an explanation of the reasons for archiving redo logs, and online Help for details on backup objects.


7. Define options, devices, and the scheduling policy to be used by following the wizard.

Refer to Data Protector online Help and the *HP OpenView Storage Data Protector Administrator’s Guide* for backup options common to all objects.

See “SAP R/3 Backup Options” on page 155 for details about SAP R/3-specific options.

Once you have defined all backup options, you have to save and name your SAP R/3 backup specification. It is recommended that you save all SAP R/3 backup specifications in the SAP group.

Figure 2-9 Saving the Backup Specification



You have now completed the creation of an SAP R/3 backup specification. After the backup specification is saved, verify that the owner of the backup specification is the specified SAP R/3 user. See “Configuring an SAP R/3 User in Data Protector” on page 140 for details about this user.

8. You can examine the newly-created and saved backup specification in the Backup context, under the specified group of backup specifications. The backup specification itself is stored in the `/etc/opt/omni/barlists/sap/<Backup_Spec_Name>` (HP-UX and Solaris systems) or `/usr/omni/config/sap/<Backup_Spec_Name>` (other UNIX systems) file on the Cell Manager system.

When the backup specification is saved, the SAP configuration, which stores information about parallelism and balancing types, is also automatically saved on the SAP R/3 Database Server.

It is recommended that you test the backup specification by clicking Start Preview. This is an interactive test which does not back up any data.

You can start a real interactive backup that includes data transfer by clicking Start Backup.

Note that you can edit backup specifications once you have specified all the backup options.

The `Use default RMAN channels` option is valid only if SAP R/3 uses RMAN for backing up the Oracle8 Target database.

NOTE

The parallelism of a backup (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the parallelism represents the sum of the device concurrencies defined in the SAP R/3 backup specification. See the *HP OpenView Storage Data Protector Administrator's Guide* for more information on load balancing.

The database system of an SAP R/3 system must operate in the ARCHIVELOG mode. This prevents the overwriting of online redo log files that have not yet been saved. To protect the archived directory from overflowing, empty the directory regularly.

Why Archive Redo Logs?

The reasons for archiving redo log files are listed below:

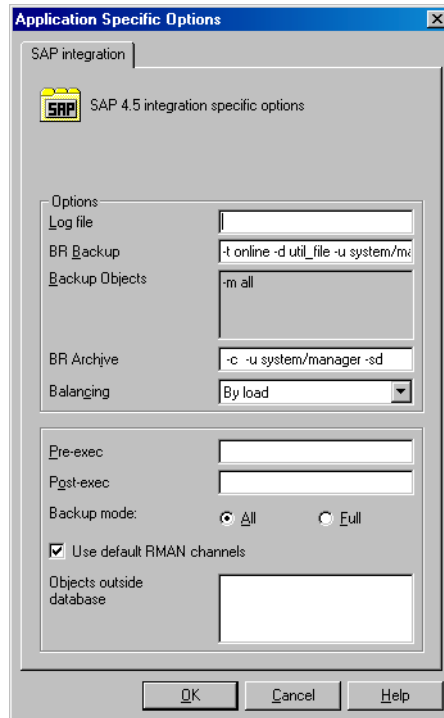
- In the event of a failure, consistent database status can only be recovered if all the relevant redo log files are available.
- An online backup of data files is useless if the related redo log files are missing. It is therefore necessary to archive the redo log files generated during the online backup immediately after running BRBACKUP.

SAP R/3 Backup Options

SAP R/3 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 SAP R/3 backup specification by clicking the Advanced button.

Figure 2-10 SAP R/3 Backup Options



Log file

Specifies the pathname of the backint log file. By default, this log file is not generated, as Data Protector stores all relevant information about backup sessions in the database. However, the user may decide to enable local logging by specifying a log file pathname.

BR Backup

Enter the BRBACKUP command options. See SAP R/3 Online Documentation for information about BRBACKUP command options. For example, type `-t online`, for online backup.

Or, type `-u <user_name>` for some other user than default user (usually the user system).

Backup Objects

When the backup specification is saved, this field lists the string passed by `omnisap.exe` to the `BRBACKUP` command.

BR Archive

Enter the `BRARCHIVE` command options. See the *SAP R/3 Online Documentation* for information about `BRARCHIVE` command options.

Balancing: By Load

Groups files in subsets by size so that the amount of data on all backup devices is approximately the same. Each subset is backed up by one Data Protector `sapback` program, thus allowing concurrent backup of all subsets.

If this option is set and your backup device uses hardware compression, the size of the backed up file on the medium will not be the same as on the disk. To make Data Protector aware of this, make sure that you specify the size of the backed up file on the medium in the compression section of the Data Protector SAP R/3 configuration file. Refer to “Data Protector SAP R/3 Configuration File” on page 132 for information on how to do this.

Balancing: By Time

Groups files in subsets so that backup to all backup devices takes approximately the same time. This depends on the file types, the speed of the backup devices, and external influences (such as mount prompts), and is therefore best for environments with large libraries of the same quality. Each subset is

backed up by one Data Protector `sapback` program, thus allowing concurrent backup of all subsets of the same type. Data Protector automatically stores backup speed information in the `speed` section of the Data Protector integration configuration file on the Cell Manager. It uses this information to optimize backup time.

This type of balancing may lead to non-optimal grouping of files in the case of online backup, or if the speed of backup devices varies significantly among devices.

Balancing: Manual

Manual balancing optimizes backups by allowing you to group files into subsets and back up these subsets using specific devices. See “Manual Balancing of Files into Subsets” on page 163 for more information.

Balancing: None

No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order use the Oracle Server Manager SQL command: `select * from dba_data_files`

Pre-exec

Specifies an object pre-exec command with options that will be started on the SAP R/3 Database Server before backup. The command/script is started by Data Protector `ldbar.exe` and has to reside in the `/opt/omni/bin` on HP-UX and Solaris systems, and `/usr/omni/bin` directory on other UNIX systems. Only the filename must be provided in the backup specification.

Post-exec	Specifies an object post-exec command with options that will be started on the SAP R/3 Database Server after backup. The command/script is started by Data Protector <code>ldbar.exe</code> and has to reside in the <code>/opt/omni/bin</code> on HP-UX and Solaris systems, and <code>/usr/omni/bin</code> directory on other UNIX systems. Only the filename must be provided in the backup specification.
Backup mode	<p>Specifies the type of RMAN backup to be used. This option is disabled if tablespaces and not the whole database are configured backup.</p> <p>If <code>All</code> is specified, RMAN backs up the complete database.</p> <p>If <code>Full</code> is specified, RMAN performs the Full backup (level 0), thus enabling RMAN incremental backups.</p>
Use default RMAN channels	Enter the concurrency value for your backup. This number overrides the parameter set in the initialization parameter file.
Objects outside database	<p>With this option, you save all non-database files of the SAP R/3 and Oracle8 environments. This means that the following directory trees can be saved:</p> <pre>/sapmnt/<ORACLE_SID> /usr/sap/<ORACLE_SID>, /usr/sap/trans,<ORACLE_HOME></pre> <p>It is recommended that you save these directories in a separate backup session.</p>

NOTE

Note that the `sapdata<n>` and `saplog` or `origlog/mirrlog` subdirectories of the `<SAPDATA_HOME>` directory should not be saved.

See online Help and the *HP OpenView Storage Data Protector Administrator's Guide* for details on other specific Data Protector backup options.

Creating or Modifying the Parameter File on the SAP R/3 Database Server

The parameter file is used by SAP R/3 to set specific SAP R/3 backup options in case these options are not yet specified using the backup command. A template for the parameter file is located on the SAP R/3 Database Server as `<ORACLE_HOME>/dbs/init<ORACLE_SID>.sap`, where `<ORACLE_SID>` represents the identifier for your database.

To link the Data Protector SAP R/3 Integration Module with the SAP R/3 backup and restore interface, modify the `backup_dev_type` parameter in the parameter file.

You can find this parameter in the following section of the parameter file:

```
# backup device type
# [disk | tape | tape_auto | pipe | pipe_auto | rman_util
| util_file_online | util_file ]
# default: tape
backup_dev_type = util_file
```

You can perform two types of online backups as well as offline backups.

- To start an offline backup, specify the `-t offline` and `-d util_file BRBACKUP` options. You can alternatively specify `backup_dev_type = util_file` and `backup_type = offline` in the SAP parameter file.
- The two types of online backups differ according to the duration in which tablespaces are in backup mode.

If the `-t online` and `-d util_file BRBACKUP` options are specified, SAP R/3 puts all tablespaces in backup mode before the backup begins, and puts them back into normal mode after the

backup. The same is achieved by specifying `backup_dev_type = util_file` and `backup_type=online` in the SAP parameter file.

If the `-t online` and `-d util_file_online` `BRBACKUP` options are specified, SAP R/3 puts individual objects in backup mode before the backup begins, and puts them back into normal mode after the backup. The same is achieved by specifying

`backup_dev_type = util_file_online` and `backup_type=online` in the SAP parameter file.

Refer to SAP R/3 documentation for more information.

Backing Up Using Recovery Manager

Benefits

Version 4.5 and higher of the SAP R/3 backup utilities allows Oracle8 data files to be backed up using RMAN mode. RMAN mode is in general transparent to the user. The User Interface remains unchanged and allows the use of new options. The most important benefit of RMAN mode is that the underlying Oracle8 database can be backed up *incrementally*.

The backup procedure using RMAN mode is very similar to the one for the underlying Oracle8 database using the Data Protector Oracle8 integration. The following restrictions must be taken into account when RMAN is used directly:

- The RMAN stores information about backups in the recovery catalog. For security reasons, this catalog should be kept in a separate database. This requires more administrative work.
- In a disaster situation (such as the loss of a production database and recovery catalog), the restoration and recovery of data is complicated. It may not be possible without the help of Oracle Support. If the Recovery Manager does not have administrative data stored in the recovery catalog, it cannot recover the database on the basis of the backups that have been made.

IMPORTANT

If the SAP R/3 integration is configured with the user `Internal`, the offline SAP R/3 backup using the RMAN fails. Configure the integration using the user `System`.

The integration of RMAN into the BRBACKUP SAP backup utility offers some important benefits:

- The recovery catalog is not used. Information about backups is saved in the control file and SAP log files. After each backup, the control file and SAP log files are saved. When data is restored, the control file is copied back first and then the data files. In case of a disaster, restore SAP log files before restoring any data files.
- Other important files will still be automatically backed up using the backint program.
- All previous SAP backup strategies can still be used with RMAN. However, RMAN cannot be used for offline redo log backups with BRARCHIVE, for standby database backups, or for split mirror backups.

Configuring the SAP R/3 Backup

In order to use SAP R/3 with the Oracle8 RMAN utility in order to back up the Oracle8 Target Database data files, you need to link the Oracle8 Server with the Data Protector Database Library. Refer to Chapter 1, “Integrating Oracle8/9 and Data Protector,” for more information.

Creating an SAP R/3 Backup Specification

To create a backup specification that enables you to use RMAN for backing up, specify the following BRBACKUP option:

`-d rman_util` (using BRBACKUP) or `backup_dev_type` (in the SAP parameter file)

NOTE

If the BRBACKUP option `-d` is not specified, the default value is taken from the SAP parameter file. In this case you have to make sure that the value for `rman_params` is set correctly in the SAP parameter file.

Before starting an incremental backup, ensure that the appropriate full backup is done using the following option:

`-m full` (using BRBACKUP) or `backup_mode=full` (in the SAP parameter file)

Incremental Backups if Using RMAN

To start an incremental backup, specify the Incr mode in the Data Protector GUI or the incr mode in the CLI, as follows:

```
omnib -sap_list <SAP_Backup_Specification> -barmode incr
```

Manual Balancing of Files into Subsets

Manual balancing allows you to precisely tailor the performance of an SAP R/3 backup by grouping files into subsets that are backed up in parallel. Make sure that:

- You use only one file from the same hard disk at a time.
- The number of files in a subset is equal to or smaller than a concurrency number, that is, the sum of concurrencies of all devices configured in the backup specification.
- If you do not specify all the files, other files that need to be backed up are added to the list automatically using the `load balance` option. Before backup, this list of files is logged in the `<ORACLE_HOME>/sapbackup/.*.lst` file.
- You specify the file subsets in the `manual_balance` section of the Data Protector integration configuration on the Cell Manager in the `/etc/opt/omni/integ/config/SAP/<client_name>%<ORACLE_SID>` file (HP-UX or Solaris Cell Manager), or in the `<Data_Protector_home>\Config\integ\config\sap\<client_name>%<ORACLE_SID>` file (Windows Cell Manager). Refer to “Data Protector SAP R/3 Configuration File” on page 132.

Creating an SAP /R3 Backup Specification for Manual Balancing

To use manual balancing, you have to edit the SAP R/3 backup specifications. The backup specifications are specified in the `/etc/opt/omni/barlist/sap` directory (HP-UX or Solaris Cell Manager) or in the `<Data_Protector_home>\Config\Barlists\SAP` directory (on Windows Cell Managers). In the backup specification, define which backup set will be backed up to which device. Use the `-restype` option followed by the ID numbers of the sets to be backed up by a specific device.

Example

To back up three subsets identified by ID numbers 1, 3, and 4, using a device named device2, specify the following:

```
DEVICE "DEVICE2"  
{  
  -restype "1 3 4"  
}
```

Note that the files in the specified subsets are thus backed up using only the specified device. To optimize backup performance, the number of sets for a device should be equal to the concurrency of the device.

Ensure that all the subsets are specified for backup using a specific device, or they will not be backed up. To ensure that all the subsets are backed up, even if you do not specify them for backup using a specific device, configure one device without the `-restype` option. All the subsets not configured for backup using a specific device will be backed up on this device.

Save the backup specification before using it.

Example of Configuration:

Suppose that you have two devices, *Device_1*, with concurrency 2, and *Device_2*, with concurrency 1. You also have the following manual balance specified in the `manual_balance` section of the Data Protector integration configuration on the Cell Manager in the `/etc/opt/omni/integ/config/SAP/<client_name>%<ORACLE_SID>` file (HP-UX or Solaris Cell Manager), or in the `<Data_Protector_home>\Config\integ\config\sap\<client_name>%<ORACLE_SID>` file (Windows Cell Manager):

```
manual_balance={  
SAP-R3={  
fileA=0;  
fileB=1;  
fileC=0;fileD=2;}}
```

Configure your backup specification `SAP-R3` to back up the files `fileA`, `fileC` and `fileD` on device *Device_1*, and `fileB` on device *Device_2*.

The backup specification then looks like:

```
BARLIST "SAP-R3"  
OWNER <user> <group> galeja.zimco.com  
DEVICE "DEVICE1"
```

```
{
  -restype "0 2"
}
DEVICE "DEVICE2"
{
  -restype "1"
}
CLIENT "ORACLE_SID" galeja.zimco.com
{
  -exec omnisap.exe
  -args "-brb -t online -m all"
}
```

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a backup.

Testing Using the Data Protector GUI

Testing Procedure The procedure consists of checking the Data Protector part of the integration to ensure that communication within Data Protector is established, that the data transfer works properly, and that transactions are recorded either in the recovery catalog (if used) or in a control file. Proceed as follows to test the integration:

1. In the Data Protector Manager switch to the Backup context.
2. In the Scoping Pane, expand Backup, then expand Backup Specifications, SAP R/3 and 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. See online Help for a description of these options.

Testing Using the Data Protector CLI

A test can be executed from the CLI on the SAP R/3 Database Server system or on any other Data Protector client system within the same cell, provided that the systems have the Data Protector User Interface installed.

You must run the `omnib` command with the `-test_bar` option. Execute the following command:

On HP-UX and Solaris systems:

Integrating SAP R/3 and Data Protector

Configuring an SAP R/3 Backup

```
/opt/omni/bin/omnib -sap_list <backup_specification_name> \  
-test_bar
```

On other UNIX systems:

```
/usr/omni/bin/omnib -sap_list <backup_specification_name> \  
-test_bar
```

What Happens?

The session messages are displayed on the screen during the command execution, while the following happens:

The `omnisap.exe` program is started, which then starts the Data Protector `testbar` command. This command then checks:

- the communication within Data Protector,
- the syntax of the SAP R/3 backup specification,
- if the devices are correctly specified,
- if the required media reside in the devices.

Backing Up an SAP R/3 Database

The Data Protector SAP R/3 integration provides both online and offline database backup. To run a backup of an SAP R/3 database, use any of the following methods:

- Backup Methods**
- Schedule a backup of an existing SAP R/3 backup specification using the Data Protector Scheduler.
 - Start an interactive backup of an existing SAP R/3 backup specification. You can start a backup using the Data Protector GUI or the Data Protector CLI.
 - Start an interactive backup on SAP R/3 Database Servers using either the `brbackup` or the `sapdba` command.

NOTE

If you use `brbackup` or `sapdba` to start a backup session, you do not receive any Data Protector messages about the progress of the session.

Messages from the Data Protector backup session are logged in the Data Protector database. SAP R/3 messages generated by the `brbackup` or `sapdba` commands are logged to the Data Protector database only if Data Protector is used to start the backup.

Duplicate SIDs Concurrent backups of systems with the same Oracle SID in the same cell are not supported.

Backup Modes Configurable backup modes that were used in Data Protector versions earlier than A.03.00 are not supported in the current version of Data Protector. However, their functionality is now supported using templates.

Incremental Backups Before starting an incremental backup, ensure that the appropriate full backup is done using the following option (note that this is valid for SAP tools version 4.5 and later):

`-m full` (using `BRBACKUP`) or `backup_mode=full` (in the SAP parameter file).

Integrating SAP R/3 and Data Protector Backing Up an SAP R/3 Database

To start an incremental backup, specify `Incremental` mode in the Data Protector GUI or `incr` mode in the CLI, as follows:

```
omnib -sap_list <SAP_Backup_Specification> -barmode incr
```

NOTE

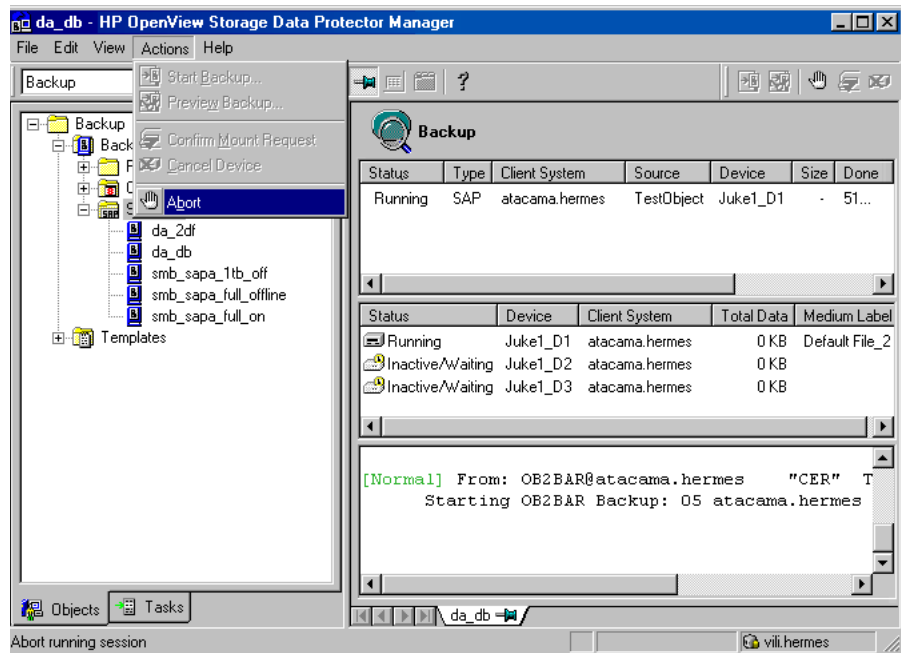
It is not possible to perform an incremental backup in `backint` mode.

Aborting a Running Session

In the Actions menu, click `Abort` to abort a running SAP R/3 backup session, and then confirm the action.

Figure 2-11

Aborting an SAP R/3 Backup Session



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 need to keep the database online continuously, then you should back it up frequently, including backup of the Archived Redo Logs, which is required in case you need a recovery to a particular point in time.

For example, you may decide to perform daily backups and make multiple copies of the online redo logs and the Archived Redo Logs to several different locations.

Some examples of scheduling backups of production databases:

- Weekly full backup
- Daily incremental backup
- Archived Log backups as needed

To schedule an SAP R/3 backup, proceed as follows:

1. In the Data Protector Manager window, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click SAP R/3.

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 Incr. The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to Full.

In the case of zero downtime backup, but only for ZDB disk or ZDB disk/tape backups (instant recovery enabled), specify the Split mirror/snapshot backup option.

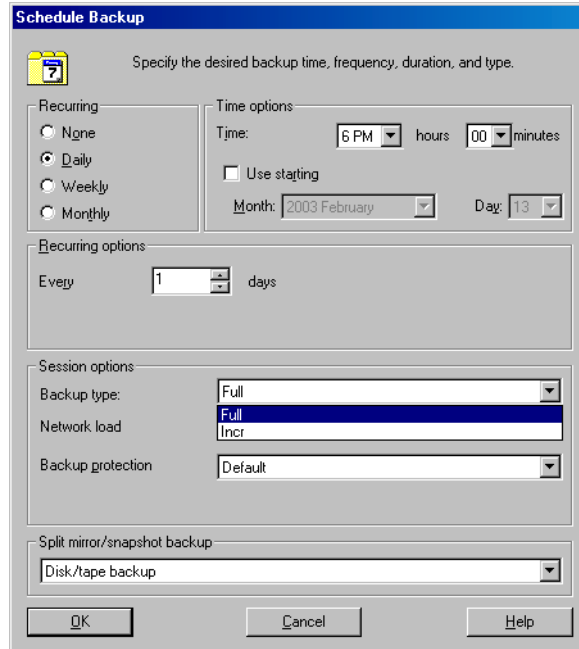
See Figure 2-12 on page 170.

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

NOTE

It is not possible to perform an incremental backup in the backint mode.

Figure 2-12 **Scheduling Backups**



NOTE

When creating an SAP R/3 backup specification, you access the Data Protector Scheduler via the Backup Wizard.

See “Creating a Data Protector SAP R/3 Backup Specification” on page 150 for information about accessing the Backup Wizard.

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, with the corresponding backup specification scheduled for a later time.

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

Using the Data Protector GUI

Follow the procedure below to start an interactive backup of an SAP R/3 backup specification:

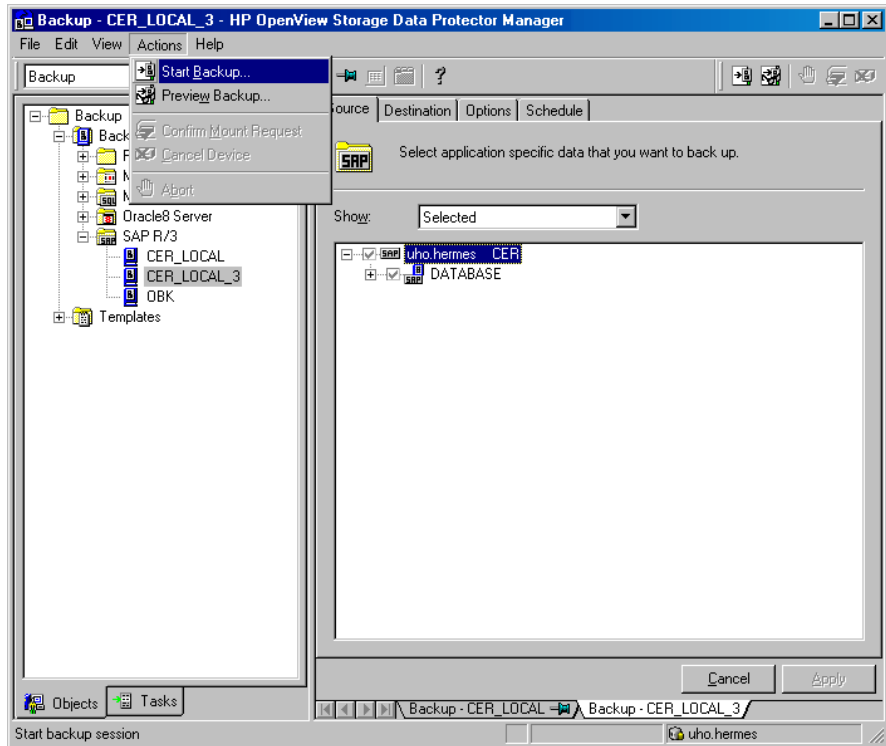
1. In the Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then expand Backup Specifications, and then SAP R/3.
3. Right-click the backup specification, then click Start Backup.

In the Start Backup dialog box, select the backup type and network load. The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to Full.

In the case of zero downtime backup, but only for ZDB disk or ZDB disk/tape backups (instant recovery enabled), specify the Split mirror/snapshot backup option.

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

Figure 2-13 Starting an Interactive Backup



Using the Data Protector CLI

An interactive backup can also be started from the CLI. Switch to the
`/opt/omni/bin` on HP-UX and Solaris systems

`/usr/omni/bin` on other UNIX systems

directory on an SAP R/3 Database Server system and run the following
command:

```
omnib -sap_list <backup_specification_name> [-barmode  
<SapMode>] [list_options]
```

You can select among the following *list_options*:

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

```
-load {low | medium | high}
-crc
-no_monitor
SapMode = { -full | -incr}
```

Example

To start a backup using an SAP R/3 backup specification, called RONA, execute the following command:

```
omnib -sap_list RONA
```

Using SAP R/3 Commands

When you interactively start a backup of your SAP R/3 object using the `brbackup` or `sapdba` commands, Data Protector uses the default SAP R/3 backup specification named `SAP-R3` for backup.

Starting a Backup Using Another Backup Specification

To start a backup using some other SAP R/3 backup specification, you must set the environment variable `OB2BARLIST` to the appropriate SAP R/3 backup specification name, and `OB2APPNAME` to the appropriate SAP R/3 backup system ID before starting the backup.

Set the environment variable by entering the following command *before* you enter the `brbackup` command or `sapdba` command:

```
export OB2BARLIST=<backup_specification_name>
export OB2APPNAME=<ORACLE_SID>
```

If you do not set this environment variable, Data Protector assumes that the SAP R/3 backup specification is named `SAP-R3`.

Restoring an SAP R/3 Database

An SAP R/3 database can be restored:

- Using the Data Protector GUI or CLI. A standard filesystem restore is performed.
- Using SAP R/3 commands

NOTE

You cannot perform a restore of backups created by the Oracle8 RMAN using the Data Protector GUI or CLI.

Before you start to restore your data using the Data Protector User Interface, you need detailed information about backed up objects. See the following section on how to find the information you need to restore your data.

NOTE

If your disk is full before a restore, restoring of a filesystem with SAP R/3 data that was backed up using the `brbackup` command will fail, because the `brrestore` command needs additional disk space for restoring the control file and online log files. How much additional disk space you need depends on the amount of the backed up data.

Finding Information Needed for Restore

To find the information needed for a restore, follow the steps below:

Execute the following commands:

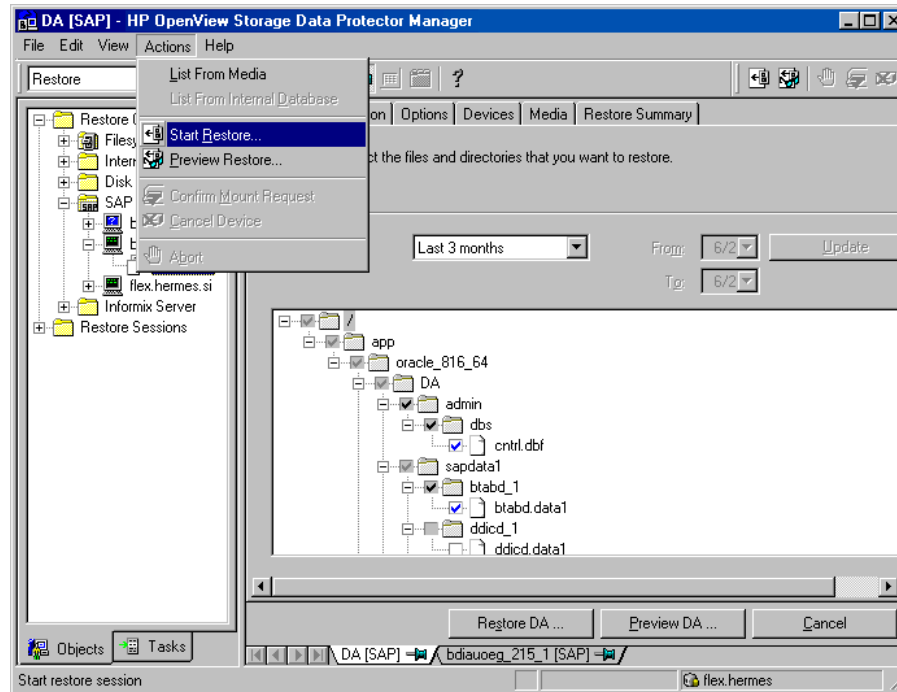
- `omnidb -sap`
to get a list of SAP R/3 objects.
- `omnidb -sap <object_name>`
to get details on a specific object, including the SessionID.

Using the Data Protector GUI

To restore your data, proceed as follows in the HP Data Protector Manager:

1. In the `Context List`, select `Restore`.
Expand `Restore`, then `SAP R/3`, and then the `SAP R/3 Database Server` from which you want to restore.
A list of backup objects is displayed in the `Results Area`.
2. Select the `SAP R/3` object you want to restore.
You can also select the search interval for browsing object versions in the `Data Protector` database by clicking the drop-down list button of the `Search Interval` option. If you select `Interval` in the drop-down list, you can set your own search interval by specifying the `From:` and `To:` options and then clicking the `Update` button.
3. Select the media and devices needed for the restore.
4. Click `Start Restore` and then `Finish` to start the restore session, or click `Next` to select the `Network Load` and `Report Level` before starting the restore session.

Figure 2-14 Restoring SAP R/3 Database Objects



When the session starts, messages are displayed in the Results Area. Upon successful completion, a message is issued in the Session Information dialog box.

Using the Data Protector CLI

The omnir Command

Using the CLI, execute the following command:

```
omnir -sap <Host:Set> -session <Session_ID> -tree  
<File_name>
```

Example

```
omnir -sap corsa.hermes.si:ABA.0 -session 2000/02/23-1 -tree  
/app/oracle805/ABA/sapdata1/btabd_1/btabd_1.dat
```


The restore session can be monitored in the Data Protector Monitor window, where mount prompts for the required media are also displayed. Refer to the man pages for more information on the Data Protector `omnir` command.

TIP

If you have a sparse file, restore using the `sparse` option to perform a faster restore.

Use any of the following methods to set the `sparse` option:

- Execute the following command: `export OB2SPARSE=sparse` if the restore is started using the SAP `sapdba` or `brrestore` commands.
- Set `Restore Sparse Files` in the `Restore Options` window if the restore is started using the Data Protector GUI.
- Set restore option `-sparse` if the restore is started using the Data Protector `omnir` command.

Using SAP R/3 Commands

The `sapdba` or `brrestore` Commands

You can use `sapdba` or `brrestore` to restore the target database. Both commands use the Data Protector `backint` interface to restore files backed up using Data Protector.

If you have backups of two different Oracle Servers with the same `ORACLE_SID` but on different SAP R/3 Database Servers, set the `OB2HOSTNAME` variable before starting restore to the name of the SAP R/3 Database Server from which you want to restore.

Example

```
export OB2HOSTNAME=<client_name>
```

```
sapdba
```

See the *SAP R/3 System Online Documentation* for instructions on how to use these utilities.

Using Another Device

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

Restoring Using the Data Protector GUI

If you are performing a restore using the Data Protector GUI, 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.

Restoring Using the Data Protector CLI or SAP R/3 Commands

If you are performing a restore using the Data Protector CLI or SAP R/3 commands, specify the new device in the `/etc/opt/omni/cell/restoredev` file in the following format:

```
“DEV 1” ”DEV 2”
```

where,

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

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

Example

Suppose you have SAP R/3 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 intended to be used as a guideline.

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

This is a general procedure for recovering an application:

1. Complete the recovery of the operating system.

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

Monitoring an SAP R/3 Backup and Restore

During a backup, system messages are sent to both the SAP R/3 Database Server and the Data Protector monitor. Thus, you can monitor a backup session from either the SAP R/3 Database Server or from any Data Protector client in the network where the User Interface is installed.

When it is detected that no more data can be backed up on the media, either because they are not in a device or because they are full, and a mount prompt is issued, the message is sent to the Data Protector monitor only, not to SAP R/3. Change the media and confirm the mount prompt in Data Protector.

Configuring the Integration as Cluster-Aware

Installation and Configuration

The Data Protector SAP R/3 integration can be configured in the MC Service Guard cluster. This means that either the Data Protector Cell Manager can be configured in a cluster, or Data Protector client can be configured in the cluster. Refer to *HP OpenView Storage Data Protector Concepts Guide* for more information on supported configurations.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* and “Configuring the Integration” on page 140 for information on how install and configure the Data Protector SAP R/3 MC Service Guard integration.

When configuring the Data Protector SAP R/3 integration, configure it only on one of the cluster nodes per one SAP R/3 server, since the Data Protector SAP R/3 configuration file resides on the Cell Manager. Use the virtual hostname when configuring the integration. However, you need to create a link to the Data Protector backint interface on all other nodes. Enter the following command on all other nodes:

```
ln -s /opt/omni/sbin/backint \  
/usr/sap/<ORACLE_SID>/sys/exe/run
```

NOTE

The environment variable `OB2BARHOSTNAME` must be set to the virtual hostname before running the configuration or backup from the command line (on the client). When the GUI is used, this is not required.

For example, if the configuration is run on the cluster physical node `physical_1.domain.com` and SAP R/3 is running on virtual host `virtual.domain.com`, the variable setting is:

```
export OB2BARHOSTNAME=virtual.domain.com
```

You also need to edit the Data Protector `omnirc` file on each cluster node and specify the name of the cluster node in the `SAPLOCALHOST` variable. Below you see an example of the `omnirc` file:

Integrating SAP R/3 and Data Protector

Configuring the Integration as Cluster-Aware

```
# SAP R/3 related entries for clustering
#
SAPLOCALHOST=<cluster_node_name>
```

NOTE

Make sure that the `SAPLOCALHOST` variable is not defined in the Environment section of the Data Protector SAP R/3 configuration file. Refer to “Data Protector SAP R/3 Configuration File” on page 132 for information on how to do that.

Add the SAP R/3 group dba user to Data Protector for the virtual server and for every node in the cluster. Refer to *HP OpenView Storage Data Protector Administrator’s Guide* for information on how to add a user to Data Protector.

For information on the Data Protector Cell Manager package configuration (if you want to install and configure an Data Protector Cell Manager in the MC/SG cluster), refer also to the *HP OpenView Storage Data Protector Administrator’s Guide*.

Backup and Restore

When creating a Data Protector SAP R/3 MC/SG cluster backup specification, always select the virtual hostname in the cluster and not a particular node.

There are some extra requirements that have to be fulfilled:

- Before you perform an offline backup, make sure that you take the Oracle Database resource offline and bring it back online after the backup.

This can be done using the `fscmd` command in the pre- exec and post- exec commands for the client system in a particular backup specification, or using the Cluster Administrator.

Refer to “Backing Up an SAP R/3 Database” on page 167 for information on how to create a Data Protector SAP R/3 backup specification and to *HP OpenView Storage Data Protector Administrator’s Guide* for information on MC/SG cluster backing up specifics.

Refer to “Restoring an SAP R/3 Database” on page 174 for information on how to restore an SAP R/3 database.

Troubleshooting

Before you start troubleshooting the Data Protector SAP R/3 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, measures have been proposed to help you work around the problem.

Using Oracle8 After Removing the Data Protector Oracle8 Integration

This section is relevant only if Oracle8 RMAN has been used to back up the SAP R/3 datafiles and you have uninstalled the Data Protector Oracle8 integration on an Oracle8 server.

After uninstalling the Data Protector Oracle8 integration on an Oracle8 server, the Oracle8 server software is still linked to the Data Protector Database Library. You have to rebuild the Oracle8 binary to remove this link. If this is not done, the Oracle8 server cannot be started after the integration has been removed.

Please refer to “Using Oracle8/9 After Removing the Data Protector Oracle8/9 Integration” on page 84 for more information on how to make the Oracle8 server functional again.

Prerequisites Concerning the Oracle Side of the Integration

The following steps should be performed to verify that Oracle is installed as required for the integration to work. These steps do not include verifying Data Protector components.

1. Verify that you can access the Oracle Target Database and that it is opened, as follows:

Export `<ORACLE_HOME>` and `<ORACLE_SID>` as follows:

- if you are using an SH - like shell enter the following commands:

```
ORACLE_HOME="<ORACLE_HOME>"
export ORACLE_HOME
ORACLE_SID = "<ORACLE_SID>"
export ORACLE_SID
```

- if you are using a CSH - like shell enter the following commands:

```
setenv ORACLE_HOME "<ORACLE_HOME>"
setenv ORACLE_SID "<ORACLE_SID>"
```

For Oracle8/8i, start the Server Manager from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl
```

For Oracle9i, start SQLPlus from the `<ORACLE_HOME>` directory

```
bin\sqlplus
```

At the SVRMGR> or SQLPLUS> prompt type:

```
connect internal
select * from dba_tablespace;
exit
```

If this fails open the Oracle Target Database.

2. **Verify that the TNS listener is correctly configured for the Oracle Target Database. This is required for properly establishing network connections:**

Export `<ORACLE_HOME>` as described on page 184 and start the listener from the `<ORACLE_HOME>` directory:

```
bin/lsnrctl start <service>
exit
```

If it fails, startup the TNS listener process and refer to Oracle documentation for instructions on how to create TNS configuration file (`LISTENER.ORA`)

Export `<ORACLE_HOME>` as described on page 184 and start the Server Manager:

```
bin/svrmgrl
At the SVRMGR prompt type
connect <Target_Database_Login>
exit
```

If it fails refer to the Oracle documentation for instructions on how to create a TNS configuration file (`TNSNAMES.ORA`).

3. **If you run backups in RMAN mode, verify that the Oracle8 Target Database is configured to allow remote connections with system privileges:**

Export `<ORACLE_HOME>` as described on page 184 and start the Server Manager from the `<ORACLE_HOME>` directory:

```
bin/svrmgrl
At the SVRMGR prompt type
connect <Target_Database_Login> as SYSDBA;
exit
```

Repeat the procedure using `SYSOPER` instead of `SYSDBA`. Set the `<ORACLE_HOME>` directory

If you use the Recovery Catalog:

```
bin/rman target <Target_Database_Login> rcvcat
<Recovery_Catalog_Login>
```

If you do not use the Recovery Catalog:

```
bin/rman target <Target_Database_Login> nocatalog
```

If this fails refer to the Oracle8 documentation for instructions on how to set up the password file and any relevant parameters in the `init<ORACLE_SID>.ora` file.

4. If you run backups in the RMAN mode, verify backup and restore directly to disk using the Recovery Manager channel type disk.

If you use the Recovery Catalog:

Export `<ORACLE_HOME>` as described on page 184 and start Recovery Manager:

```
bin/rman target <Target_Database_Login> rcvcat  
<Recovery_Catalog_Login> cmd_file=rman_script
```

If you do not use the Recovery Catalog:

Export `<ORACLE_HOME>` as described on page 184 and start Recovery Manager:

```
bin/rman target <Target_Database_Login> nocatalog  
cmd_file=rman_script
```

An example of the `rman_script` is listed below:

```
run {allocate channel 'dev0' type disk;  
backup (tablespace <tablespace_name>  
format '<ORACLE_HOME>/tmp/<datafile_name>');}
```

After a successful backup, try to restore the backed up tablespace by running the following restore script:

```
run {  
allocate channel 'dev0' type disk;  
sql 'alter tablespace <tablespace_name> offline immediate';  
restore tablespace <tablespace_name>;  
recover tablespace <tablespace_name>;  
sql 'alter tablespace <tablespace_name> online'  
release channel 'dev0';}
```

If one of the above procedures fails, refer to the Oracle8 documentation to learn how to execute backup and restore directly to disk using the Recovery Manager.

Prerequisites on the SAP R/3 Side of the Integration

The following verification steps must be performed in order to verify that SAP R/3 is installed as required for the integration to work. These steps do not include Data Protector components.

1. Verify backup directly to disk as follows:

```
brbackup -d disk -u <user>/<password>
```

If this fails, check the error messages and resolve possible problems before you continue.

2. Verify restore directly to disk as follows:

```
brrestore -d disk -u <user>/<password>
```

If this fails, check the error messages and resolve possible problems before you continue.

3. If you are running backups in RMAN mode, verify backup and restore directly to disk using Recovery Manager channel type disk as follows:

a. Relink the Oracle8 software with the Database Library provided by SAP R/3 (`libobk.sl`).

b. Use the same procedure as described for linking the Data Protector Database Library.

Refer to *HP OpenView Storage Data Protector Integration Guide* for information on how to do this.

IMPORTANT

Before you can use Data Protector again in the RMAN mode, you have to relink the Oracle8 again with the Data Protector Database Library.

c. You have to define the parameter `init` in the initialization file `init<ORACLE_SID>.ora`.

Run the following commands:

```
brrestore -d pipe -u <user>/<password> -t online -m all
```

```
brrestore -d disk -u <user>/<password>
```

If this fails, refer to the SAP R/3 Online Help to learn how to execute backup and restore directly to disk using the SAP R/3 backup utility.

Check the error message and resolve this issues before you continue.

4. **Verify that the SAP R/3 backup tools correctly start backint (which is provided by Data Protector):**

Move the original backint and create a test script named backint in the directory where the SAP R/3 backup utility resides, with the following entries:

```
#!/usr/bin/sh  
echo "Test backint called as follows:"  
echo "$0 $*"   
echo "exiting 3 for a failure"  
exit 3
```

Then start the following commands as the SAP R/3 user; see “Configuring an SAP R/3 User in Data Protector” on page 140:

```
brbackup -t offline -d util_file -u <user>/<password> -c
```

If you receive backint arguments, this means that SAP R/3 is properly configured for backup using backint; otherwise you have to reconfigure SAP R/3.

Refer to “Configuring an SAP R/3 Database Server” on page 142.

Configuration Problems

IMPORTANT

The procedure described in the previous sections must be performed before you start checking the Data Protector configuration.

1. **Verify that the Data Protector software has been installed properly.**

See “Installing and Upgrading the Data Protector SAP R/3 Integration” on page 138 for details.

2. Verify that the Data Protector Database Library is linked with the Oracle8 executable:

Use the following command to check if the `libob2oracle8.so` on Solaris and `libob2oracle8.sl` (`libob2oracle8_64bit.sl`) on HP-UX is linked with the Oracle8 executable.

Export the `<ORACLE_HOME>` and the `<ORACLE_SID>` as described on page 184.

HP-UX platform:

```
/usr/bin/chatr <ORACLE_HOME>/bin/oracle
```

Solaris platform:

```
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle
```

The output has to state that the respective Data Protector library is required by Oracle8 executable.

The following is an extract of the command output on HP-UX:

```
bin/oracle:
    shared executable
    shared library dynamic path search:
        SHLIB_PATH  enabled second
        embedded path disabled first Not Defined
    shared library list:
        static
/opt/omni/lib/libob2oracle8.sl(libob2oracle8_64bit.sl)
    dynamic /usr/lib/librt.2
    dynamic /usr/lib/libnss_dns.1
    dynamic /usr/lib/libdld.2
```

The line starting with `SHLIB_PATH` should be returned as in the example above. If this line is different, then enable the Data Protector Database Library dynamic path as follows:

```
/usr/bin/chatr +s enable <ORACLE_HOME>/bin/oracle
```

3. Perform a filesystem backup of the SAP R/3 Database Server:

Perform a filesystem backup of the SAP R/3 Database Server system so that you can eliminate any potential communication problems between the SAP R/3 Database Server and the Data Protector Cell Manager system.

Do not start troubleshooting an online database backup unless you have successfully completed a filesystem backup of the SAP R/3 Database Server system.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for details about how to do a filesystem backup.

4. Examine the environment variables:

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP R/3 configuration file on the Cell Manager. Refer to “Data Protector SAP R/3 Configuration File” on page 132.

5. Verify the permissions of the currently used user account:

Your user account has to enable you to perform backup or restore using Data Protector. Use the `testbar2` utility to check the permissions:

```
/opt/omni/bin/utilns/testbar2 -perform:checkuser
```

If the user account holds all required permissions, you will receive only `NORMAL` messages displayed on the screen.

See also “Configuring an SAP R/3 User in Data Protector” on page 140.

6. Examine system errors:

System errors are reported in the

```
/var/opt/omni/log/debug.log (HP-UX and Solaris systems) or  
/usr/omni/log/debug.log (other UNIX systems) file on the SAP R/3  
Server.
```

Backup Problems

At this stage, you should have performed all the verification steps described in the previous sections. If backup still fails, proceed as follows:

1. Check your SAP R/3 Server configuration:

To check the configuration, start the following command on the SAP R/3 Server system:

```
/opt/omni/lbin/util_sap.exe -CHKCONF <ORACLE_SID> (HP-UX  
and Solaris systems) or
```

```
/usr/omni/bin/util_sap.exe -CHKCONF <ORACLE_SID> (other  
UNIX systems)
```

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> (HP-UX and  
Solaris systems) or
```

```
/usr/omni/bin/omnigetmsg 12 <Error_number> (other UNIX  
systems)
```

The *RETVAL*0 indicates successful configuration.

2. Verify Data Protector internal data transfer using the testbar2 utility.

Before you run the testbar2 utility, verify that the Cell Manager name is correctly defined on the SAP R/3 Database Server. Check the /etc/opt/omni/cell/cell_server (HP-UX and Solaris systems) or /usr/omni/config/cell/cell_server (other UNIX systems) file, which contains the name of the Cell Manager system. Then run the following command:

```
/opt/omni/bin/utilns/testbar2 -type:SAP  
-appname:<ORACLE_SID> -bar:<backup_specification_name>  
-perform:backup (HP-UX and Solaris systems)
```

```
/usr/omni/bin/utilns/testbar2 -type:SAP  
-appname:<ORACLE_SID> -bar:<backup_specification_name>  
-perform:backup (other UNIX systems)
```

Examine the errors reported by the testbar2 utility by clicking the Details button in the Data Protector Monitor context.

If the messages indicate problems concerning the Data Protector side of the integration, proceed as follows:

- a. Check that the owner of the backup specification is the SAP R/3 backup owner as described in the “Configuring an SAP R/3 User in Data Protector” on page 140 and that this user belongs to the Data Protector operator or admin group.
- b. Check that the respective Data Protector user group has the See private objects user right enabled.
- c. Create an SAP R/3 backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices.

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

If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=<barlist_name>
```

```
export OB2APPNAME=<ORACLE_SID>
```

```
/opt/omni/lbin/backint -f backup -t file -u <ORACLE_SID>  
-i <input_file> (HP-UX and Solaris systems)
```

```
/usr/omni/bin/backint -f backup -t file -u <ORACLE_SID> -i  
<input_file> (other UNIX systems)
```

where <input_file> is a file with a list of full pathnames for backup.

Backint expects the list of files in the following format:

```
<pathName_1>
```

```
<pathName_2>
```

```
<pathName_3>
```

Backup fails at the beginning with the message Internal heap ERROR 17112

Problem

When using SAP 4.6D kernel on HP-UX 11.11, backup fails immediately after it was started due to a BRBACKUP core dump. A line similar to the following can be found at the beginning of the message:

```
Internal heap ERROR 17112 addr=0x800003ffff7f3660
```

Action

1. Login to the SAP server as the user who is owner of the backup specification.

2. Run the command

```
env | grep NLS_LANG
```

The output is similar to the following:

```
NLS_LANG=AMERICAN_AMERICA.US7ASCII
```

3. Add the NLS_LANG variable to the backup specification. For more details, refer to “Setting, Retrieving, Listing and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI” on page 135.
4. Restart the backup.

Restore Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. **Verify a user for the restore:**

Verify that user specified for the restore session is the user of backup session and that he/she belongs to the Data Protector operator or admin group.

See “Configuring an SAP R/3 User in Data Protector” on page 140

2. **Verify that a backup object exists on the backup media and in the IDB:**

This can be done by executing the command

```
/opt/omni/bin/omnidb -SAP "<object_name>" -session  
<Session_ID>" -media (HP-UX and Solaris systems) or
```

```
/usr/omni/bin/omnidb -SAP "<object_name>" -session  
<Session_ID>" -media (other UNIX systems)
```

on the SAP R/3 Database Server system.

The output of the command lists detailed information about the specified backup object, session IDs of the backup sessions containing this object, and a list of the media used.

For detailed syntax of the omnidb command, run:

```
/opt/omni/bin/man omnidb (HP-UX and Solaris systems)
```

```
/usr/omni/bin/man omnidb (other UNIX systems)
```

You can also do this using the SAP R/3 utilities:

Use `backint`, so that `SAPDBA` will also use this command to query:

```
/opt/omni/lbin/backint -f inquiry -u <ORACLE_SID> -i  
<input_file> (HP-UX and Solaris systems)
```

```
/usr/omni/bin/backint -f inquiry -u <ORACLE_SID> -i  
<input_file> (other UNIX systems)
```

where the specified `<input_file>` is queried.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

Backint anticipates a list of files of the following format:

```
<backup_ID_1> <pathName_1> [<targetDirectory_1>]  
<backup_ID_2> <pathName_2> [<targetDirectory_2>]  
<backup_ID_3> <pathName_3> [<targetDirectory_3>]
```

To retrieve the `<backup_ID>` numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u <ORACLE_SID>
```

or, alternatively, you can just specify `#NULL` as `<backup_ID_1>` in the `<input_file>`. In this case, the latest backup session for the file is used for the restore.

3. Verify the restore using the Data Protector User Interface

This test is possible if the objects have been backed up by `backint`.

Refer to “Restoring an SAP R/3 Database” on page 174.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

4. Simulate a Restore Session

Once you know the information about the object to be restored, you can simulate a restore using the Data Protector `testbar2` utility.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the SAP R/3 Database Server.

Check the `/etc/opt/omni/cell/cell_server` (HP-UX and Solaris systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the testbar2 utility:

```
/opt/omni/bin/utilns/testbar2 -type:SAP
  -apname:<ORACLE_SID>
  -perform:restore
  -object:<object_name>
  -version:<object_version>
  -bar:<backup_specification_name> (HP-UX and Solaris
systems) or
```

```
/opt/omni/bin/utilns/testbar2 -type:SAP
  -apname:<ORACLE_SID>
  -perform:restore
  -object:<object_name>
  -version:<object_version>
  -bar:<backup_specification_name> (other UNIX systems)
```

You should see only NORMAL messages displayed on your screen, otherwise examine the errors reported by the testbar2 utility by clicking the Details button in the Data Protector Monitor context.

5. Verify the restore using backint

Run the following command:

```
/opt/omni/lbin/backint -f restore -u <ORACLE_SID> -i
<input_file> (HP-UX and Solaris systems)
/usr/omni/bin/backint -f restore -u <ORACLE_SID> -i
<input_file> (other UNIX systems)
```

where the contents of the <input_file> will be restored.

If this fails, check if the session was performed successfully and if the restore was started under the appropriate user account.

Backint anticipates a list of files in the following format:

```
<backup_ID_1> <pathName_1> [<targetDirectory_1>]
<backup_ID_2> <pathName_2> [<targetDirectory_2>]
<backup_ID_3> <pathName_3> [<targetDirectory_3>]
```

Troubleshooting

To retrieve the *<backup_ID>* numbers, enter the following command:

```
echo #NULL #NULL | backint -f inquiry -u <ORACLE_SID>
```

Example of SAP R/3 Database Restore

This section describes some examples of how you can restore an SAP R/3 database. The following examples of restore are given:

- Full database restore and recovery
- Restore of lost files
- Archive log restore

IMPORTANT

The restore of an SAP R/3 database can be performed using SAP R/3 utilities, which are not a part of Data Protector. This section only describes *examples* of how you can perform a restore using the BRRESTORE utility from SAPDBA. The examples provided do not apply to all situations, where the restore is needed. For additional information on how you can restore an SAP R/3 database using the BRRESTORE utility, refer to the SAP R/3 documentation.

Preparing the SAP R/3 Database for Restore

If you are performing a full database restore, you need to know how the backup was performed, whether you have used the Oracle8 RMAN channels or only BRBACKUP tools. If you used RMAN, then you need to use Oracle8 `svrmgr1` and RMAN commands to perform the restore. If you have used BRBACKUP utility, then you need to use SAPDBA to perform the restore.

If you are performing a partial restore, you can use BRRESTORE tools that come with the SAP R/3 BRBACKUP utility.

The following environment variables must be set before performing the restore:

- ORACLE_SID: system ID of the database instance

Example: P01

SAPSID refers to the name of the SAP R/3 system, while the DBSID refers to the name of the database instance. When a single instance is installed, SAPSID and DBSID are the same.

- ORACLE_HOME: home directory of the Oracle software
Default location is: /opt/oracle/<DBSID>
- SAPDATA_HOME: home directory of the database files
Default location is: /opt/oracle/<DBSID>

IMPORTANT

The environment variables ORACLE_SID, ORACLE_HOME and SAPDATA_HOME must always be set.

The following environment variables must only be set if the corresponding paths are different from the default locations:

- SAPARCH: directory for the BRARCHIVE logs
Default location: <SAPDATA_HOME>/saparch
- SAPBACKUP: directory for the BRBACKUP logs
Default location: <SAPDATA_HOME>/sapbackup
- SAPCHECK: directory for the sapdba -check/analyze logs
Default location: <SAPDATA_HOME>/sapcheck
- SAPREORG: directory for all other SAPDBA logs, as well as shell and SQL scripts. It is also the standard directory for export and unload dump files, if the parameter exireo_dumpdir in the profile init<DBSID>.dba is not set.
Default location: <SAPDATA_HOME>/sapreorg
- SAPTRACE: directory for Oracle8 trace files and the alert file
Default location: <SAPDATA_HOME>/saptrace
- SAPDATA1: directory of the database data files
Default location: <SAPDATA_HOME>/sapdata1
Syntax for SAPDATA<n> is: n=1, . . . , 99. The environment variables SAPDATA<n> must only be defined if directories are on a location other than the default.
- TWO_TASK: identification of a remote database system
This environment variable must not be set.

Other optional environment variables that can be set:

- LINES: definition of the screen height
- COLUMNS: definition of the screen width
- SAPDBA_DEBUG: setting the trace function for error analysis

Examples of SAP R/3 Database Restore

Full Database Restore and Recovery

To perform a full database restore and recovery, follow the steps below:

1. Login to the SAPDBA utility. In the SAPDBA select the **m** to display User and Security option. Select the **Expert** mode and enter the Expert's password.

Figure 2-15

Starting the SAPDBA in Expert Mode

```
SAPDBA V4.6A - SAP Database Administration

ORACLE version: 9.0.5.0.0
ORACLE_SID      : ABA
ORACLE_HOME     : /app/oracle805/product
DATABASE        : open
SAPR3           : not connected

a - Startup/Shutdown instance   h - Backup database
b - Instance information        i - Backup offline redo logs
c - Tablespace administration   j - Restore/Recovery
d - Reorganization             k - DB check/verification
e - Export/import              l - Show/Cleanup
f - Archive mode               m - User and Security
g - Additional functions        n - SAP Online Help

q - Quit

Please select ==> m

User and Security

a - Expert mode
b - User information
c - Role information
d - Restricted mode
p - Change password

q - Return

Please select ==> a
```

2. When the menu appears, select the Restore/Recovery option.

Figure 2-16

Selecting the Restore/Recovery Option

```
SAPDBA V4.6A - SAP Database Administration

ORACLE version: 8.0.5.0.0
ORACLE_SID      : ABA
ORACLE_HOME     : /app/oracle805/product
DATABASE        : open
SAPR3           : not connected

a - Startup/Shutdown instance   h - Backup database
b - Instance information         i - Backup offline redo logs
c - Tablespace administration   j - Restore/Recovery
d - Reorganization              k - DB check/verification
e - Export/import               l - Show/Cleanup
f - Archive mode                m - User and Security
g - Additional functions         n - SAP Online Help

q - Quit

Please select ==> j
```

3. When the new menu appears, you can select between different types of restore. Select Full restore and recovery option. SAPDBA will check if your database is up and running.

Figure 2-17

Selecting Full Restore and Recovery

```
Restore / Recovery (2001-10-09)

a - Partial restore and complete recovery (Check and repair,
redo logs and control files are prerequisites)
b - Full restore and recovery
(excl. redo logs, control files incl. if required)
c - Reset database
(incl. redo logs and control files)

d - Restore one tablespace
e - Restore individual file(s)

h - Help
q - Return

Please select ==> b
```

4. After the SAPDBA checks the status of the database, a new window displaying the results appears. Specify the Select a backup of type option to select the backup version you want to use to perform the restore.

Figure 2-18

Selecting the backup type and version for restore

```
a0y0C _____
                                     Full Restore and Recovery (2001-10-09)
-----
DATABASE STATE   : open
RESTORE / RECOVER: disallowed (see status)

                                     Current setting
A - Select a backup of type
   full online/offline (level 0) or
   whole online/offline (all)         <not selected>

c - Recover until                       now
d - Show status
e - Options
g - Restart restore/recover operation

S - Start restore and recover
q - Return

Please select ==> █
```

5. Afterwards, enter the full pathname name for the backup tool parameter file.
6. Select the Start restore and recover option to start the restore session.

Figure 2-19

Starting the Restore Session

```
a0y0C _____
                                     Full Restore and Recovery (2001-10-09)
-----
DATABASE STATE   : open
RESTORE / RECOVER: allowed

                                     Current setting
A - Select a backup of type
   full online/offline (level 0) or
   whole online/offline (all)         bdgjpvla.anf
                                     2001-10-08 14.51.06
c - Recover until                       now
d - Show status
e - Options
f - Show/Delete datafiles younger than 2001-10-08 14.51.06
g - Restart restore/recover operation

S - Start restore and recover
q - Return

Please select ==> S█
```

7. Select the Return to restore procedure and continue, if you want to specify or modify the restore parameters.

Figure 2-20

Selecting Return to restore process and continue option

```
a0y0C _____
                                     Specify Restore Parameters for Backup Files
-----
Selected  bdgjp1a.anf 2001-10-08 14.51.06

a - BRBACKUP profile                Current value
b - Use (choose) former restores   initABA.sap
c - Clear list of former restores   rdgjwtty.rsb
g - Backup utility parameter file  util_file
i - Language                        English

q - Return to restore process and continue
r - Cancel restore process

Please select ==> q
```

IMPORTANT

After the restore and recovery are performed, the command ALTER DATABASE OPEN RESETLOGS is always executed. For security reasons, perform a backup of the database, before you open it.

After you have opened the database, the current LOG SEQUENCE NUMBER = 1 and its operations overwrite the old redo log files. Back up the offline redo log files before the function is executed.

Partial Restore

To perform a partial restore and recovery, you need to determine whether you need to restore a backup file or an archive redo log. The task of the SAPDBA recovery function is to fix certain media and user errors. When such errors occur, they usually involve the loss of database files, which contain many various types of objects: Oracle8 Dictionary segments, temporary segments, rollback segments, or user segments (tables and indexes).

SAPDBA utility supports restoring the database after the loss of the following files:

- SAP tablespaces data file (PSAP<name>D/I)
- System tablespace files (SYSTEM)
- Rollback tablespace files (PSAPROLL)
- Temporary tablespace files (PSAPTEMP)

The menu option `Check (and repair) database only` enables the recovery of the database up to the present time.

Restoring lost files

To restore the lost files, follow the steps below:

1. Define the time period within which you want SAPDBA to search for the backup files. The default value is 30 days. Then select the `Start finding backup files` menu option. SAPDBA utility uses the BRBACKUP log files to find the backup files.

If the SAPDBA utility finds backup files, the necessary log sequence number is determined by SAPDBA as follows: SAPDBA searches for the most recent BRBACKUP file for each lost file and then selects the lowest of the respective log sequence numbers.

2. Select the `Show the list of damaged files` to determine the files that need to be restored.

The SAPDBA utility lists all the lost files and their backup files. Each file shown in the list contains one of the following comments:

- Backup file: `<name> on <tape/disk>`

`Backed up by <name of the external backup program>`

This means that the file was backed up using the specific program. This comment appears when the parameter `backup_util_name` of the profile `init<DBSID>.dba` contains the name of the external backup program. Otherwise, the comment is displayed as, for example: `ext. backup utility`.

- No restore of a backup file required

This means that the existing file can be used.

- No backup file found

This means that no backup was found for this file in the specified period of time.

3. Select the `Show the list of backup files` option to specify the lost files for which you would like to see the available backup files. Each file that has been lost can have several backup files.
4. Select the `Select a backup file for restore` if you would like to change the proposed backup file, that should be restored. The file that is selected for the restore is flagged with `Selected for restore`.

5. Select the `Select a BRBACKUP run for restore` if you want to change the newest found backup file for each individual file from which the requested files can be restored. You can change this setting, for example, if all the files for restore were backed up in the same backup session and you want to specify only that backup session. The following information is listed:
 - Sequential number of the backup file found
 - Coded timestamp, date and time of the backup
 - The medium on which the backup was performed
 - The number of files found in this backup which are to be restored
6. Select the `Return` option to continue with the recovery process.
The lost files are restored using the SAP utility `BRRESTORE`.
7. Select the `Start restore of backup files`.
SAPDBA checks if the files that are to be restored are still available. If these files are still available, an error message is displayed. Confirm that SAPDBA may overwrite these files. If you do not allow SAPDBA to overwrite these files, the restore procedure is terminated at this point.

SAPDBA checks if there is a backup file for each data file that was lost. If a backup file is missing, the restore procedure is terminated at this point.

SAPDBA displays the restore parameters. The SAP utility `BRRESTORE` is started in order to restore the files.

Archive Log Restore

To restore the archive log files, follow the steps below:

1. Select the `Restore archive files` option.
Archive log files are restored using the SAP `BRRESTORE` utility. If SAPDBA determines that the archiving directory `<Oracle_home>/saparch` does not have enough space to restore all the necessary redo log files, the redo log files that have already been used will be deleted and the next required redo logs are restored during the subsequent recovery.
2. Select the `Start restore of archive files` option.

This option is mandatory when the recovery requires offline redo log files that are no longer in the archiving directory. The recovery cannot be started until the necessary archived redo logs are restored.

SAPDBA displays the following information on the screen:

- The log sequence number of the first archived file to be restored.
- The archived files that were found.
- The maximum size of the archived redo log files.
- The configured restore parameters which you can change using the `Specify restore parameters` option.

The SAP `BRRESTORE` utility restores the required files. If the redo logs are still available on the disk, they do not have to be restored.

3. Select `Return` to continue with the recovery process.

Integrating SAP R/3 and Data Protector
Example of SAP R/3 Database Restore

In This Chapter

This chapter explains how to install, configure, and use the Data Protector IBM DB2 UDB integration.

The chapter is organized into the following sections:

“Overview” on page 209

“Prerequisites and Limitations” on page 211

“Integration Concept” on page 213

“Installing the DB2 Integration” on page 219

“Configuring the Integration” on page 220

“Backing Up a DB2 Database” on page 230

“Restoring a DB2 Database” on page 236

“Monitoring a DB2 Backup and Restore” on page 248

“Troubleshooting” on page 250

Overview

The Data Protector integration with IBM DB2 Universal Database (UDB) Server (hereafter referred to as DB2 Server) allows you to perform online and offline backups of the DB2 database objects.

Refer to the *HP OpenView Storage Data Protector Software Release Notes* for information about platforms and devices supported by the Data Protector DB2 integration.

Backup Types

You can perform the following types of backup using the Data Protector DB2 integration:

- Backup of one or several databases
- Backup of one or several table spaces
- Backup of several table spaces from different databases
- Incremental database/table space backup
- Delta database/table space backup

Each backup type may be performed both online and offline.

Restore Types

The Data Protector DB2 integration supports the following restore types:

- Offline database restore and rollforward operations.
- Offline and online table space restore and rollforward operations.
- Automatic restore from incremental or delta backups. In this case, DB2 7.2 with Fixpack 7 must be installed.
- Restore to a new database (for database-level backups only).
- Restore to a different system.

Advantages

Integrating Data Protector with DB2 Server offers several advantages over using the internal DB2 backup and restore functionality:

- Central Management for all backup operations:

The administrator can manage backup operations from a central point.

Overview

- **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 to organize and manage devices in media pools.

- **Scheduling:**

Data Protector has a scheduler that allows the DB2 administrator to automate backups to run periodically. Using the Data Protector Scheduler, you can configure the backups to run unattended, at specified times, if the devices and media are set properly.

- **Device Support:**

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

- **Reporting:**

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

- **Monitoring:**

Data Protector has a feature that allows the DB2 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 internal database (IDB), which provides the administrator with a history of activities that can be queried later.

Prerequisites and Limitations

This is a list of prerequisites and limitations for the Data Protector DB2 integration:

Prerequisites

- It is assumed that you are familiar with the DB2 database administration and basic Data Protector functionality.
- Before you begin, make sure that you have correctly installed and configured DB2 Server and Data Protector systems. Refer to the:
 - ✓ *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 general information on how to configure and run backups.
 - ✓ *DB2 Administration Guide*
 - ✓ *DB2 Server Books Online* for online information on DB2 Server.

For an up-to-date list of supported versions, platforms, devices, and other information, refer to the:

- ✓ *HP OpenView Storage Data Protector Software Release Notes* or
- ✓ http://www.openview.hp.com/products/datapro/spec_0001.html
- To perform an online backup of DB2 objects you need to set the DB2 `logretain` and `userexit` parameters to ON. The backup will fail if the database does not have these parameters set correctly.
- Fixpack 7 is required for Data Protector integration with DB2 7.2. After you have installed the fixpack, update the DB2 instances by running the DB2 command `db2iupdt`.

Limitations

The following is not supported:

- Backup or restore using the DB2 Command-Line Processor or the DB2 Control Center.
- Backup of a partitioned database.
- Table or datafile backup and restore.

Prerequisites and Limitations

- Restore to a new database is supported for database-level backups only.
- DB2 temporary table spaces can only be backed up during the full database backup.
- Rollforward recovery of system catalog can only be performed if no other table spaces from the same DB2 database are being restored from the same session.

Integration Concept

DB2 Components The DB2 part of the integration provides an open interface for backing up and restoring DB2 objects. This allows you to use Data Protector as a data transferring module for database backup and restore operations.

Data Protector Components The Data Protector integration software consists of the following components:

- The `db2bar` module, which is installed on the DB2 Server system. It controls the activities between DB2 Server and the Data Protector backup and restore processes.
- The `libob2db2` component, which is the actual data transferring module, called by DB2 Server.
- The `db2arch` program, which performs backup and restore of the DB2 log files. It is called automatically if the `DB2 logretain` and `userexit` parameters are set to ON.
- The `util_db2` utility, which is used by Data Protector to configure a DB2 instance and check the instance configuration.
- The `testbar2` utility, which checks the Data Protector part of the integration.

From the perspective of DB2 Server, Data Protector is seen as media management software. On the other hand, DB2 Server is a Data Protector client from the Data Protector Cell Manager's point of view.

The concept of the Data Protector DB2 integration is described in Figure 3-2 on page 218.

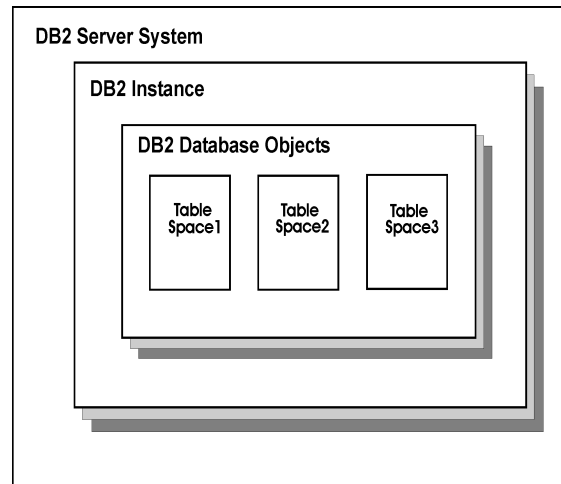
DB2 Database Concepts The following DB2 database concepts are important from the backup perspective:

- **Instance**, which controls the operations performed on data and manages system resources assigned to it. Each instance has its own *databases*, which other instances cannot access.
- **Database**, which presents data as a collection of tables, where each table consists of a defined number of columns and any number of rows. A database is organized into the parts called *table spaces*.

- **Table spaces**, which are the places for storing tables. A table space can be spread over to one or more physical storage devices.

Refer to Figure 3-1 on page 214 for the DB2 database structure overview.

Figure 3-1 DB2 Database Structure



Recovery Methods Two recovery methods, used to restore DB2 database objects, are version and rollforward recovery. They are closely connected with the DB2 database logs, which keep records of database changes. If a database needs to be restored to a point beyond the last full offline backup, logs are required to roll the data forward to the point of failure.

- **Version recovery** is the restore of a previous version of the database using an image created during a backup operation. A database restore operation rebuilds the entire database using a database backup performed earlier. This allows you to restore a database to the state identical to the one at the time that the backup was made.

A “ring” of online log files is used to provide recovery from transaction failures and system crashes. This behavior is called **circular logging**. With this type of logging, only full offline backups of the database are possible. Circular logging does not allow you to roll forward a database through prior transactions from the last full

backup. Recovery from media failures and disasters is done by restoring from a full offline backup. Every unit of work from the time of a full backup to the time of failure is lost.

- **Rollforward recovery** is the restore of a database or a table space to its state at a specified point in time. A log is closed and becomes archived when changes in the active log are no longer needed for normal processing. This behavior is called **archived logging**. At the end of the restore operation, the database is in the rollforward pending state that allows the rollforward recovery to take place.

The archived logs can be online, meaning that a log is stored in the database log path directory, and offline, meaning that a log is no longer found in the database log path directory.

The archived logs are backed up and restored using the DB2 User Exit program, which is called whenever a new offline redo log appears, but no sooner than the previous backup or restore is completed.

There are two types of rollforward recovery:

- ✓ Database rollforward recovery. If this type of rollforward recovery is used, the transactions, recorded in database logs, are applied following the database restore operation. The database logs record all changes made to the database. This method completes the recovery of the database to its state at a particular point in time, or to its state immediately before the failure.
- ✓ Table space restore and rollforward. If the database is enabled for rollforward recovery, you may also back up, restore and roll forward table spaces. To perform a table space restore and rollforward, you must have the following:
 - Backup image of either the entire database (that is, all table spaces), or one or more individual table spaces.
 - The log records affecting the table spaces that are to be recovered.

You can roll forward through the logs to one of the following two points:

- the end of logs
- a particular point in time (point-in-time recovery).

NOTE

All DB2 timestamps in messages during rollforward recovery are by DB2 design in Universal Coordinated Time (UCT) format.

Refer to “Restore Options” on page 241 for information on the rollforward restore options.

Backup Flow

The basic backup unit is a table space. It means that only table spaces or databases can be selected for backup.

A backup session is started by the Data Protector Backup Session Manager (BSM). The BSM invokes `db2bar`, which, using the DB2 API, starts the backup applying the backup options defined in the backup specification. After that, DB2 Server calls the sequence of functions from the `libob2db2` shared library, which performs the backup. At this point, the BSM starts Media Agents, which write the data to the backup devices.

If case of an online backup, DB2 Server closes the log files, and then calls the `db2arch` module, which is the User Exit program responsible for backing up log files. After a successful backup of a log file, the file is automatically deleted by DB2 Server.

Messages from the backup session are sent to the BSM, which writes them and the information regarding the respective session to the IDB.

Backup Types

Three types of backup supported by the Data Protector DB2 integration are **full**, **incremental**, and **delta**.

A full backup is a backup of all selected database objects regardless of whether they have been changed after the last backup was made. Some data, such as database configuration, history file, etc, which is important for restore, is included into the full backup automatically. An incremental backup selects all the changes made to the database after the last full backup. A delta backup is a backup containing all the changes made to the database from the last backup of any type.

Restore Flow

A restore session is started by the Data Protector Restore Session Manager (RSM). The RSM invokes the `db2bar` utility, which starts the restore using the DB2 API. After that, DB2 Server calls the sequence of functions from the `libob2db2` shared library, which performs the restore. At this point, the RSM starts Media Agents, which read the data

from the backup devices and send it to the DB2 Server through the processes `libob2db2` library. The DB2 Server processes write the data to disks.

In case you are performing a recovery from an incremental or delta backup, Data Protector will first restore the selected backup session in order to get information about DB2 backup chain history. Then it will restore the last full backup (of the selected incremental/delta backup) and finally the last incremental backup and/or all subsequent delta backups.

In case of a restore from an online backup, the rollforward operation is performed. DB2 Server calls the `db2arch` module, which is the User Exit program responsible for restoring log files, to restore the logs, needed for rollforward, one by one.

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

The architecture of the Data Protector DB2 integration is presented in Figure 3-2 on page 218.

Figure 3-2 Data Protector DB2 Integration Concept




Table 3-1

Legend:

SM	The Data Protector Session Manager, which is the Data Protector Backup Session Manager during a backup session, and the Data Protector Restore Session Manager during a restore session.
MA	The Data Protector Media Agent, which reads and writes data from and to media devices.
IDB	The Data Protector internal database where all the information about Data Protector sessions, including session messages, objects, data, used devices and media is written.

Installing the DB2 Integration

Installation

It is assumed that your DB2 Server is up and running.

Install the Data Protector DB2 Integration software on your DB2 Server either locally, from the CD-ROM, or remotely, by using the Data Protector GUI.

You need to install the following Data Protector software components:

- DB2 Integration
- Disk Agent
- Media Agent (if you have devices connected to the system)

It is recommended that you install:

- User Interface

Install this component to have access to the Data Protector GUI and the Data Protector CLI on this system.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details about the installation (Cell Manager on HP-UX systems, Client on AIX systems).

NOTE

If you perform a local installation of the Data Protector software on the DB2 Server system, select the DB2 Integration software component during the setup procedure. Other required components are selected by default.

Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

By now, you should have installed the Data Protector DB2 integration software on your DB2 Server system. The DB2 Server system has become a Data Protector client. At this point, you are ready to proceed to configuration described in “Configuring the Integration” on page 220.

Configuring the Integration

It is assumed that the installation of the Data Protector software components on the DB2 Server system was successful.

Prerequisite

The `logretain` and `userexit` DB2 database parameters must be set to `ON` before you start a configuration procedure. Otherwise, online backups will fail. For the information on how to set these database parameters, refer to the *DB2 Administration Guide*.

Configuration Overview

The following list gives an overview of the global tasks for configuring the DB2 integration:

1. “Configuring a DB2 User” on page 220.
2. “Configuring a DB2 Instance” on page 220
3. “Configuring a DB2 Backup” on page 222.

Configuring a DB2 User

In order to perform backup and restore-related operations, the DB2 user must be registered in the operating system and have the `SYSADM`, `SYSCTRL`, or `SYSMAINT` authorities. The root user has to be a member of the Data Protector admin group to be able to browse the DB2 database structure.

Configuring a DB2 Instance

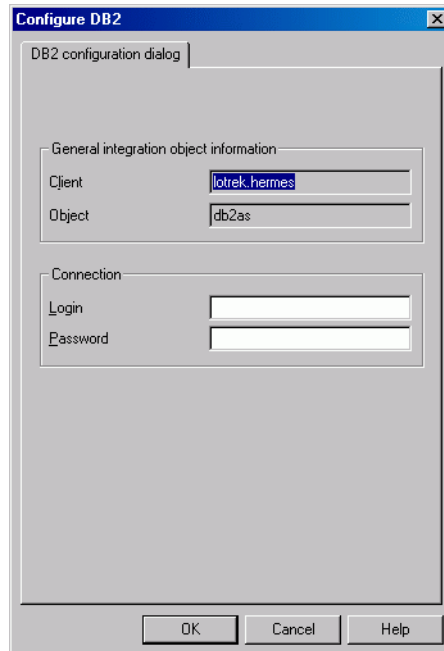
The parameters that need to be specified during the configuration of a DB2 instance are the username and the password of the DB2 user. These parameters will be also used for establishing the connection to the DB2 Server system if the user starts non-backup and non-restore-related operations, such as listing of objects for backup.

The configuration is performed during the creation of a new backup specification, or by modifying an existing backup specification. For a step-by-step procedure on creating a backup specification, refer to “Creating a Backup Specification” on page 222.

The procedure below describes the configuration of a DB2 instance using an existing 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 DB2 Integration. Click an existing backup specification.
3. Right-click the name of the DB2 Integration listed in the Source property page, and then select Configure from the pop-up window.
4. In the Configure DB2 window, specify the username and the password of the operating system user.

Figure 3-3 DB2 Configuration



Click OK to confirm the configuration.

If properly configured, the DB2 user is allowed to back up or restore the DB2 database objects. In order to start a backup of a DB2 database object using Data Protector, the user must also be the owner of the Data Protector backup specification.

Refer to the DB2 documentation for further information on different types of connections, the roles and authorities of DB2 database administrators and security issues that must be considered.

Configuring a DB2 Backup

To configure a DB2 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.

The Data Protector backup specification is stored on the Cell Manager system and contains instructions on how to perform a backup using Data Protector.

Once the backup specification is created and saved, it can be scheduled so that unattended backups can be performed.

Creating a Backup Specification

To create a backup specification for backing up the DB2 database objects, 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 DB2 Integration and then select Add Backup. The Create New Backup dialog box is displayed.
4. Select one of the templates described below:

Database_Backup

Used for online and offline backups of DB2 database objects only. You cannot perform an archive log backup using this template.

Archived_Logs_Backup

Used for backing up archived logs only. You cannot perform a database objects backup using this template. The backup

specification created using this backup template can only be saved, and not started or scheduled. It will be used any time the User Exit program starts the backup of archived logs.

You should not create a new archived logs backup specification if an older archived logs backup specification already exists. You must erase the old one first.

IMPORTANT

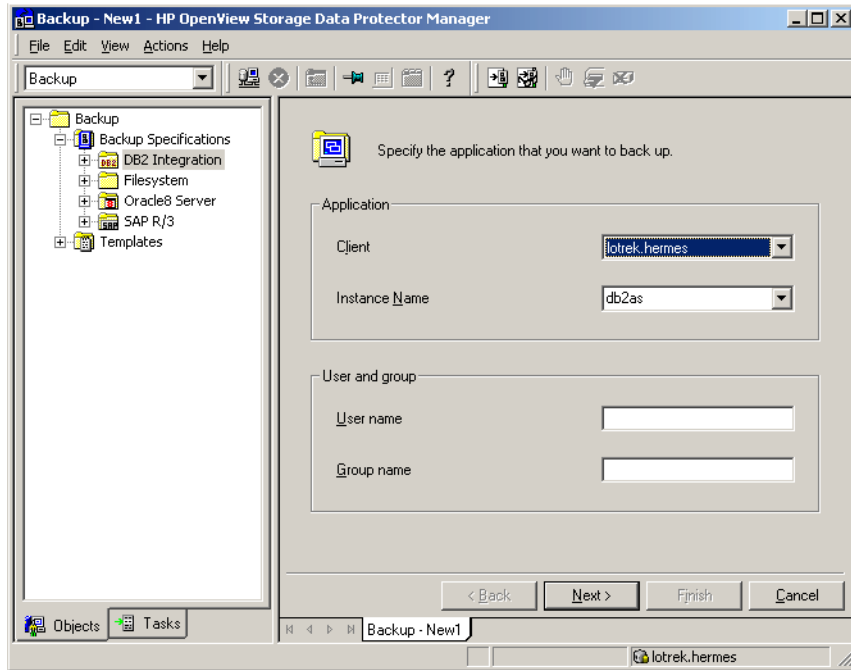
Whenever an online database backup is started, DB2 also tries to back up archive logs, therefore you must create an archive logs backup specification prior to running online db2 backup. Since the backup of archived logs is started automatically at the time a new offline archived log appears, you must always have a device that will be used only for the backup specification created using the `Archived_Logs_Backup` template.

Click OK.

5. Select the client on which DB2 Server is running, and the application database. Data Protector lists all the configured DB2 instances located on this system. If the instance has not been configured yet, enter the instance name. Then enter the username and the password for the user, under whose account the backup session will be started.

If the application database you have entered had not been configured yet, the configuration window appears. Refer to “Configuring a DB2 Instance” on page 220.

Figure 3-4 **Selecting Client System and Application Database**

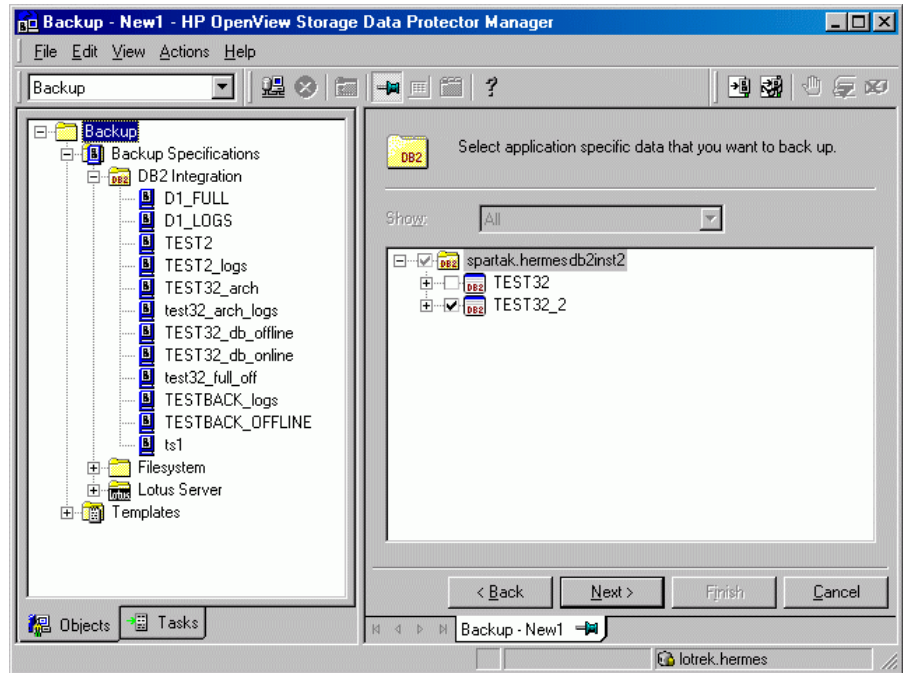


6. Click Next.
7. In the next step of the wizard, select the database objects you want to back up.

To perform an offline backup of the object, right-click the object and select Properties. The Properties window appears. In this window, select the Offline backup option.

If you used the Archive_Log_Backup template, the Archived Logs item is selected by default and you cannot deselect it.

Figure 3-5 **Selecting Backup Objects**




8. Follow the wizard to define options, devices, and schedule.

Refer to Data Protector online Help and the *HP OpenView Storage Data Protector Administrator's Guide* for a description of the backup options common to all backup objects.

Refer to “DB2 Specific Backup Options” on page 226 for details about the DB2 backup options.

9. Once you have defined all the backup options, name and save your DB2 backup specification. It is recommended that you save all the DB2 backup specifications in the DB2 Integration group.

Figure 3-6 Saving the Backup Specification



After the backup specification is saved, it can be started either from the Data Protector GUI or the Data Protector CLI, or can be scheduled to run automatically using the Data Protector Scheduler. Refer to “Backing Up a DB2 Database” on page 230 for information on how to perform a backup using the Data Protector GUI or the Data Protector CLI and on how to schedule a backup specification.

You can examine the newly created and saved backup specification in the Backup context. The backup specification itself is stored in the `<Data_Protector_home>\config\barlists\db2\<backup_specification_name>` file on Windows Cell Manager systems and in the `/etc/opt/omni/barlists/db2/<backup_specification_name>` file on UNIX Cell Manager systems.

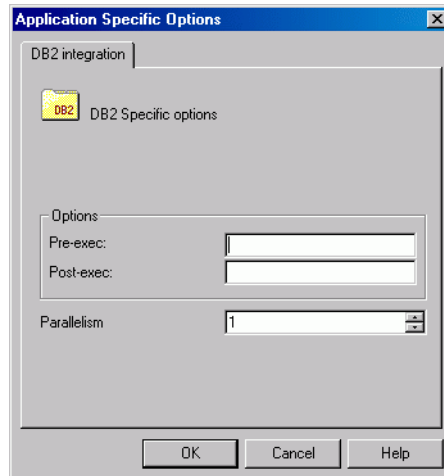
It is recommended that you test the backup specification by clicking the Start Preview button. This is an interactive test that does not back up any data. However, as a result of this test, the file `<Data_Protector_home>\tmp\<Backup_Specification_Name>_TEST_FILE` is created on the DB2 Server system. It should be deleted after the test. Refer to “Testing the Integration” on page 228 for a step-by-step procedure.

You can start an interactive backup that includes data transfer by clicking the Start Backup button.

DB2 Specific Backup Options

The DB2 specific backup options are specified using the Data Protector GUI by clicking the Advanced button next to Application Specific Options.

Figure 3-7 Backup Options



The following are the DB2 specific backup options:

Pre-exec Specifies a command with arguments or a script that will be started on DB2 Server before the backup. This command/script is started by the Data Protector db2bar module and must reside in the /opt/omni/lbin directory on HP-UX systems or in the /usr/omni/bin directory on AIX systems. Only the filename, relative to the directory named above, must be provided in the backup specification. For more information on pre-exec commands, refer to “Pre- and Post-Exec Commands” in the *HP OpenView Storage Data Protector Administrator’s Guide*.

Post-exec Specifies a command with arguments or a script that will be started on DB2 Server after the backup. This command/script is started by the Data Protector db2bar module and must reside in the /opt/omni/lbin directory on HP-UX systems or in the /usr/omni/bin directory on AIX systems. Only the filename, relative to the directory named above, must be provided in the backup specification. For more information on post-exec commands, refer to

“Pre- and Post-Exec Commands” in the *HP OpenView Storage Data Protector Administrator’s Guide*.

Choose the `Parallelism` value to set how many streams of data will be backed up. The default value is 1. The DB2 parallel backup must be made on different devices with concurrency set to 1, otherwise restore even from successful backup can be impossible.

Testing the Integration


Once you have created and saved a backup specification, you should test it before running a backup.

Testing Using the Data Protector GUI

Testing Procedure The testing procedure consists of checking the Data Protector part of the integration to ensure the communication within Data Protector is established and the data transfer works properly. Proceed as follows to test the integration:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, then Backup Specifications, DB2 Integration and 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. See online Help for a description of these options.

Figure 3-8 Previewing a Backup



Testing Using the Data Protector CLI

To test a backup specification, run the `omnib` command with the `-test_bar` option.

On HP-UX systems, execute the following command:

```
/opt/omni/sbin/omnib -db2_list <backup_specification_name> \  
-test_bar
```

On AIX systems, execute the command given below:

```
/usr/omni/bin/omnib -db2_list <backup_specification_name> \  
-test_bar
```

What Happens?

The session messages are displayed on the screen during the command execution.

The `db2bar` program is started, which then starts the Data Protector `testbar2` command. This command checks the following:

- if the communication within Data Protector works properly.
- if the syntax of the DB2 Integration backup specification is correct.
- if the devices are correctly specified.
- if the required media reside in the devices.

After that, the DB2 part of the preview is started, which checks if all the backup objects are present and are in a correct state for a backup.

Backing Up a DB2 Database

There are two modes for backing up DB2 database objects: the online and the offline database backups.

During an online backup, the database is open and available for the other applications. During an offline backup, the database is closed and unavailable for use.

IMPORTANT

Before you perform an online backup of DB2 objects, set the DB2 `logretain` and `userexit` parameters to `ON`. The backup of archive log files will not be possible if the database does not have these parameters set correctly.

To run a backup, use any of the following methods:

- Schedule a backup of an existing DB2 Integration backup specification using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or the Data Protector CLI.

To back up archived logs, you have to create a backup specification using the `Archive_Log_Backup` template. Refer to “Creating a Backup Specification” on page 222. Note that you can only save this type of backup specification, but not execute or schedule it. The backup of archived logs will be started automatically as soon as a new offline archived log file appears.

To back up DB2 temporary table spaces, you have to perform full database backup. Individual restore of temporary table spaces is possible only from full database backup.

To enable incremental or incremental delta online backups you must first enable modification tracking. To do so, you have to perform the following steps:

1. Run the following command to activate modification tracking:

```
db2 update db cfg for <DatabaseName> USING TRACKMOD ON
```
2. Restart the database.

3. Perform a full database backup.

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 DB2 backup specification, follow the steps described in “Creating a Backup Specification” on page 222.

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 DB2 Integration.

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 3-9 on page 232.
6. Click OK to return to the Schedule property page.
7. Click Apply to save the changes.

NOTE

You cannot schedule a DB2 backup specification created using the Archive_Log_Backup template.

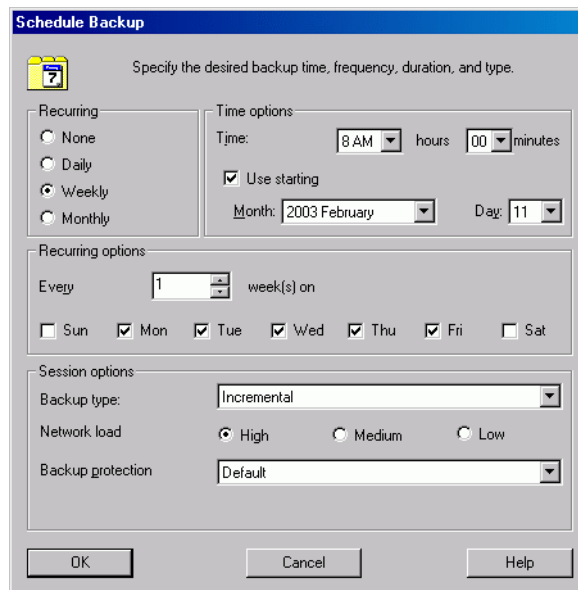
Scheduling Example

To schedule a backup specification so as to back up table spaces 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 Incremental backup type. Click OK.

See Figure 3-9 on page 232.

Figure 3-9 Scheduling the 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.

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

NOTE

When creating a DB2 backup specification, you access the Data Protector Scheduler through the Backup Wizard. See “Creating a Backup Specification” on page 222 for information about accessing the Backup Wizard.

Running an Interactive Backup Using the Data Protector GUI

An interactive backup can be run any time after the backup specification has been created and saved.

Backup Procedure To start an interactive backup of a DB2 backup object using the Data Protector GUI, perform the following steps:

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

Expand DB2 Integration. A list of backup specifications appears.

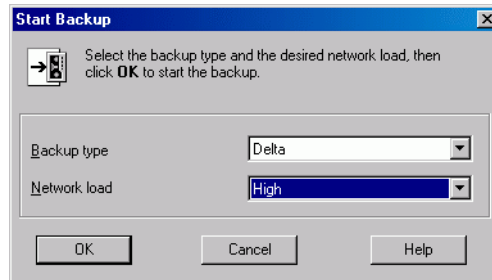
3. Right-click the backup specification you want to back up, and then select Start Backup from the pop-up menu.

The Start Backup dialog box appears.

Select the backup type and network load. Refer to online Help for a description of network load.

Click OK.

Figure 3-10 Starting an Interactive Backup



Messages appear in the Results Area as the backup session proceeds. Upon successful completion of the backup session, the `Session completed successfully` message and backup size are displayed. Backup size is by DB2 design calculated as size of full database backup + size of incremental/delta backup.

Running an Interactive Backup Using the Data Protector CLI

You can start an interactive backup from the Data Protector CLI. Switch to the `/opt/omni/lbin` directory on HP-UX systems or to the `/usr/omni/bin` directory on AIX systems, and run the following command:

```
omnib -db2_list <ListName> [-barmode <db2mode>]  
[<list_options>] [-preview]
```

The `<ListName>` parameter is the name of a backup specification.

The `<db2mode>` parameter specifies the type of the backup.

The `<list_options>` parameter sets the level of the protection, the level of the network traffic generated by the session, enables writing a CRC checksums, and disables monitoring of the backup session.

You can select among the following `<db2mode>`: {full | incr | delta}

You can select among the following `<list_options>`:

```
-protect {none | weeks n | days n | until date | permanent}  
-load {low | medium | high}  
-crc  
-no_monitor
```

Example

To start a full backup using an existing DB2 backup specification called TEST, and to set data protection to 10 weeks, execute the following command:

```
omnib -db2_list TEST -barmode full -protect weeks 10
```

Restoring a DB2 Database

You can restore a DB2 object using either the Data Protector GUI or the Data Protector CLI.

You can restore a previous version of the database using an image created during a backup operation, that is, to perform a version recovery. Also, you can restore the database(s)/table space(s) to their state at a specified point in time, that is, to perform a rollforward recovery. The rollforward operation ensures that all the changes made to the database during the online backup are captured and reapplied.

NOTE

Rollforward recovery of system catalog can only be performed if no other table spaces from the same DB2 database are being restored from the same session. Rollforward recovery of a system catalog can only be performed to the end of logs.

For more information on recovery methods, refer to “Integration Concept” on page 213.

NOTE

If you use the version recovery method, make sure that you perform full offline database backups on a regular basis.

IMPORTANT

The database restore and rollforward operations must always be performed offline. The table space (except for DB2 System Catalog) restore and rollforward can be performed online; though the table space itself is not available until the operation completes, you still can access data in other table spaces.

Restoring a DB2 Object Using the Data Protector GUI


Use the following procedure to restore a DB2 database object using the DP GUI:

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand DB2 Integration, and then the name of the client system from which the data you want to restore was backed up.
3. In the Source property page, browse for and select the backed up DB2 database objects you want to restore. The top-level elements are databases, and the second-level elements are table spaces.

The latest backed up version of each log file is restored automatically during the rollforward operation. However, in some cases, you may want to restore the version of a log file other than the latest. If you want to restore specific log files, select Archive logs option and then select log files you want to restore.

You may check the properties of each object by right-clicking the object name.

Figure 3-11 Restore Objects

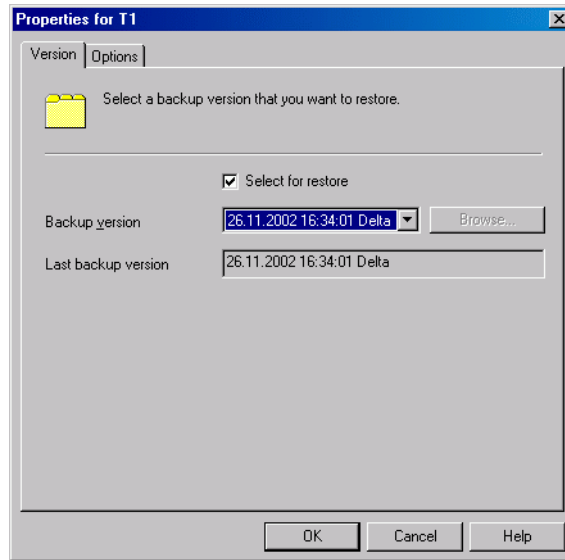


Under the Version tab in the Properties window, select the version from which you want to restore the data. The version is identified by the date and time of a backup and the backup type. If you are restoring the log files, only the date and time of a backup is displayed.

NOTE

By default, the latest version of the object is selected. If you want to restore some other version, select it from the Backup version drop-down list.

Figure 3-12 **Selecting a Version**



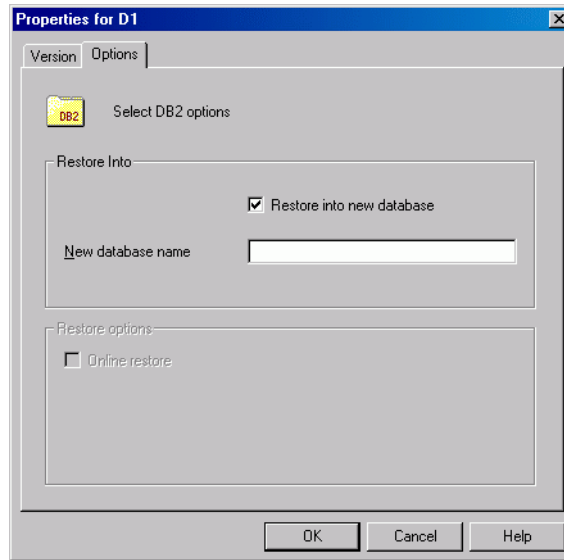
Under the Options tab in the Properties window, select the destination to which you want to restore the data. If the whole database was selected for restore, you can specify whether you want to restore it into a current or into a new database. To restore to a new database, enter the name of the new database. Refer to “Restoring into a New Database” on page 244

You can also specify whether you want to restore a table space offline or online.

NOTE


The default restore mode is online.

Figure 3-13 **Selecting a Destination**



4. Select the restore options from the `Options` property page. Refer to “Restore Options” on page 241.
5. In the `Devices` property page, the names of devices used for backup are displayed. If you want to restore from a device different from the one used for backup, select the device you want to change, and click `Change`. The list of all configured devices is displayed. From this list, select the device you want to use for restore, and then click `OK`.

Figure 3-14 **Selecting a Device**



Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on how to perform a restore using another device.

NOTE

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

6. Click `Restore DB2` to start the restore procedure.


The restore session messages are displayed in the *Results Area*.

Restore Options

The following restore options are specific to the Data Protector DB2 integration:

- User name** The name of the user who must have SYSADM, SYSCTRL, or SYSMAINT privilege to perform a restore to an existing database, and SYSADM or SYSCTRL privilege to perform a restore to a new database.
- User group** Group account of the above user.
- Password** Password of the above user.
- Restore to client** By default, the data is restored to the original backup client. However, you may restore your data to another client. In this case, enter the name of this client in the `Restore to client` text box. This option is valid for the whole database restore only. The instance with the name, specified in the `Restore to instance` text box, must be already created on the specified client. To perform a restore to a new client, you must have SYSADM or SYSCTRL privilege.
- Restore to instance** If you want to perform a restore to an instance other than the original, specify its name here. This instance must be created and configured prior to starting a restore session. Refer to “Restoring to Another Instance” on page 246. By default, the original instance name is displayed.
- Rollforward** This option is selected by default. To perform a version recovery, deselect the `Rollforward restore` option. However, if you deselected the `Rollforward` option and you are restoring from an online backup, the database goes into the rollforward pending state and is unavailable for use. To make the database available, start the rollforward operation using the DB2 Command-Line Processor or the DB2 Control Center.
- Rollforward to the end of logs** Rollforward is performed to the end of logs. This option is available only if the `Rollforward` option is selected.
- Rollforward to date** Rollforward is performed to a particular point in time. This option is available only if the `Rollforward` option is selected.

Figure 3-15 Restore Options



Restoring a DB2 Object Using the Data Protector CLI

You can also start a restore session from the Data Protector CLI. Switch to the `opt/omni/lbin` directory on HP-UX systems or to the `usr/omni/bin` directory on AIX systems, and run the following command:

```
omnir -db2
-barhost <ClientName>
[-destination <ClientName>]
-instance <InstName>
-dbname <DBName> [-session <SessionID>] [-newdbname
<NewDBName>] ...
[-frominstance <InstName>] ...
-tsname <TSName> [-session <SessionID>] [-offline] ...
-logfile <LogFileName> [-session <SessionID>] ...
[-rollforward [time: <YYYY-MM-DD.hh.mm.ss>]]
```

Restoring a DB2 Database

The `-barhost <ClientName>` parameter is the name of the DB2 Server system from which you are restoring; the `[-destination <ClientName>]` parameter is the name of the target DB2 Server.

The `<DBName>` parameter is the name of the DB2 database you want to restore (in case of a database restore); `<SessionID>` is the ID of the respective session. The `<NewDBName>` parameter is the name of a new database and it needs to be specified if you are restoring to a database (instance) other than the original.

The `-instance <InstName>` parameter is the name of the instance in which you want to restore the data.

The `-frominstance <InstName>` is the name of the instance from which you want to restore the data in case of restoring to a new instance.

The `<TSName>` parameter is the name of a table space you want to restore (in case of a table space restore); `<SessionID>` is the ID of the respective session.

The `<LogFileName>` parameter is the name of a log file you want to restore; `<SessionID>` is the ID of the respective session.

Example

To start an online restore of a DB2 database called `TEMP` on host `degas` and rollforward it till the 10th of January 2003, 9:15 a.m., execute the following command:

```
omnir -db2 -barhost degas -dbname TEMP -rollforward time:  
2003-01-10.09.15.00
```

Restoring into a New Database

Before you can start restoring a database into a new database, you must define the new table space containers for non-system table spaces. A container is the directory, file, or raw disk that stores table spaces.

To find all the containers needed for redirection, you can use two following commands:

- To list all table spaces in database:

```
db2 list tablespaces
```

- To find the name of the container for table space:

```
db2 list tablespace containers for <number of table space>
```

To redefine table space containers, you have to put additional options for redirection to the DB2 configuration file using `util_cmd` utility. Execute the following command:

```
util_cmd -putopt DB2 <instance name> "<old_container>" "<new  
container>" -sublist Redirection/<dbname>
```

This command has to be executed for every container.

After that you can start restoring to a new database using GUI. The new database will be in rollforward pending state. In the case of restore from offline backup, execute the following command:

```
db2 rollforward db <dbname> stop
```

In the case of restore from online backup you will have to restore log files using Data Protector GUI and then perform rollforward from CLI with the `OVERFLOW LOG PATH` option equal to a log path of the original database:

```
db2 rollforward db <dbname> to <time> OVERFLOW LOG \  
PATH "(<original database log path>)"
```

Example

The following steps show how to perform online restore of the `db2db_old` database, which resides in `db2inst` instance, into the `db2db_new` database. The log files for `db2db_old` database are in the `/db2_db/db2inst/NODE0000/SQL00003/SQLLOGDI` directory. In this example, one of the table spaces resides in `"/tmp/db2cont1"` container.

1. Define new container `"/tmp/db2cont2"` for table space using the `util_cmd` utility:

```
/util_cmd -putopt DB2 db2inst "/tmp/db2cont1" \  
"tmp/db2cont2" -sublist Redirection/db2db_old
```

2. Using Data Protector GUI start restoring the `db2db_old` database into the new `db2db_new` database, or using CLI, start restore using `omnir` command with the following parameters:

```
omnir -db2 -barhost <ClientName> -instance db2inst \  
-dbname db2db_old -newdbname db2db_new
```

3. Using Data Protector GUI restore all log files needed for rollforward.
4. Using Data Protector CLI perform rollforward to the end of logs executing the following command:

```
db2 rollforward db db2db_new to end of logs OVERFLOW \  
LOG PATH "(<original database log path>)"
```

Restoring to Another Instance

Before you can start restoring a database to an instance different from the original instance, you must configure the target instance by defining the new table space containers for non-system table spaces. A container is the directory, file, or raw disk that stores table spaces.

To find all the containers needed for redirection, you can use two following commands:

- To list all table spaces in database:

```
db2 list tablespaces
```

- To find the name of the container for table space:

```
db2 list tablespace containers for <number of table space>
```

To redefine table space containers, you have to put additional options for redirection to the DB2 configuration file using `util_cmd` utility. Execute the following command:

```
util_cmd -putopt DB2 <instance name> "<old_container>" "<new container>" -sublist Redirection/<dbname>
```

The *<instance name>* is the name of the target instance. This command has to be executed for every container.

After configuring the target instance you can start restoring to it using Data Protector GUI. The restored database will be in rollforward pending state. In the case of restore from offline backup, execute the following command:

```
db2 rollforward db <dbname> stop
```

In the case of restore from online backup you will have to restore log files using Data Protector GUI, and then perform rollforward from CLI with the `OVERFLOW LOG PATH` option equal to a log path of the restored database:

```
db2 rollforward db <dbname> to <time> OVERFLOW LOG \  
PATH '('<restored database log path>')
```

DB2 logs will be restored to the same directory where they resided at the time of backup. Set the write permissions of that directory to be able to restore the logs. After log files have been restored, check if the instance you are restoring into, has all the needed permissions for the log files you are restoring.

Example

The following steps show how to perform restore of the db2db database, which resides in inst1 instance, to the db2db database in the inst2 instance.

1. Define new container `"/tmp/db2cont2"` for table space using the `util_cmd` utility:

```
/util_cmd -putopt DB2 inst2 "/tmp/db2cont1" \  
"tmp/db2cont2" -sublist Redirection/db2db
```

2. Using Data Protector GUI start restoring the db2db database to the inst2 instance, or using CLI, start restore using `omnir` command with the following parameters:

```
omnir -db2 -barhost <ClientName> [-destination \  
<destination client name>] -instance inst2 - dbname \  
db2db -frominstance inst1
```

Use the `-destination` option if you want to make restore to another host only.

NOTE

If you are restoring a DB2 database to another instance on another host, use `db2 list tables for all` command to list tables.

Monitoring a DB2 Backup and Restore

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

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

Monitoring Current Sessions

Follow these steps to monitor a current session using the Data Protector GUI:


1. In the HP OpenView Storage Data Protector Manager, switch to the Monitor context.

Sessions that are currently in progress are displayed in the Results Area. If there are no sessions currently in progress, the screen is empty.

2. Select the session you want to monitor.

Figure 3-16

Monitoring a Current Session

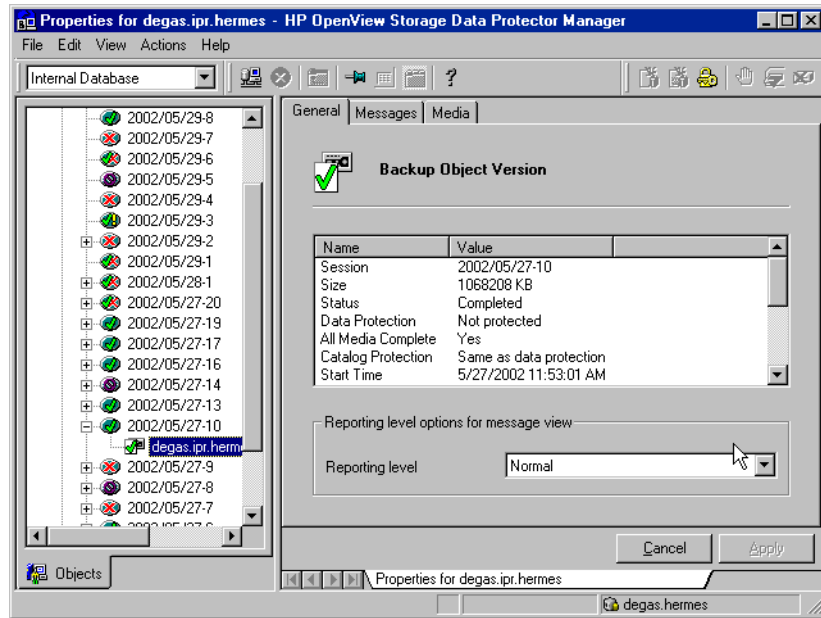


Viewing Previous Sessions

Use the following procedure to view a previous session using the Data Protector GUI:

1. In the HP OpenView Storage Data Protector Manager, switch to the Internal Database context.
2. Expand Sessions. A list of sessions appears in the Scoping Pane.
3. Select the session you want to preview.

Figure 3-17 Viewing a Previous Session



Troubleshooting

The following section provides some testing procedures you should perform before calling the Data Protector support. Following these guidelines, you may either resolve the problem yourself or identify the area where the problems occur.

Should you fail when performing a troubleshooting procedure, actions are proposed to help you work around the problem.

The section is divided into the following subsections:

- General troubleshooting
- Backup problems
- Restore problems

General Troubleshooting

1. Ensure that the latest official Data Protector patches are installed. Refer to “Verifying Which Data Protector Patches Are Installed” section 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, problems and workarounds, as well as for a list of related Data Protector patches.
3. Try to run a backup and restore without using Data Protector. Use the DB2 Command-Line Processor or the DB2 Control Center to back up and restore the DB2 database objects. For details, refer to the DB2 administration reference.

Backup Problems

General Backup Troubleshooting

Start a preview of the Data Protector DB2 Integration backup specification.

- If the DB2 Server part of the preview fails, refer to the DB2 documentation.
- If the Data Protector part of the preview fails, proceed as follows:

Create a DB2 Integration backup specification to back up to a null or file device and run the backup. If the backup succeeds, the problem may be related to the backup devices. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

- If the preview succeeds, proceed as follows:
 1. Check if the filesystem backup of the problematic client works. It is much easier to troubleshoot a filesystem backup.
 2. Examine the errors reported in the `/var/opt/omni/log/debug.log` and `/var/opt/omni/log/db2.log` files on HP-UX systems or in the `/usr/omni/log/debug.log` and `/usr/omni/log/db2.log` files on AIX systems.
 3. Try to restart DB2 Server and start the backup again.

Online Backup Is Not Allowed

Problem

DB2 reports that online backup is not allowed because either `logretain` or `userexit` option for rollforward is not activated or that a backup pending condition is in effect for the database.

Action

After configuring the DB2 database for the rollforward recovery (`userexit` and `logretain ON`), the database has to be first backed up offline. If online backup is started first, the above error will be reported.

Archived Logs Are Not Backed Up

- Problem** If you create several backup specifications and the last one is removed, the older backup specifications are not used and archived logs are not backed up.
- Action** Create a new backup specification.

Incremental Backup Is Not Enabled For the Database

- Problem** The following error message is displayed by Data Protector if incremental backup is attempted but no full backup has been performed: Incremental backup is not enabled for this database.
- Action** Perform the following steps:

1. Run the following command to activate modification tracking:

```
db2 update db cfg for <DatabaseName> USING TRACKMOD ON
```
2. Restart the database.
3. Perform a full database backup.

Error Occurred While Accessing an Object

- Problem** DB2 reports: SQL2048N An error occurred while accessing object *<object>*. Reason code: *<CodeNumber>*
- The following can be a reason (code number) for this error message:
1. An invalid object type is encountered.
 2. A lock object operation failed. The lock wait may have reached the lock timeout limit specified in the database configuration.
 3. An unlock object operation failed during the processing of a database utility.
 4. Access to an object failed.
 5. An object in the database is corrupted.
 6. The object being accessed is a table space. Either the table space is in such a state that the operation is not allowed or one or more containers of the table space is not available. (LIST TABLESPACES will list the current table space state.)

7. A delete object operation failed.
8. Trying to load/quiesce into a table that is not defined on this partition.

Action If a lock object operation failed, ensure that the lock timeout limit in the database configuration is adequate and resubmit the utility command. You may also consider using the QUIESCE command to bring the database to a quiesced state to ensure access.

Cannot List Table Spaces

Problem Data Protector reports: Cannot list table spaces.

Action

- Make sure that the database is not in a backup/restore/rollforward pending state.
- Make sure that the root user is in the db2admin group.

Restore Problems

General Restore Troubleshooting

- To restore to another instance, make sure that this instance is configured in Data Protector and running.
- Ensure that the filesystem restore of the problematic client works. It is much easier to troubleshoot a filesystem restore. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on troubleshooting filesystem restores.
- Examine the errors reported in the `/var/opt/omni/log/debug.log` and `/var/opt/omni/log/db2.log` files on HP-UX systems or in the `/usr/omni/log/debug.log` and `/usr/omni/log/db2.log` files on AIX systems.

Restore Finishes Successfully, But Rollforward Fails

Problem When you are performing a rollforward recovery from an online backup, restore finishes successfully, but rollforward fails.

Action Archived logs must be available in order to perform rollforward recovery. Check if they are available. If they are not available, restore them from the last backup.

In This Chapter

This chapter explains how to install, 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.

Organization of This Chapter

The chapter is organized into the following sections:

- “Overview” on page 257
- “Prerequisites” on page 259
- “Limitations” on page 260
- “Integration Concepts” on page 261
- “Installing and Upgrading the Integration” on page 263
- “Configuring the Integration” on page 265
- “Testing the Integration” on page 291
- “Backing Up a Sybase Database” on page 294
- “Restoring a Sybase Database” on page 303
- “Monitoring a Sybase Backup and Restore Session” on page 319
- “Sybase Character Sets” on page 320
- “Configuring the Integration as Cluster-Aware” on page 321
- “Troubleshooting” on page 323

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 **isql** commands:

- 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.

- 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.

- 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

This section provides you with a list of prerequisites you must be aware of before using the integration.

- You need a license to use the Data Protector Sybase integration. See 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, the UNIX operating system, and basic Data Protector functionality. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for Data Protector details.

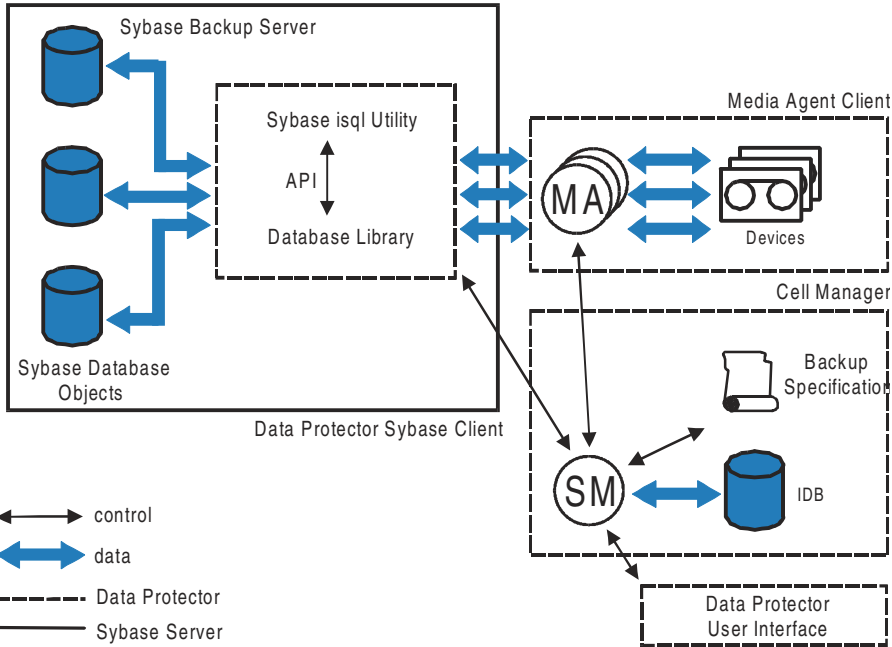
Limitations

See 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.

Integration Concepts

Figure 4-1 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 (**Backup And Restore**). 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. Refer to Figure 4-1 for the architecture of the Data Protector Sybase integration.

Table 4-1 Legend

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

Table 4-1

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 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 Media Agents (MA). The Data Protector `ob2sybase.exe` command starts the `sybase_<SYBASESERVERNAME>.sh` script, 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.

Installing and Upgrading the Integration

Upgrading

The Sybase integration is upgraded automatically during the client upgrade. For the upgrade procedure, refer to “Upgrading to Data Protector A.05.10” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Installation

You can install Data Protector software on your Sybase Backup Server locally, from a CD-ROM, or remotely, using the Data Protector User Interface.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for installation details.

Install the following software components:

Which Components Should I Install?

- Sybase Integration
- User Interface

Install this component to have access to the Data Protector GUI and the Data Protector CLI on this system.

- Disk Agent

Data Protector requires a Disk Agent to be installed on Backup Servers (clients with (filesystem) data to be backed up). Install the Disk Agent for two reasons:

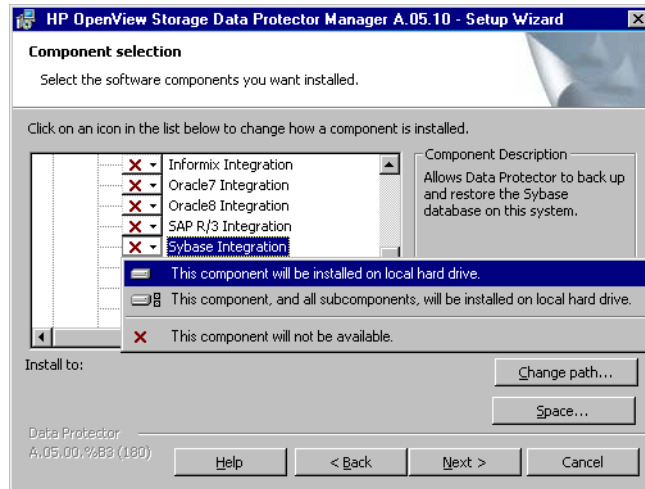
Why to Install the Disk Agent?

- ✓ To run a filesystem backup of Sybase Backup Server. Make this backup *before* configuring your Data Protector Sybase integration and resolve all problems related to Sybase Backup Server and Data Protector.
- ✓ To run a filesystem backup of important data that *cannot* be backed up using Sybase Backup Server.
- Media Agent

Install this component on clients with connected devices.

Integrating Sybase and Data Protector

Installing and Upgrading the Integration



Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

You have successfully installed Data Protector software on the Sybase Backup Server system. This system is now a client in the Data Protector cell. The integration is not yet ready for use. The next section gives you instructions on the procedure to configure the integration to make it ready for use.

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.

To configure the integration, follow these steps:

Configuration Overview

1. Configure a Sybase user

This is a user with appropriate rights in both Data Protector and Sybase environments as shown in “Configuring a Sybase User in Data Protector” on page 268.

2. Configure a Sybase Server

This is a client running Sybase Backup Server. Refer to “Configuring a Sybase Server” on page 270 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). Refer to “Configuring a Sybase Backup” on page 279 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. See “Installing and Upgrading the Integration” on page 263 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 language other than English, see “Sybase Character Sets” on page 320.

**Checking if
Sybase Server Is
Running**

Proceed as follows:

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

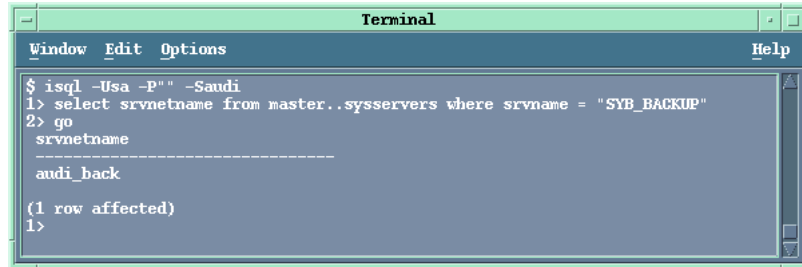
```
select srvnetname from master..syservers where srvname  
= "SYB_BACKUP"
```

and in the second line type `go`.

The name of Sybase Backup Server is returned.

In the following example, the name of Sybase SQL Server is `audi`.

Figure 4-2 Checking if Sybase Backup Server is Up



The name of Sybase Backup Server is *audi_back*.

- ✓ 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
- Full pathname of the Sybase `isql` command, for example,
/applications/sybase/bin/isql
- 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*.

Configuring a Sybase User in Data Protector

To start a Sybase backup session, a user needs 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 user *sybase* in 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 that 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, user *sybase* in group *sybase*) must be added to the Data Protector *admin* and *operator* groups.

Table 4-2 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 4-2 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 root 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.

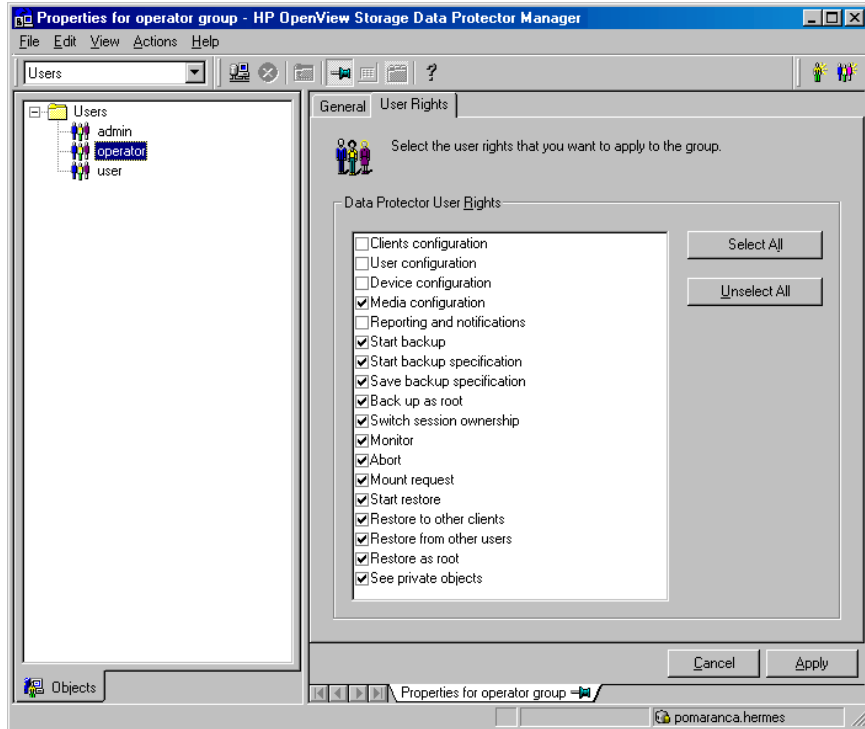
Data Protector User Rights

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 right, 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 4-3 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.

IMPORTANT

Do not proceed until you configure and run a filesystem backup of Sybase Backup Server as stated in “Before You Begin Configuring” on page 265. Refer to the *HP OpenView Storage Data Protector Administrator’s Guide* for instructions.

IMPORTANT

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 268 has permissions to read the files in the `/etc/opt/omni/sybase/<instance_name>` (HP-UX and Solaris systems) or in the `/usr/omni/config/sybase/<instance_name>` (other UNIX systems) directory on the Sybase Server.

Before you configure a Sybase Server, ensure that Sybase Backup Server is running. Refer to “Before You Begin Configuring” on page 265 for instructions. You can configure a Sybase Server using either the Data Protector GUI or the Data Protector CLI.

Using the Data Protector CLI

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> (HP-UX and Solaris systems) or  
  
/usr/omni/bin/util_sybase.exe -CONFIG <SYBASE_SERVERNAME>  
<SYBASE_HOME> <ISQL_PATH> <SYBASE_USER> <SYBASE_PASSWORD>  
<SYBASE_ASE> <SYBASE_OCS> (other UNIX systems)
```

NOTE

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

Integrating Sybase and Data Protector

Configuring the Integration

Where:

- `<SYBASE_SERVERNAME>`
is the name 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
```

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/sbin/omnigetmsg 12 <error_number>` for HP-UX and Solaris systems or `/usr/omni/bin/omnigetmsg 12 <error_number>` for other UNIX systems.

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 4-4 Getting an Error Description



Using the Data Protector GUI

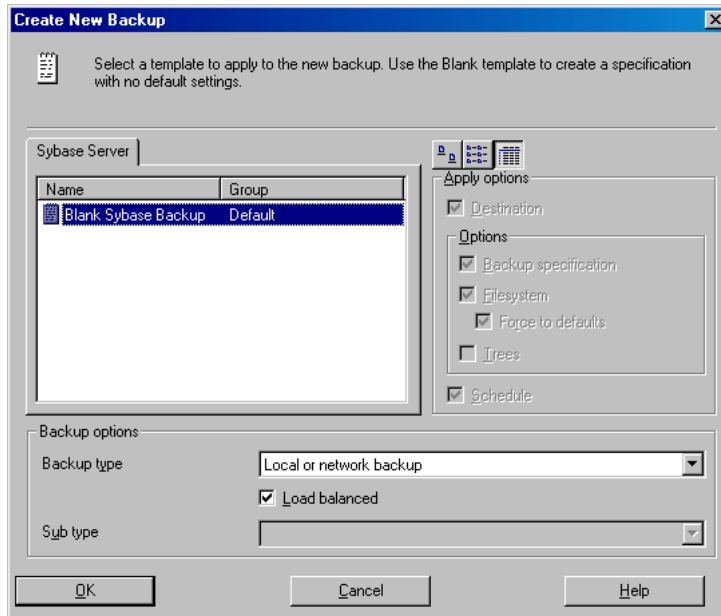
To configure a Sybase Server, perform the following steps in the HP OpenView Storage Data Protector Manager:

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 4-5 **Creating a New Sybase Backup**

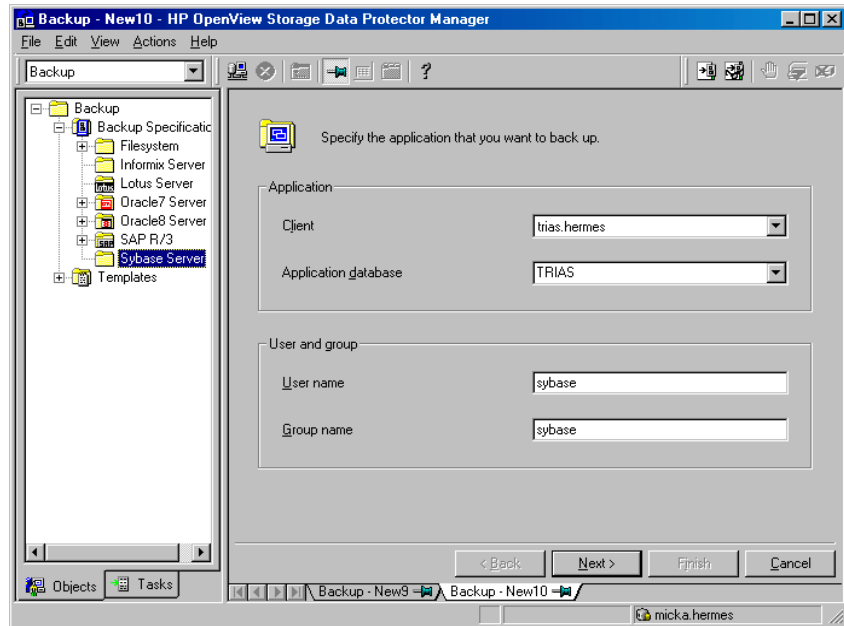


- Select the **Load balanced** option, which enables Data Protector to automatically balance the usage of devices that you select for the backup.
3. Click **OK**.

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 268. For example, user `sybase` in group `sybase`.

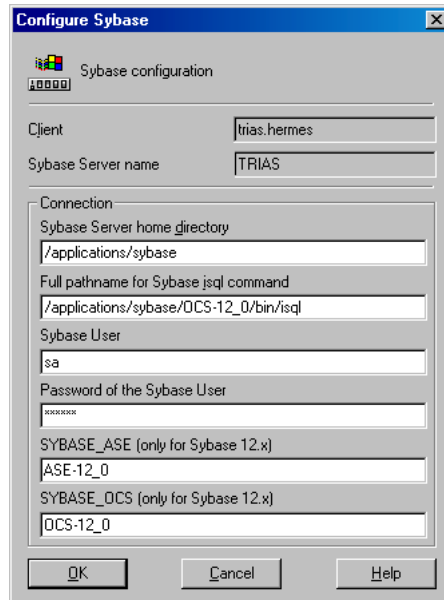
Figure 4-6 Configuring a Sybase Server



Click Next.

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

Figure 4-7 Configuring a Sybase Server



4. Enter the Sybase SQL Server home directory, for example, /applications/sybase/, the full pathname of the Sybase isql command, for example, /applications/sybase/bin/isql, 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
5. Click OK.

NOTE

If you receive a message that your Sybase Backup Server is not running, leave the configuration, run 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.

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 repository. On HP-UX and Solaris systems the repository is in the `/etc/opt/omni/sybase/<SYBASESERVERNAME>/` directory and on other UNIX in the `/usr/omni/config/sybase/<SYBASESERVERNAME>/` directory.
2. It creates the `sybase_<SYBASESERVERNAME>.sh` script.
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

You can check the Sybase configuration 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.

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

Using the Data Protector CLI

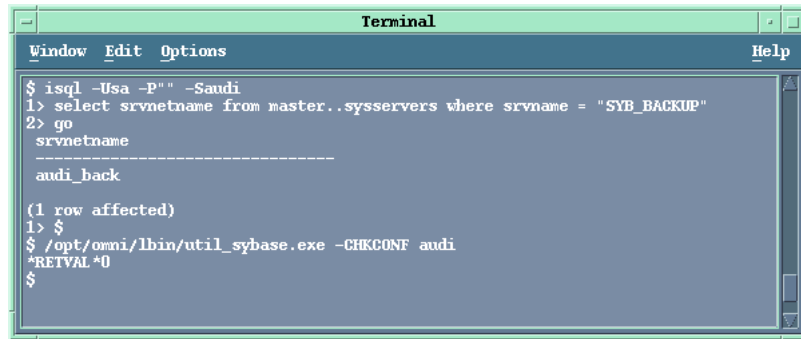
```
/opt/omni/sbin/util_sybase.exe -CHKCONF <SYBASESERVERNAME>  
(HP-UX and Solaris systems) or,
```

```
/usr/omni/bin/util_sybase.exe -CHKCONF <SYBASESERVERNAME>  
(other UNIX systems)
```

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

Figure 4-8 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 `*RETVAL*0` message was returned.

Figure 4-8 **Checking the Sybase Configuration**



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> for HP-UX and
Solaris systems or */usr/omni/bin/omnigetmsg 12 <error_number>*
for other UNIX.

Using the Data Protector GUI

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

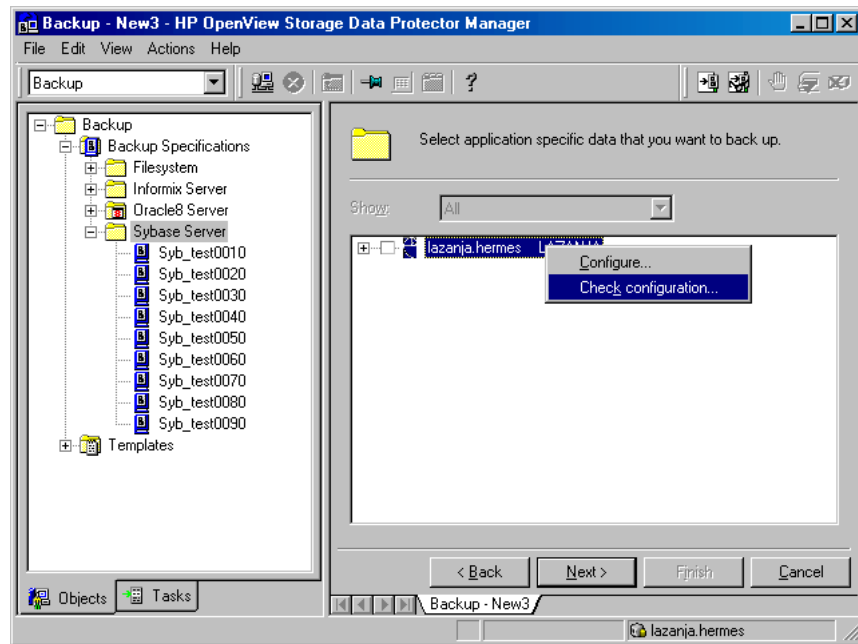
1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications, and then Sybase.
3. Go over the configuration procedure described in “Configuring a Sybase Server” on page 270 and the proceed with the procedure below.

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 4-9 Checking the Sybase Configuration



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. This is the last step before you run your first backup of Sybase data!

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

Configuring the Integration

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:

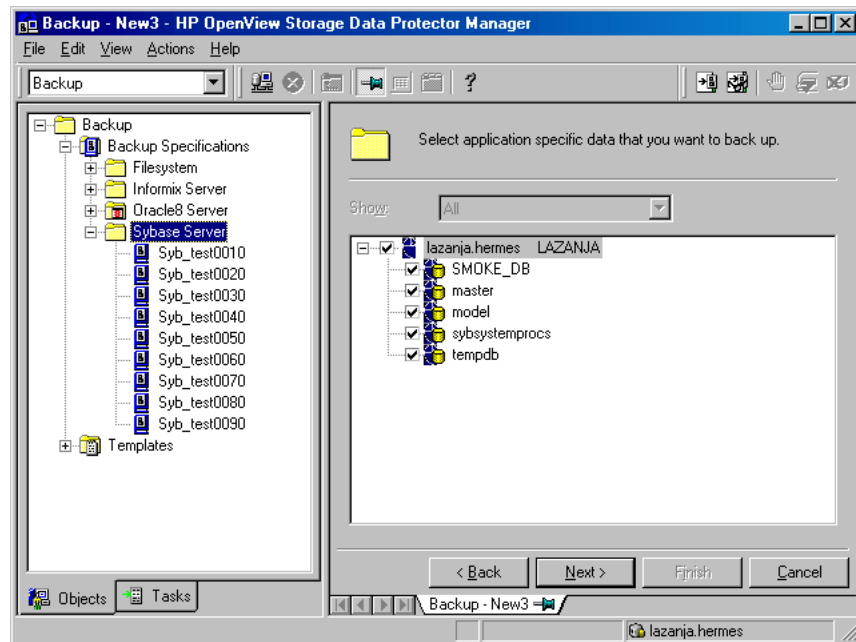
`/etc/opt/omni/barlists/sybase` (HP-UX or Solaris Cell Manager) or
`<Data_Protector_home>\config\barlists\sybase` (Windows Cell Manager).

A Sybase backup specification is created using the Data Protector GUI. Ensure that you have appropriate privileges. See “Configuring a Sybase User in Data Protector” on page 268 for more information.

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 270.

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 4-10, all databases were selected for backup.

Figure 4-10 Selecting Databases For Backup



Click Next.

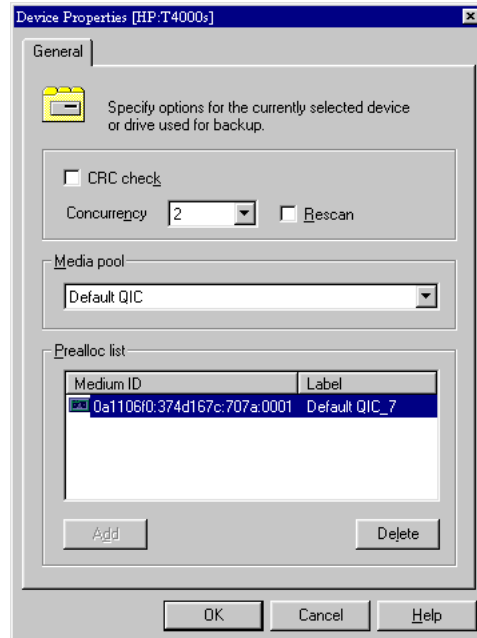
2. Select devices and media to which you will make your backup. See online Help for details.

NOTE

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

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

Figure 4-11 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 4-11, 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.

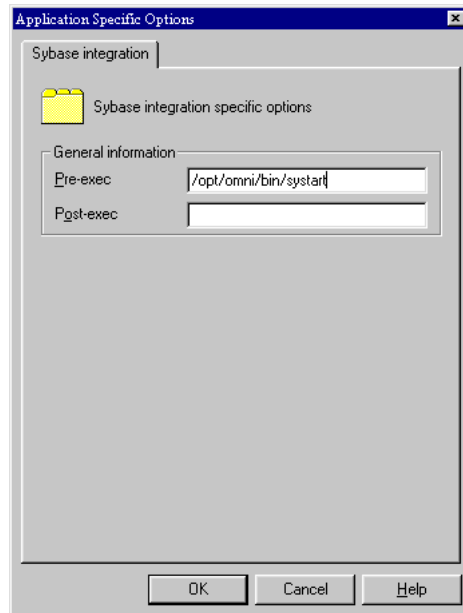
Under General Information in Figure 4-12, optionally specify the following:

- Pre-exec
a command that will be started on the Sybase Server before the backup. The command is started by the `ob2sybase.exe` command. 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. 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 4-12 Object-Specific Pre-Exec and Post-Exec Commands

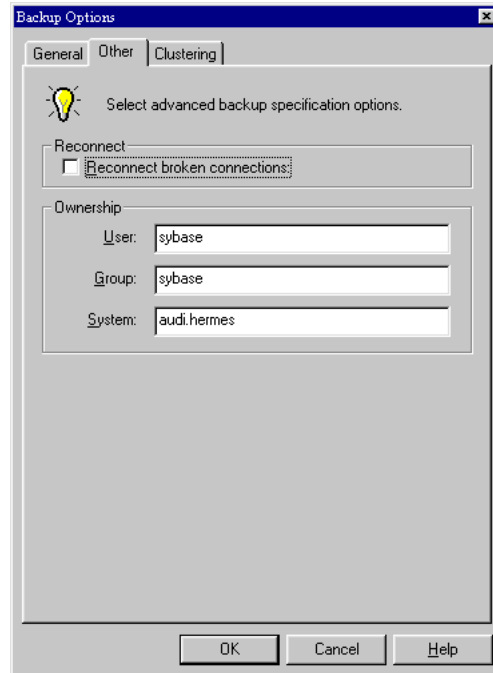


Changing the Sybase User

The Sybase user is the person 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 Other and edit the Ownership group box

Figure 4-13 Changing the Sybase User



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.

Scheduling Example

A full backup will be scheduled to start at 9.00 p.m. every Friday.

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.

Integrating Sybase and Data Protector Configuring the Integration

Click Add to open the Schedule Backup dialog box and specify the options as shown in Figure 4-14.

Figure 4-14 Scheduling a Weekly Full Backup

Schedule Backup

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

Recurring

None
 Daily
 Weekly
 Monthly

Time options

Time: 9 PM hours 00 minutes
 Use starting
Month: 2003 February Day: 13

Recurring options

Every 1 week(s) on
 Sun Mon Tue Wed Thu Fri Sat

Session options

Backup type: Full
Network load: High Medium Low
Backup protection: Default

OK Cancel Help

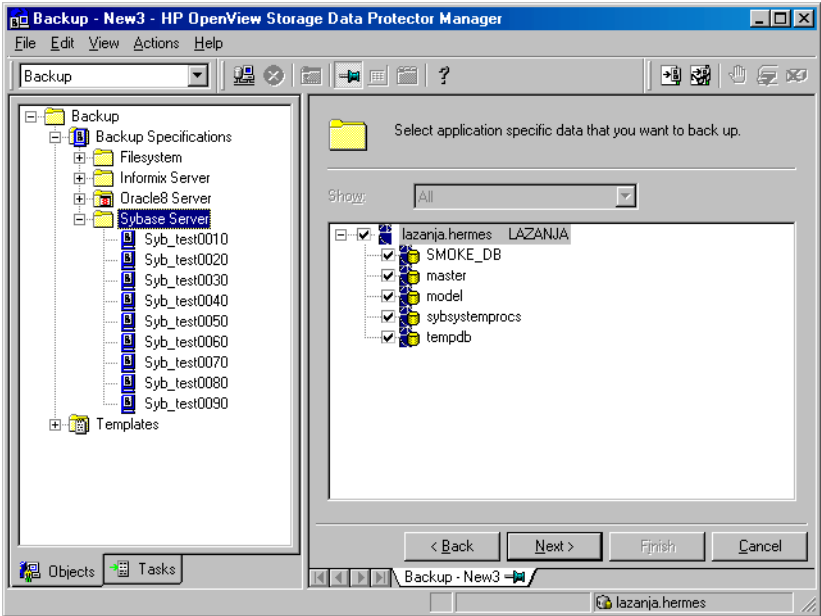
Click Next.

All the objects you selected for backup are displayed.

NOTE

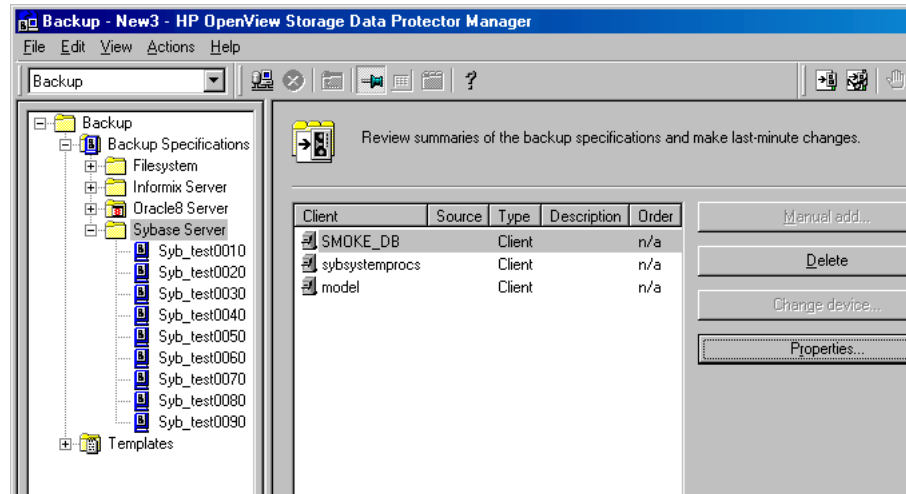
If you chose to back up the whole client by selecting the client as shown in Figure 4-15, then only the client is displayed in the Backup Specification Summary dialog box and not individual databases:

Figure 4-15 Selecting the Whole Client For Backup



In the example shown in Figure 4-16, individual objects were selected for backup. Hence, individual objects are displayed in the Backup Specification Summary dialog box.

Figure 4-16 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.

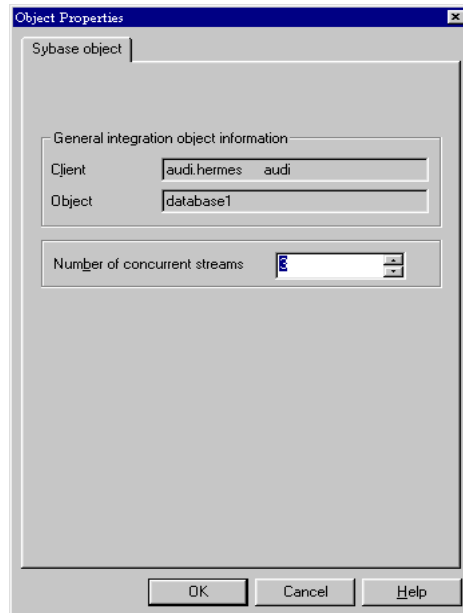
NOTE

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 4-17 **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. Refer to “Testing the Integration” on page 291.

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. See online Help for a description of these options.

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

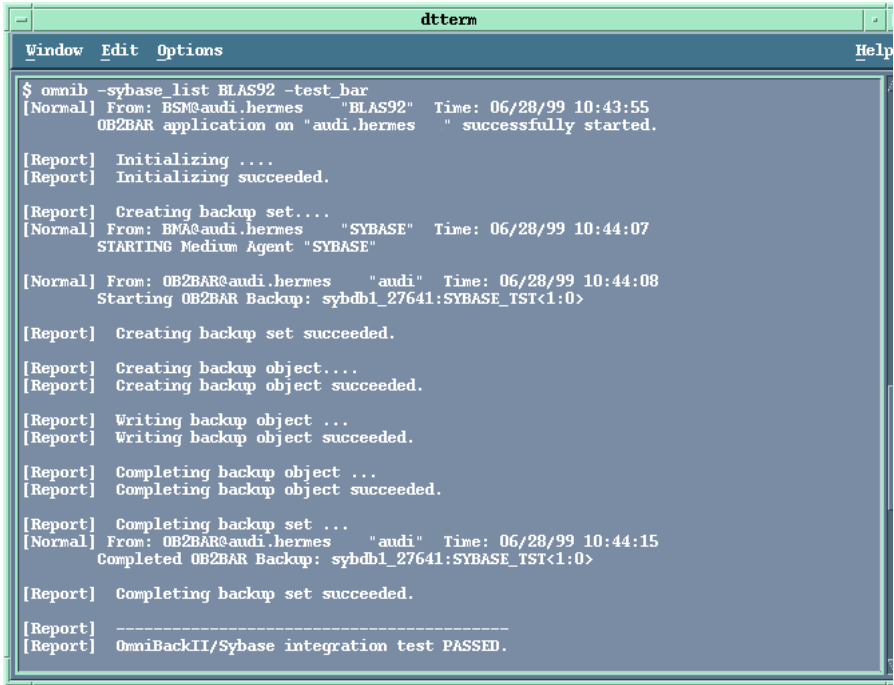
**The omnib
Command**

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

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

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

Figure 4-18 Testing the Configuration of the *BLAS92* Backup Specification



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

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

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

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

[Report] Creating backup set succeeded.

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

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

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

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

[Report] Completing backup set succeeded.

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

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

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

To get the error description, start the command,
/opt/omni/lbin/omnigetmsg 12 *<error_number>* for HP-UX and Solaris systems or /usr/omni/bin/omnigetmsg 12 *<error_number>* for other UNIX systems.

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 `testbar` command only tests 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 housekeeping 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. Refer to “Scheduling an Existing Backup Specification” on page 296.
- 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. Refer to “Running an Interactive Backup” on page 299.
- Start a backup using the Sybase CLI. Refer to “Backing Up Using Sybase Commands” on page 301.

Backup Types

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

Table 4-3 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.

See 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 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.

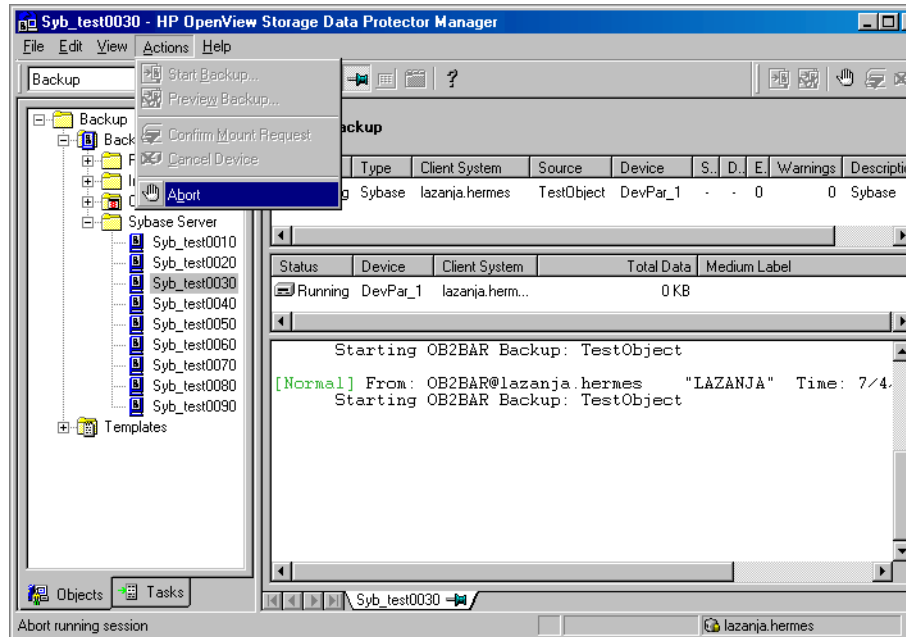
Messages from the Data Protector backup session and messages generated by Sybase are logged to the IDB. Upon successful completion of the backup, the “Session completed successfully” message is displayed in the `Session Information` dialog box.

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 4-19, the backup session of the backup specification `FullSybase` is being aborted.

Figure 4-19 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 280.

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

1. In the Context List, select Backup.

Scheduling Procedure

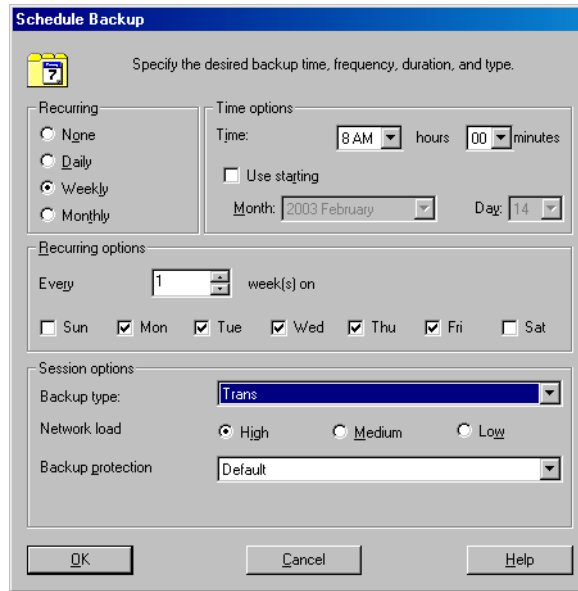
2. In the Scoping Pane, expand Backup, then Backup Specifications. Click Sybase Server.
A list of backup objects is displayed in the Results Area.
3. Double-click the backup specification you want to schedule and click the Schedule tab to open the Schedule property page.
4. In the Schedule property page, 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 4-20 on page 298.
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 4-20 on page 298.

Figure 4-20 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.

See 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 280 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`}. See online Help for a description of these options.

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

The `omnib` Command

You can also start an interactive backup of a Sybase database using the `omnib` command located in the Data Protector home 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 4-4

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 one 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.

```
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 291.

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 Data Protector home 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 Data Protector home 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 initialized 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:

```
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*

Backing Up a Sybase Database

<*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.

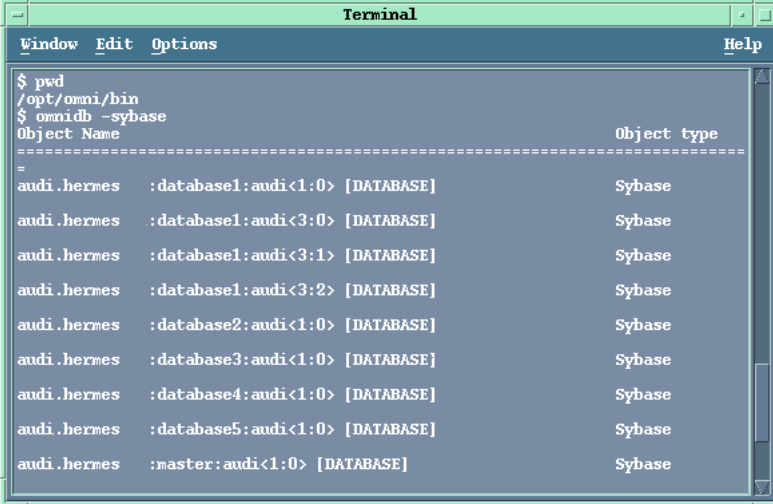
To restore a corrupted database, you need to find the right media and the sessionID of the last backup session with a full backup. If you have backed up a database with several streams, you need to know the number of streams. This information can be found using the Data Protector `omnidb` command. Refer to “The Data Protector `omnidb` Command” on page 303 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. Refer to “The Data Protector `syb_tool` Command” on page 310 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.

The Data Protector `omnidb` Command

To find the information needed to restore your data, execute the following commands in the `/opt/omni/bin/` directory:

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

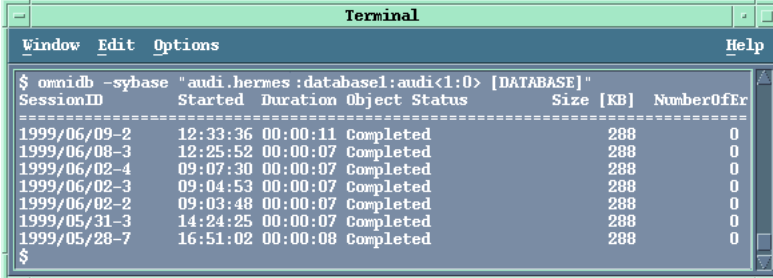
Figure 4-21 List of Sybase Objects



```
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
```

- omnidb -sybase "object_name"
to get details on a specific object, including the SessionID. Figure 4-22 shows how you get details about the object called `audi.hermes:database1:audi<1:0> [DATABASE]`.

Figure 4-22 Details about a Specific Session

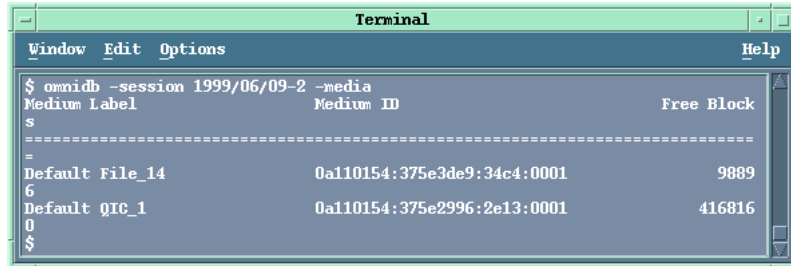


```
Terminal
Window Edit Options Help
$ omnidb -sybase "audi.hermes:database1:audi<1:0> [DATABASE]"
SessionID    Started Duration Object Status      Size [KB]  NumberOfEr
-----
1999/06/09-2 12:33:36 00:00:11 Completed 288        0
1999/06/08-3 12:25:52 00:00:07 Completed 288        0
1999/06/02-4 09:07:30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


- `omnidb -session <SessionID> -media`

to display media needed for restore. In Figure 4-23, media used for session `1999/06/09-2` are displayed.

Figure 4-23 Finding Media Needed for Restore



See the man pages for detailed information on the `omnidb` command.

Using the load Command

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

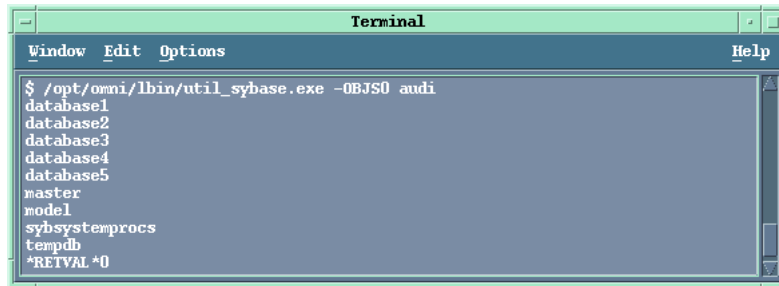
Starting the `isql` Command

1. Log on to your Sybase SQL Server as user `sybase`
2. Type in the following command in the Sybase SQL Server home directory:

```
bin/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.

Before you restore, you want to find out the databases on Sybase SQL Server. The `/usr/omni/bin/util_sybase -OBS0 <SYBASESERVERNAME>` (other UNIX systems) or `/opt/omni/lbin/util_sybase.exe -OBS0 <SYBASESERVERNAME>` command lists Sybase databases on the defined Sybase SQL Server, shown in Figure 4-24.

Figure 4-24 List Sybase Database Names



The Sybase load database and load trans commands are covered in detail in the *Sybase SQL Server System Administration Guide*. In this chapter, 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::::<old_db_name>[:<old_db_servername>]"
stripe on
"ob2syb::::<old_db_name>[:<old_db_servername>]"
where:
```

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

<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 “The Data Protector `omnidb` Command” on page 303.

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 “Starting the `isql` Command” on page 305.

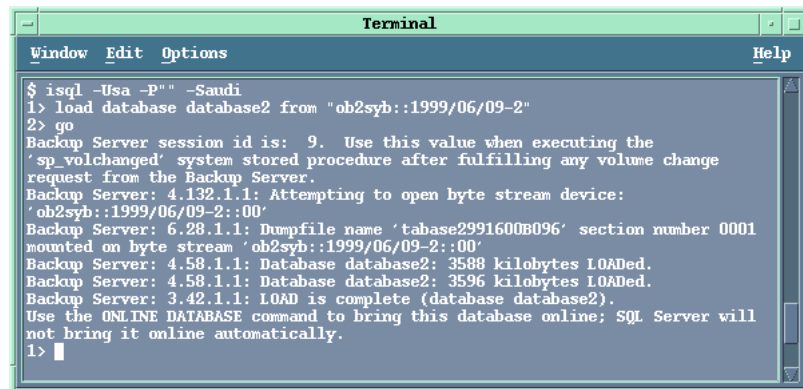
Example 1

To restore a database named `database2`, backed up in a session with sessionID `1999/06/09-2`, start the Sybase `isql` command and execute the following command, also shown in Figure 4-25:

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

Figure 4-25

Restoring `database2`, Backed Up in Session `1999/06/09-2`



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


Example 2

To restore 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 the following example:

Figure 4-26 **Creating a Database Device**




2. Create the empty database using the create database command, as shown in Figure 4-27. 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 4-27 **Creating an Empty Database**



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

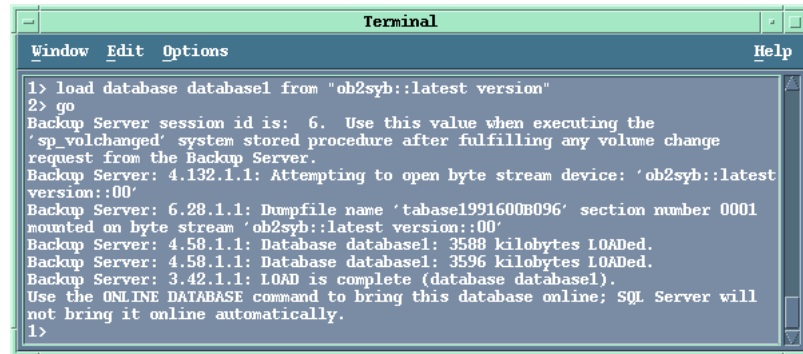
To restore *Sybdata* into *Sybdata1*, start the Sybase *isql* utility and execute the following commands:

```
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 4-28 Restoring the Latest Version of *database1*



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.

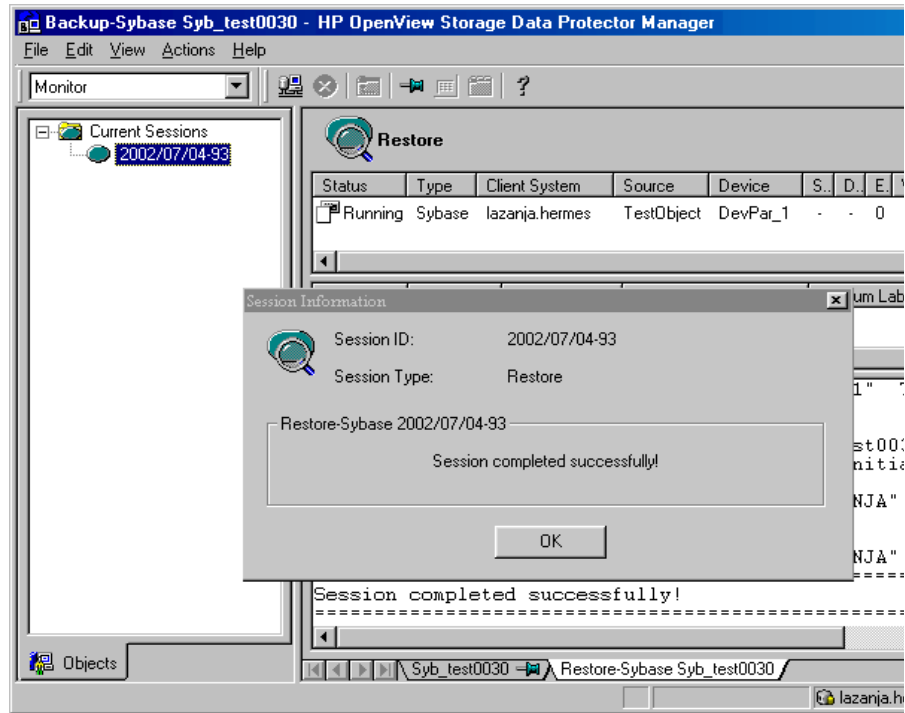
For example, to restore the latest version of a database named *database3*, backed up with three streams, start the Sybase *isql* utility and execute the following commands:

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

Monitoring Restore Sessions

Though you cannot start your restore sessions from the Data Protector GUI, use the Data Protector GUI to monitor your sessions. Refer to “Monitoring a Sybase Backup and Restore Session” on page 319 for instructions about monitoring your restore sessions. You get a “Session completed successfully” message at the end of a successful restore session, as shown Figure 4-29:

Figure 4-29 Monitoring Restore Sessions



The Data Protector `syb_tool` Command

The Data Protector `syb_tool` command, located in the `/opt/omni/bin` (HP-UX and Solaris systems) or `/usr/omni/bin` (other UNIX systems) directory, creates an `isql` load command that is used to restore a database on a specified date. To start this 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:

`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 name of Sybase SQL Server (required)
<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)
<code>media</code>	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.

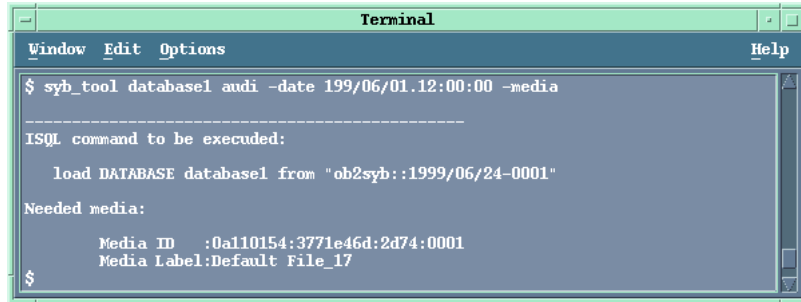
Example 1

To return an `isql` command to restore the first backup of the database `database1` on Sybase Server `audisherlock` that was made after 12:00 on August 12 June 1, 1999 and also the media on which the backup was made, type in the following command in the Data Protector home directory:

```
syb_tool database1 audisherlock -date 1999/06/01/08/12.12:00:00  
-media
```

The required `isql` command sequence as well as the media that you need for the restore are returned, as shown in Figure 4-30.

Figure 4-30 A load Command Including the Required Media

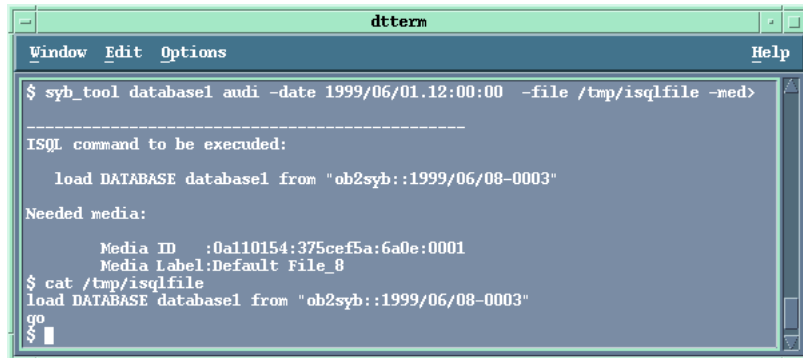


Example 2 To return the results of the above command to a file, */tmp/isqlfile*, type in the following command:

```
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 4-31.

Figure 4-31 A load Command Including the Required Media to a File

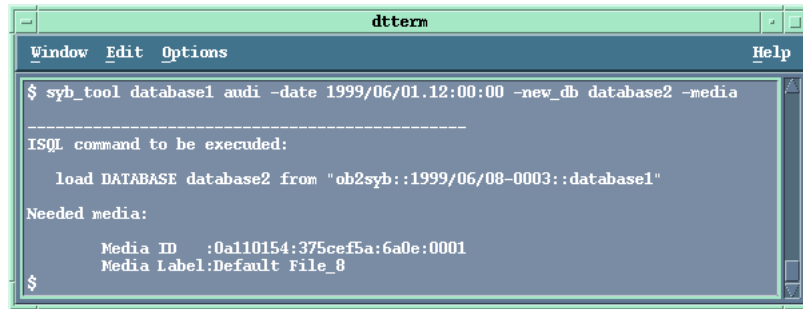


Example 3 To return a load command that restores a database *database1* to *database2*, perform the following command:

```
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 4-32.

Figure 4-32 A load Command to a Different Database

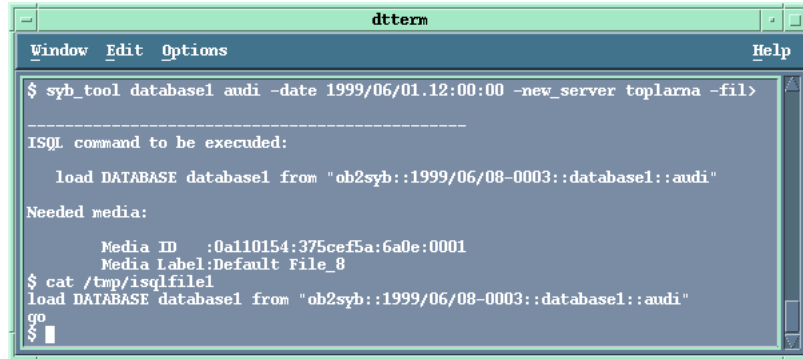


Example 4 To return a load command that restores a database *database1* backed up using the server *audi* to server *toplarna*, perform the following command:

```
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 4-33.

Figure 4-33 A load Command to a Different Server



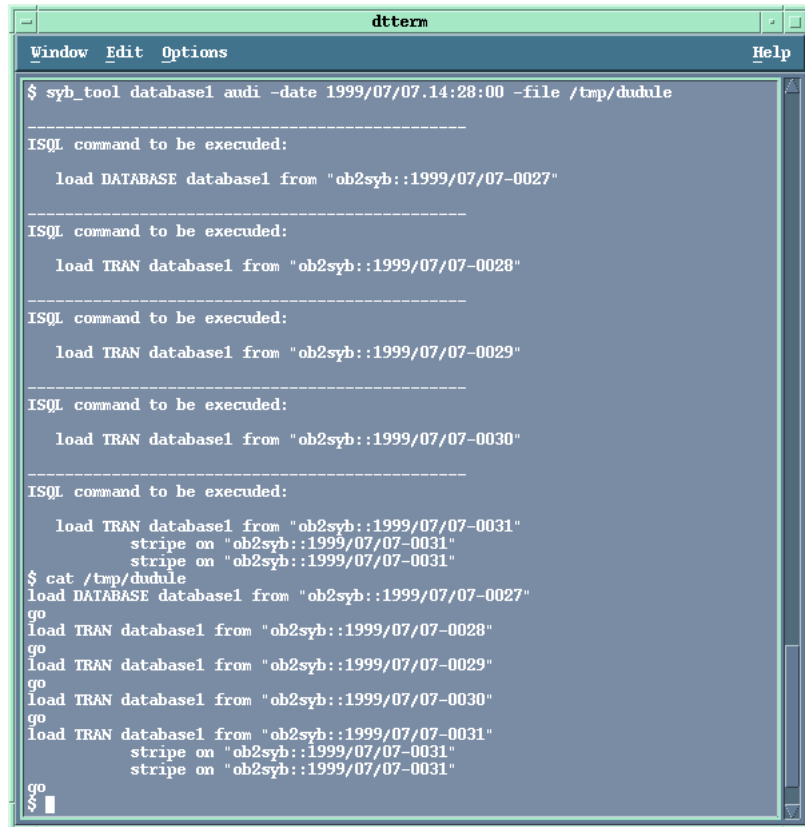
```
dtterm
Window Edit Options Help
$ syb_tool database1 audi -date 1999/06/01.12:00:00 -new_server toplarna -fil>
-----
ISQL command to be executed:
    load DATABASE database1 from "ob2syb::1999/06/08-0003::database1::audi"
Needed media:
    Media ID      :0a110154:375cef5a:6a0e:0001
    Media Label:Default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

Example 5 To return a 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 4-34.

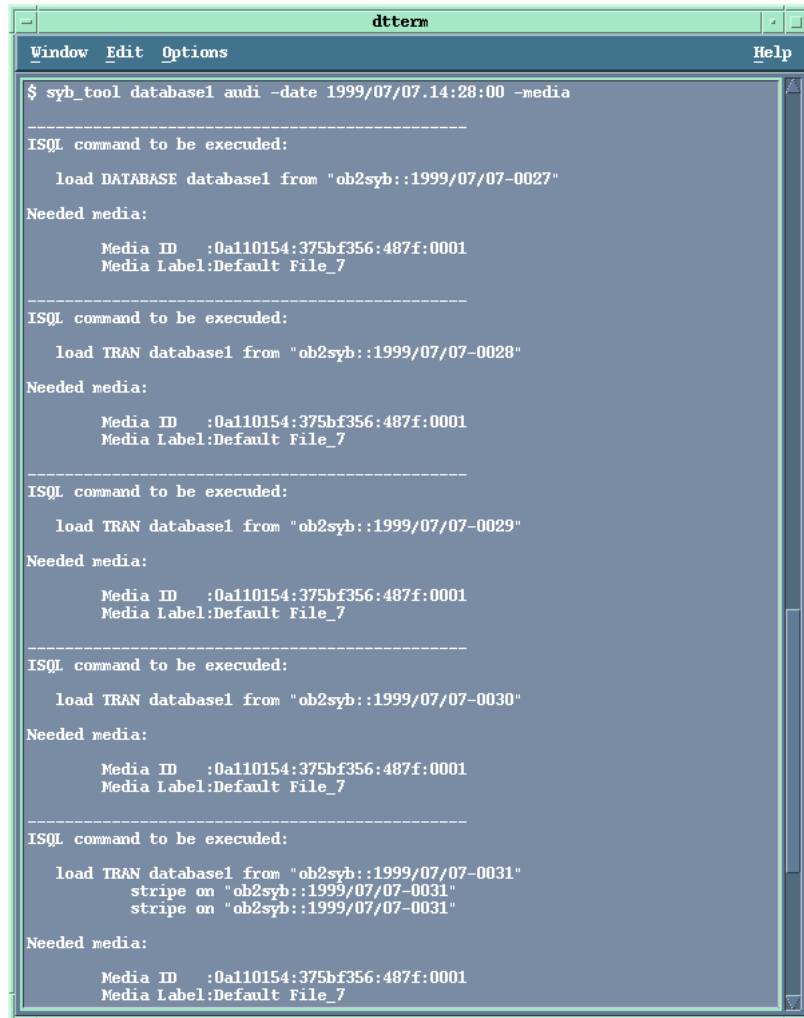
Figure 4-34 Loading Transaction Logs from Multiple Devices



```
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 above are shown in Figure 4-35.

Figure 4-35 Media Required for “Example 5” on page 314



```
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 the restore of Sybase database objects from devices other than those on which the database objects were backed up.

Specify these devices in the `/etc/opt/omni/cell/restoredev` 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.

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 see 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.
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.

Restoring a Sybase Database

4. Start restore. When 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 allows you to monitor currently running and view past backup and restore sessions. When you run an interactive backup or restore session, a monitor window opens 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.

To monitor a currently running session, proceed as follows in the HP OpenView Storage Data Protector Manager:

Monitoring Procedure

To get session details, double-click the running session you want to monitor.

At the end of the session a message is displayed indicating the success or failure of the session.

All actions and messages are logged to both Data Protector and Sybase log files. Error messages from the last backup are logged in the `/var/opt/omni/log/sybase.current.log` file. Mount prompt requests are displayed on the Data Protector monitor.

Sybase Character Sets

Sybase SQL server supports various language environments. Refer to *Sybase SQL Server Utility Programs* for more information. To enable this support in the Data Protector Sybase integration, you need to export the environmental variable `OB2_ISQL_OPTS` by adding the following line(s) to the

```
/etc/opt/omni/sybase/<SYBASE_SERVERNAME>/.profile (HP-UX and Solaris systems) or  
/usr/omni/config/sybase/<SYBASE_SERVERNAME>/.profile (other UNIX systems) file:
```

```
export OB2_ISQL_OPTS="-J<char_set>" (HP-UX and Solaris systems)  
or
```

```
OB2_ISQL_OPTS="-J<char_set>"
```

```
export OB2_ISQL_OPTS (other UNIX systems)
```

where `<char_set>` is the character set to be used. For example, adding the following line(s) to the `.profile` file:

```
export OB2_ISQL_OPTS="-Jsjis" (HP-UX and Solaris systems) or
```

```
OB2_ISQL_OPTS="-Jsjis"
```

```
export OB2_ISQL_OPTS (other UNIX systems)
```

What Happens?

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

Configuring the Integration as Cluster-Aware

Installation and Configuration

The Data Protector Sybase integration can be configured in the MC Service Guard cluster. This means that either the Data Protector Cell Manager can be configured in a cluster, or Data Protector client can be configured in the cluster. Refer to *HP OpenView Storage Data Protector Concepts Guide* for more information on supported configurations.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* and “Configuring the Integration” on page 265 for information on how install and configure the Data Protector Sybase MC Service Guard integration.

When configuring the Data Protector Sybase integration, configure it only on one of the cluster nodes per one Sybase server. Use the virtual hostname when configuring the integration. However, if the Cell Manager is outside the cluster, you need copy/append the Data Protector Sybase configuration files to all other nodes *after the integration has been configured on one node*.

Copy the following directory(ies) and file(s) to the same position on all other nodes:

```
/etc/opt/omni/sybase/<sybase_server_name>/ and  
/opt/omni/lbin/sybackup_<sybase_server_name>.sh
```

Append the following file to the same file on the same position on all other nodes:

```
/etc/opt/omni/sybase/SERVERLIST
```

For information on the Data Protector Cell Manager package configuration (if you want to install and configure a Data Protector Cell Manager in the MC/SG cluster), refer also to the *HP OpenView Storage Data Protector Administrator’s Guide*.

Backup and Restore

When creating a Data Protector Sybase MC/SG cluster backup specification, always select the virtual hostname in the cluster and not a particular node.

Refer to “Backing Up a Sybase Database” on page 294 for information on how to create a Data Protector Sybase backup specification and to *HP OpenView Storage Data Protector Administrator’s Guide* for information on MC/SG cluster backing up specifics.

Refer to “Restoring a Sybase Database” on page 303 for information on how to restore a Sybase database.

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.

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 4-36, <PASSWORD> is the Sybase Administrator  
password and <SA> is the Sybase user.
```

Figure 4-36 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

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 4-37, <PASSWORD> is the Sybase  
Administrator password and <SA> is the Sybase user.
```

Figure 4-37 Checking if Sybase Backup Server Is Running



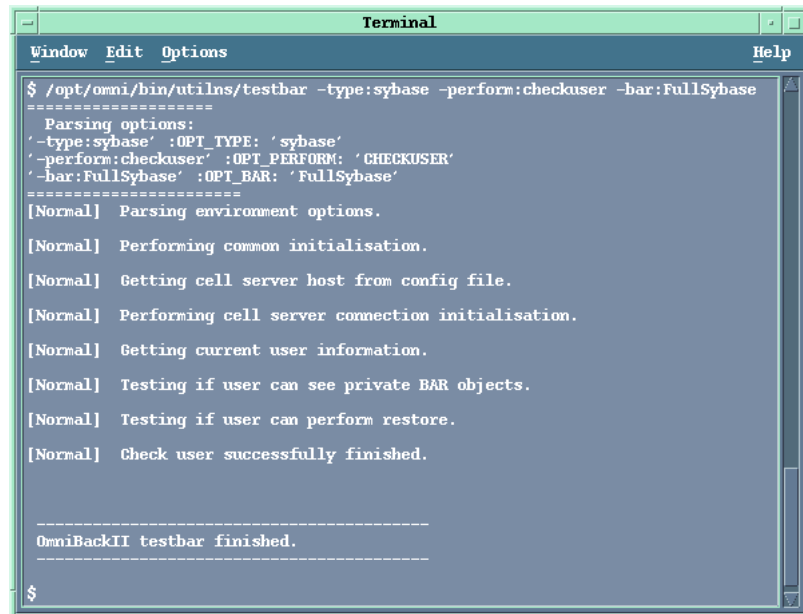
If the server is not running, then perform the following command in the Sybase home directory, to start it:

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

4. Examine system errors reported in the `/var/opt/omni/log/debug.log` (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the Sybase server.
5. If you have any non-default Sybase settings, ensure that they are registered in the `/etc/opt/omni/sybase/<SYBASESERVERNAME>/ .profile` (HP-UX and Solaris systems) or `/usr/omni/config/sybase/<SYBASESERVERNAME>/ .profile` (other UNIX systems) file.
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 4-38 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 4-38, 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.
```

Refer to “Configuring a Sybase User in Data Protector” on page 268 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 4-36, `<PASSWORD>` is the Sybase Administrator password `<SA>` is the Sybase user.

Figure 4-39

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


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 4-37, <PASSWORD> is the Sybase
Administrator password <SA> is the Sybase user.
```

Figure 4-40

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>` for HP-UX and Solaris systems or `/usr/omni/bin/omnigetmsg 12 <error_number>` for other UNIX systems.

Using the Data Protector CLI

Using the Data Protector GUI

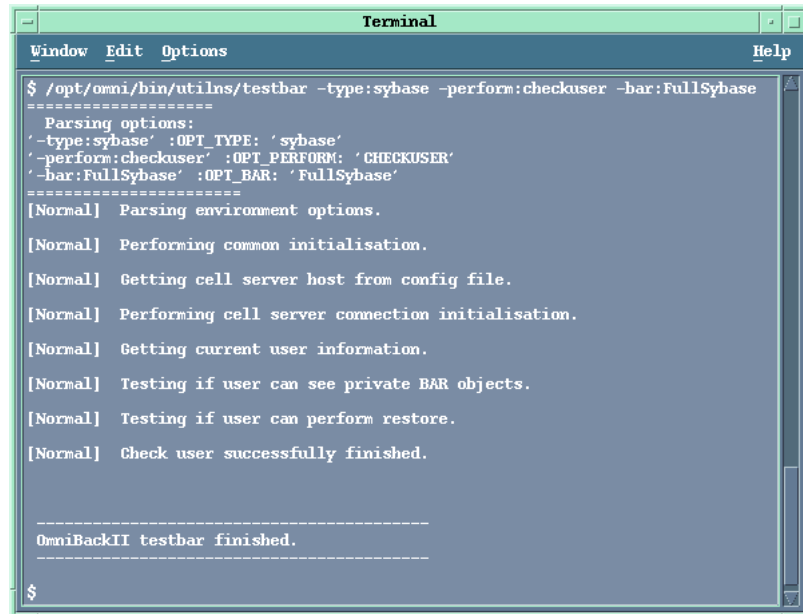
You can also check the configuration of your Sybase Server by performing the following steps in the HP OpenView Storage Data Protector Manager:

- 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 (HP-UX and Solaris systems)  
  
/usr/omni/bin/utilns/testbar -type:Sybase  
-perform:checkuser -bar:FullSybase (other UNIX systems)
```


Figure 4-41 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 4-41, 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.
```

Refer to “Configuring a Sybase User in Data Protector” on page 268 for information about the right privileges.

6. Verify that the Data Protector Cell Manager is correctly set on the Sybase Server by checking the `/etc/opt/omni/cell/cell_server` (HP-UX and Solaris systems) `/usr/omni/config/cell/cell_server` (other UNIX systems) file.

The name of the Data Protector Cell Manager is stored in this file. In the example shown in Figure 4-42, the configured Data Protector Cell Manager is called *audi.hermes*.

Figure 4-42 Verifying the Cell Manager



7. Test the Data Protector Sybase configuration as per instructions in “Testing the Integration” on page 291.

Example

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

```
/opt/omni/bin/omnib -sybase_list FullSybase -test_bar  
(HP-UX and Solaris systems), or
```

```
/usr/omni/bin/omnib -sybase_list FullSybase -test_bar  
(other UNIX systems).
```

- 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 301 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.

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 (for HP-UX and Solaris)
-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. Note, that for other UNIX systems, the correct path is /usr/omni/bin/utilns/testbar.

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 (HP-UX and Solaris systems) or
/usr/omni/log/debug.log (other UNIX systems) file on the Sybase Server.
9. If you have any non-default Sybase settings, ensure that they are registered in the
/etc/opt/omni/sybase/<SYBASESERVERNAME>/.profile (HP-UX and Solaris systems) or
/usr/omni/config/sybase/<SYBASESERVERNAME>/.profile (other UNIX systems) file.

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` (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) 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 4-36, <PASSWORD> is the Sybase Administrator
password and <SA> is the Sybase user.
```

Figure 4-43

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 4-44, <PASSWORD> is the Sybase
Administrator password and <SA> is the Sybase user.
```

Figure 4-44 **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 Data Protector Cell Manager is correctly set on the Sybase Server by opening the `/etc/opt/omni/cell/cell_server` (HP-UX and Solaris systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file.

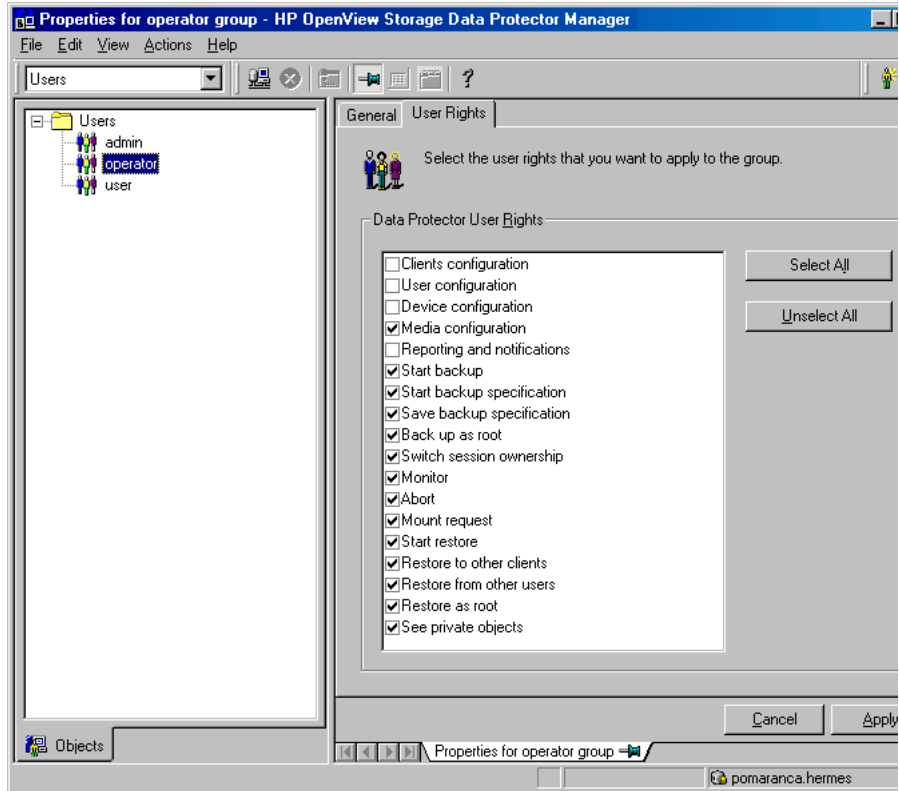
The name of the Data Protector Cell Manager is stored in this file. In the example shown in Figure 4-42, the configured Data Protector Cell Manager is called `audi.hermes`.

Figure 4-45 **Verifying the Cell Manager**



6. 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.
7. 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 4-46 **Setting the See Private Objects User Right**



- c. If the See private objects user right is selected, click Apply.
8. 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 (HP-UX  
and Solaris systems)
```

```
/usr/omni/bin/utilns/testbar -perform:checkuser (other  
UNIX systems)
```

Figure 4-47 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.

```

Refer to “Configuring a Sybase User in Data Protector” on page 268 for information about the right privileges.

9. 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 (HP-UX and Solaris)
-type:Sybase
-appname:<SYBASESERVERNAME>
-bar:<backup_specification_name>
-perform:backup

```

Troubleshooting

where `<SYBASESERVERNAME>` is the name of Sybase SQL Server and `<backup_specification_name>` the name of the Data Protector backup specification. Note, that for other UNIX systems, the correct path is `/usr/omni/bin/utilns/testbar`.

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` (HP-UX and Solaris systems) or
`/usr/omni/log/debug.log` (other UNIX systems) file on the Sybase Server.

In This Chapter

This chapter explains how to install, configure, and use the Data Protector Informix integration. It explains the concepts and methods that you need to understand to back up and restore Informix **dbobjects**. Information pertaining to Informix refers to OnLine Dynamic Server.

Organization of This Chapter

The chapter is organized into the following sections:

- “Overview” on page 339
- “Prerequisites” on page 342
- “Limitations” on page 343
- “Integration Concepts” on page 344
- “Installing and Upgrading the Integration” on page 346
- “Installing and Upgrading the Integration” on page 346
- “Configuring the Integration” on page 348
- “Testing the Integration” on page 374
- “Backing Up an Informix Database” on page 378
- “Restoring an Informix Database” on page 390
- “Monitoring an Informix Backup and Restore” on page 403
- “Configuring the Integration as Cluster-Aware” on page 404
- “Troubleshooting” on page 406

Overview

Data Protector integrates with **OnLine Server** to offer the online backup of your Informix database objects, hereinafter referred to as dbobjects.

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 integration.

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

Backup Types

You can perform the following types of backups of your Informix dbobjects using the Data Protector User Interface:

- Interactive backup using any of the following backup types:
 - Full, at which a baseline (level-0) backup of the selected dbobjects is made
 - Incr1, which backs up all changes since the last full (level-0) backup
 - Incr2, which backs up all changes since the last incremental (level-1) backup

NOTE

The Informix terms level-0, level-1, and level-2 backup are equivalent to Data Protector terms full, incr1, and incr2 backup, respectively.

- Scheduled backup of selected Informix dbobjects. You can schedule the same backup types as for interactive backups. 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.

NOTE

You can also back up Informix dobjects using the Informix `onbar` command.

A backup is always executed on OnLine Server via the Informix **ON-Bar** system. The **onbar utility** communicates backup and restore requests to OnLine Server.

You can restore your Informix dobjects using the Data Protector GUI or the Informix `onbar` command. Data Protector allows you to perform various types of restores, giving you all the flexibility you need to recover your mission-critical data.

Why Use the Data Protector User Interface?

Backing up and restoring using the integration offers various advantages over backing up using OnLine 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.

- 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.

- 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

This section provides you with a list of prerequisites you must be aware of before using the integration.

- Before you begin, ensure that you have correctly installed and configured OnLine 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
 - ✓ *INFORMIX-OnLine Dynamic Server: Backup and Restore Guide* for more information on configuring and using INFORMIX-OnLine Dynamic Server

Audience

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

Limitations

See the *HP OpenView Storage Data Protector Software Release Notes* for a list of general Data Protector limitations. 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.

Integration Concepts

Data Protector integrates with Informix 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 **XBSA interface**, the Informix onbar utility. Refer to Figure 5-1 for the architecture of the Data Protector Informix integration.

Figure 5-1

Informix Backup Concept

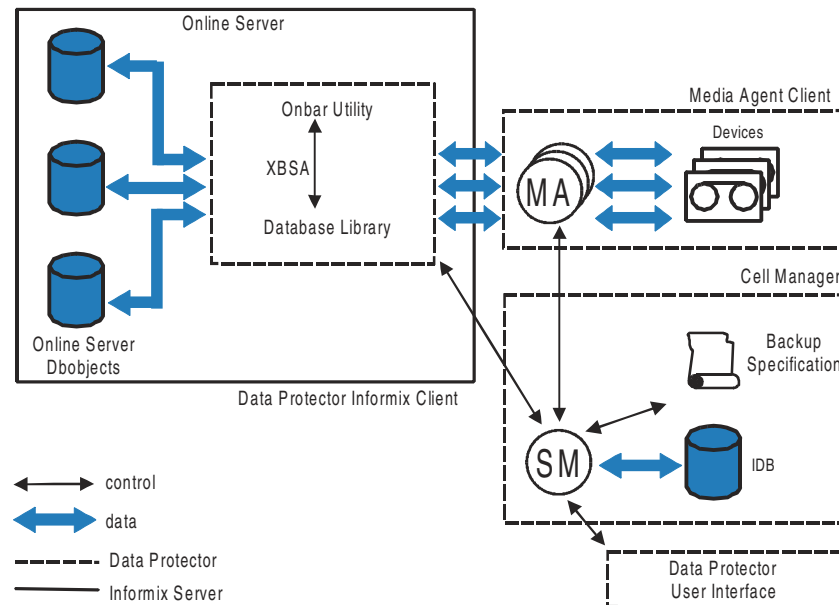


Table 5-1

Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
XBSA	X/Open Backup Services Application Programmer's Interface

Table 5-1 Legend

Database Library	The Data Protector set of routines that enables the data transfer between the OnLine Server and Data Protector.
MA	Data Protector Media Agent

Backup Specification

The onbar utility executes backup and restore requests coming from the Informix command-line and from Data Protector. 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.

XBSA

The onbar utility and Data Protector exchange control as well as backup and restore data via the X/Open Backup Services Application Programmer's Interface (XBSA). When a request to execute a backup or restore is received, the onbar utility initiates a session with both OnLine Server and Data Protector.

Backup Flow

For a backup, the onbar utility requests dbobjects from OnLine Server and passes them to Data Protector. Data Protector then writes the data to devices.

Restore Flow

For a restore, Data Protector retrieves the requested dbobjects from media and sends them through the XBSA interface to the onbar utility, which sends the data to OnLine Server for writing to disk.

While OnLine 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.

Installing and Upgrading the Integration

Upgrading

The Informix integration is upgraded automatically during the client upgrade. For the upgrade procedure, refer to “Upgrading to Data Protector A.05.10” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Installation

You can install Data Protector software on your OnLine Server locally, from a CD-ROM, or remotely, using the Data Protector User Interface.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Install the following software components:

Which Components Should I Install?

- Informix Integration
- User Interface

Install this component to have access to the Data Protector GUI and the Data Protector CLI on this system.

- Disk Agent

Data Protector requires a Disk Agent to be installed on Backup Servers. Install the Disk Agent for two reasons:

Why Install the Disk Agent?

- To run a test backup of any filesystem on the system where the OnLine Server is running. Make this backup *before* configuring your Data Protector Informix integration and resolve all problems related to OnLine Server and Data Protector.
- To run a filesystem backup of important data (Informix **ONCONFIG** file, Informix **sqlhosts** file, ON -Bar **emergency boot file**, `oncfg_<INFORMIXSERVER>.<SERVERNUM>`, HP-UX, Solaris configuration files, etc.) that *cannot* be backed up using OnLine Server.

- Media Agent

Install this component on Drive Servers (clients with connected devices).

Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

This system is now a client in the Data Protector cell. The integration is not yet ready for use. The next section gives you instructions on the procedure to configure the integration to make it ready for use.

Configuring the Integration

After the installation, the integration is not yet ready for use. The following subsections provide instructions for configuring the integration.

To configure the integration follow these steps:

Configuration Overview

1. Configure an Informix user

This is a user with appropriate rights in both Data Protector and Informix environments as shown in “Configuring an Informix User in Data Protector” on page 350.

2. Configure an OnLine Server

This is a client running OnLine Server. Refer to “Configuring an OnLine Server” on page 353.

3. Configure an Informix backup

Configure the devices and media needed for your backup, and create a Data Protector backup specification. Refer to “Configuring an Informix Backup” on page 360.

Before You Begin Configuring

Check the following before you start configuring:

✓ OnLine Server is up and running

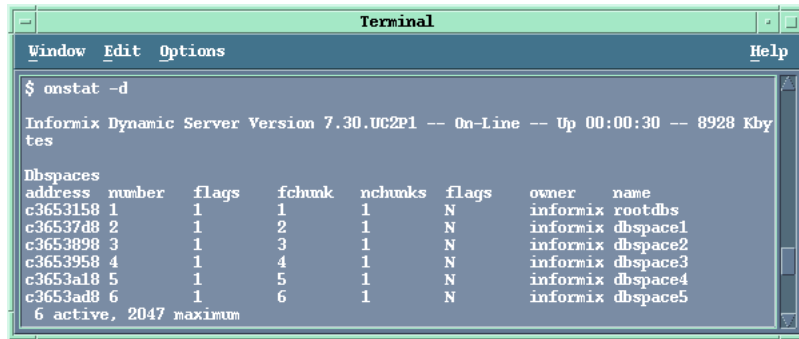
Before you start installing your integration software, ensure that your OnLine Server is running:

1. Log on to your OnLine Server as UNIX user `informix`
2. Type in the following command:

```
<INFORMIXDIR>/bin/onstat -d where <INFORMIXDIR> is the  
home directory of OnLine Server.
```

If OnLine Server is up and running, the `-- On-Line --` message is displayed as shown in Figure 5-2:

Figure 5-2 Checking if OnLine Server is Up



If OnLine Server is not up and running, proceed as follows to run it:

1. Log on to your OnLine Server as UNIX user `informix`
2. Type in the following command:

```
<INFORMIXDIR>/bin/oninit
```

where `<INFORMIXDIR>` is the home directory of OnLine Server.

- ✓ You can successfully run a test backup of any filesystem on the system where the OnLine Server is running.

Configure and run a Data Protector backup of any filesystem on the system where the OnLine Server is running for test purposes. By doing this, you check whether OnLine 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 OnLine 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.

- ✓ The Informix Integration Module has been installed on all OnLine Servers you want to back up.
- ✓ For each client running OnLine Server, you will need to provide the following information:

**Information
Needed to
Configure a
Backup**

- The Informix home directory, `<INFORMIXDIR>`, for example, `/applications/informix`.
- Filename of OnLine Server ONCONFIG configuration file, for example, `onconfig`. The ONCONFIG file is located in the `<INFORMIXDIR>/etc/` directory, where `<INFORMIXDIR>` is the Informix home directory.
- Full pathname of the OnLine Server sqlhosts configuration file, for example, `/applications/informix/etc/sqlhosts`. The sqlhosts file is located in the `<INFORMIXDIR>/etc/` directory, where `<INFORMIXDIR>` is the Informix home directory.
- Name of OnLine Server. This is stored in the `INFORMIXSERVER` shell variable.

Log on to OnLine Server as user `informix` in group `informix` and type in the following:

```
echo $INFORMIXSERVER
```

You should get the name of OnLine Server returned. In our case, the server is called `ODS730`.

Figure 5-3 Finding the Name of OnLine Server



For more information, refer to the *INFORMIX-OnLine Dynamic Server: Backup and Restore Guide*.

Configuring an Informix User in Data Protector

To start an Informix backup session, a user needs an operating system logon with sufficient privileges on the system where OnLine Server is running.

**Who Is the
Informix User?**

To find an Informix user with sufficient backup and restore privileges, run the following command on OnLine Server:

```
$ ls -l <INFORMIXDIR>/bin/onbar
```

(for OnLine Server 7.2x)

or

```
$ ls -l <INFORMIXDIR>/bin/onbar_d
```

(for OnLine Server 7.3x)

where <INFORMIXDIR> is the home directory of OnLine Server.

OnLine Server returns the user, in this case user root in group informix with the following permissions:

```
-rwsr-sr-x 1 root informix 1569592 June 10 1999
/applications/informix73/bin/onbar_d
```

where <INFORMIXDIR> is /applications/informix73.

Owner of an Informix Backup Specification

Using that logon the user must be able to back up and restore Informix dbobjects. To start a backup of an Informix dbobject using Data Protector, the user must then become the owner of a Data Protector Informix backup specification.

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

Table 5-2 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 5-2 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 root 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.

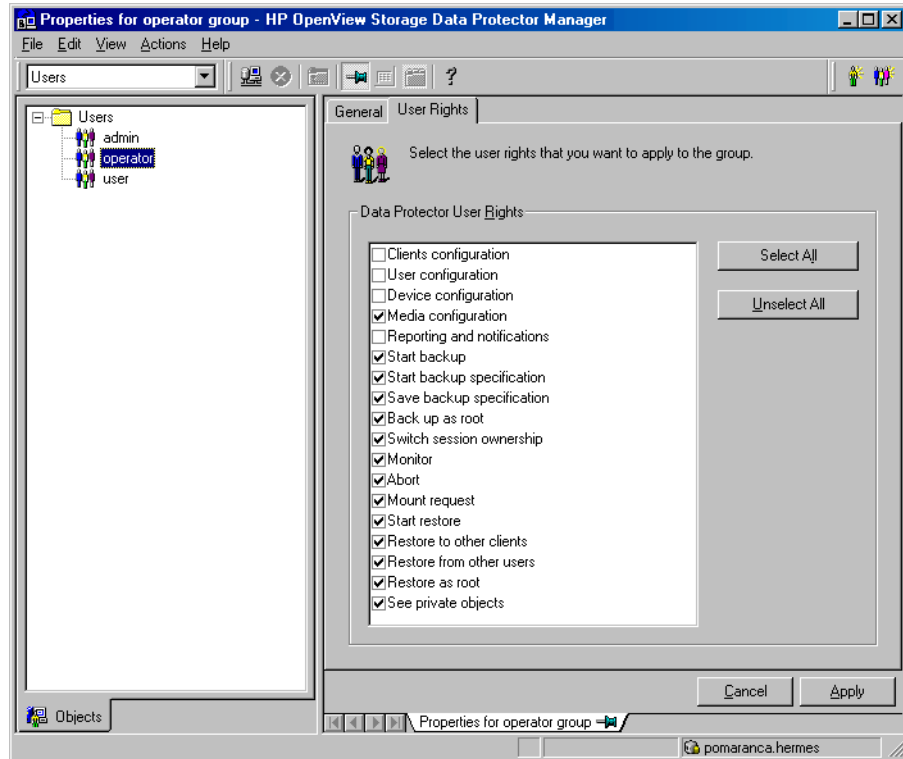
Data Protector User Rights

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 right, proceed as follows:

Configuration Procedure

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

Figure 5-4 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 an Informix user, who has appropriate rights in both the Data Protector and Informix environments. You are now ready to configure your OnLine Server.

Configuring an OnLine Server

Each system running OnLine Server must be configured for proper integration with Data Protector.

IMPORTANT

Do not proceed until you configure and run a test backup of any filesystem on the system where the OnLine Server is running, as stated in “Before You Begin Configuring” on page 348. Refer to the *HP OpenView Storage Data Protector Administrator’s Guide* for instructions.

IMPORTANT

After you have configured the Online Server, using either the CLI or GUI as described further on, make sure that the Informix user configured as described in the “Configuring an Informix User in Data Protector” on page 350 has permissions to read the files in the `/etc/opt/omni/informix/<instance_name>` (HP-UX and Solaris systems) or in the `/usr/omni/config/informix/<instance_name>` (other UNIX systems) directory on the Online Server.

Before you configure an OnLine Server, ensure that OnLine Server is running. You can configure an OnLine Server using either the Data Protector CLI or the Data Protector GUI.

Using the Data Protector CLI

Configuring an OnLine Server

To configure an OnLine Server, log in as user `root` on OnLine Server and execute the following command on the client you want to configure:

```
/opt/omni/sbin/util_informix.exe -CONFIG <INFORMIXSERVER>  
<INFORMIXDIR> <ISQL_PATH> <ONCONFIG> (HP-UX and Solaris systems)  
or /usr/omni/bin/util_informix.exe -CONFIG <INFORMIXSERVER>  
<INFORMIXDIR> <ISQL_PATH> <ONCONFIG> (other UNIX systems),
```

where:

Informix Configuration Options

<INFORMIXSERVER> is the name of OnLine Server.

<INFORMIXDIR> is the Informix home directory

<ISQL_PATH> is the full pathname of the OnLine Server `sqlhosts` file

Integrating Informix and Data Protector

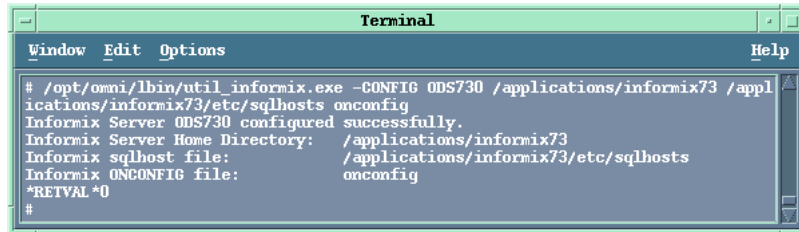
Configuring the Integration

<ONCONFIG> is the filename of the OnLine Server ONCONFIG file

Configuration Example

```
/opt/omni/lbin/util_informix.exe -CONFIG ODS730 /applications/informix73 /applications/informix73/etc/sqlhosts onconfig
```

Figure 5-5 OnLine Server CLI Configuration



```
Terminal
Window Edit Options Help
# /opt/omni/lbin/util_informix.exe -CONFIG ODS730 /applications/informix73 /appl
ications/informix73/etc/sqlhosts onconfig
Informix Server ODS730 configured successfully.
Informix Server Home Directory: /applications/informix73
Informix sqlhost file: /applications/informix73/etc/sqlhosts
Informix ONCONFIG file: onconfig
*RETVAl*0
#
```

If the configuration is successful, the message, *RETVAl*0, is displayed. Otherwise, the error number is displayed in the form *RETVAl* <error number>.

To see more details about the error, run the following command on OnLine Server:

```
/opt/omni/lbin/omnigetmsg 12 <error_number> (HP-UX and Solaris systems) or /usr/omni/bin/omnigetmsg 12 <error_number> (other UNIX systems).
```

or

Check the /var/opt/omni/log/informix.log file (HP-UX and Solaris systems) or /usr/opt/omni/log/informix.log file (other UNIX systems).

For more information, refer to the *INFORMIX-OnLine Dynamic Server: Backup and Restore Guide*.

Using the Data Protector GUI

Log in as user root and perform the following steps in the HP OpenView Storage Data Protector Manager:

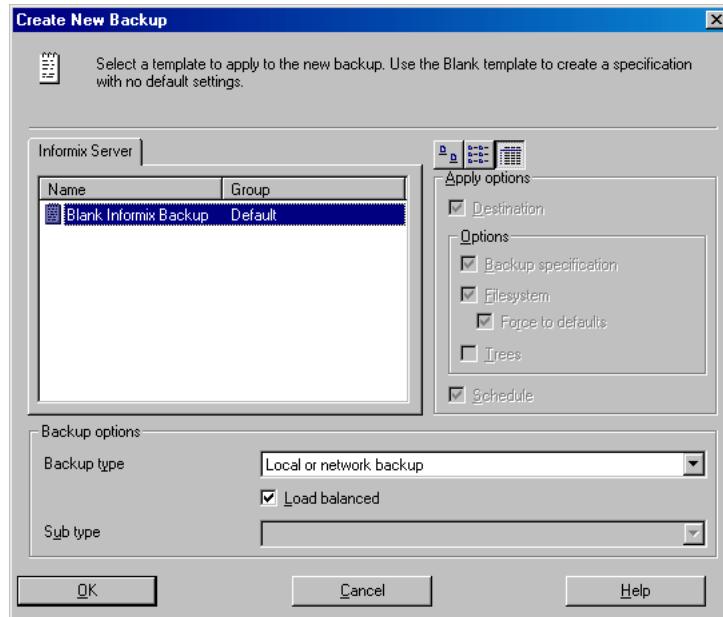
Configuring an OnLine Server

1. In the Context List, select Backup.

2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click Informix Server and click Add Backup.

The Create New Backup dialog box is displayed.

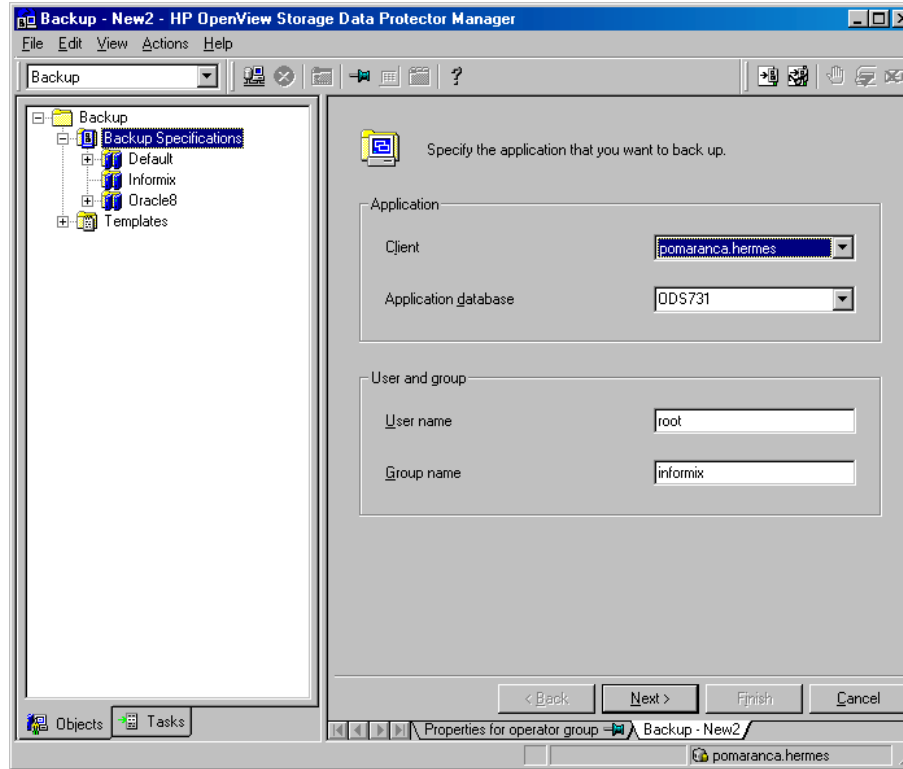
Figure 5-6 Selecting an Informix Backup Template



Select the Load balanced option, which enables Data Protector to automatically balance the usage of devices that you select for the backup, thus ensuring equal usage of the devices. For more information about Data Protector load balancing, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

3. Click OK.

Figure 5-7 **Configuring OnLine Server**



In the Results Area, enter the following information:

- The hostname of the OnLine Server you want to configure, for example, `hase.hermes`
- The name of OnLine Server, for example, `ODS731`
- The UNIX user name and user group of the Informix user, referred to in the section, “Configuring an Informix User in Data Protector”, for example, user `root` in group `informix`.

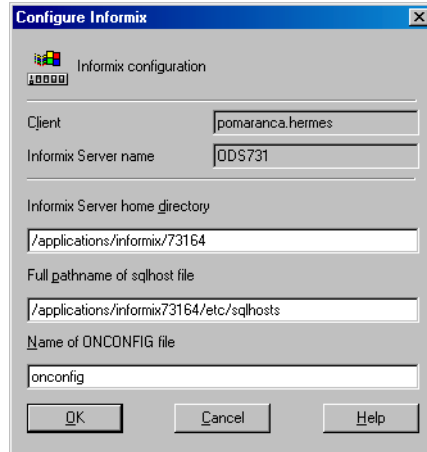
Click Next.

A message is displayed stating that the OnLine Server instance has not yet been configured.

Click OK.

The Configure Informix dialog box is displayed.

Figure 5-8 **Configuring Informix**

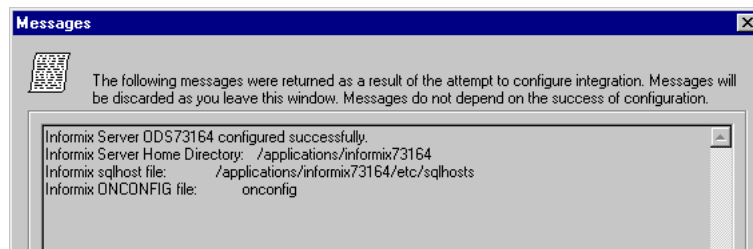


4. Enter the Informix home directory, for example, /applications/informix73, full pathname of the sqlhosts file, for example, /applications/informix73/etc/sqlhosts, and the name of the ONCONFIG file, for example, onconfig.

Click OK.

Upon successful configuration, a message like the one depicted in Figure 0-6 is returned:

Figure 5-9 **Success of the Configuration**



The message states that the configuration was successful and also returns the name of the configured OnLine Server, name of the OnLine Server home directory, full pathname of the sqlhosts file,

and the name of the Informix ONCONFIG file.

Click OK.

Otherwise, an error number is displayed in the form
**RETVAL* <error number>*.

To see more details about the error, run the following command on OnLine Server:

```
opt/omni/sbin/omnigetmsg 12 <error_number> (HP-UX and Solaris systems) or /usr/omni/bin/omnigetmsg 12 <error_number> (other UNIX systems).
```

or

Check the `/var/opt/omni/log/informix.log` file (HP-UX and Solaris systems) or `/usr/omni/log/informix.log` file (other UNIX systems).

What Happens?

The following happens after saving the configuration.

Data Protector starts the file, `util_informix.exe`, on the OnLine Server, which performs the following:

1. It saves the configuration parameters in the Data Protector repository. On HP-UX and Solaris systems, the repository is in the `/etc/opt/omni/informix/<INFORMIXSERVER>/` directory, and on other UNIX systems in the `/usr/omni/config/informix/<INFORMIXSERVER>` directory.
where `<INFORMIXSERVER>` is the name of OnLine Server.
2. It verifies connections to OnLine Server.

What's Next?

Before you start configuring your Data Protector Informix backup specifications, check the configuration, as per instructions in “Checking the Informix Configuration”.


Checking the Informix Configuration

To check the Informix configuration, start the following command:

Using the Data Protector CLI

```
/opt/omni/sbin/util_informix.exe -CHKCONF  
<INFORMIXSERVER>(HP-UX and Solaris systems) or  
/usr/omni/bin/util_informix.exe -CHKCONF <INFORMIXSERVER>  
(other UNIX systems).
```

Figure 5-10 Checking the Informix Configuration Using the CLI



If the configuration is successful, the message, `*RETVAL*0`, is displayed. Otherwise, the error number is displayed in the form `*RETVAL*error number`.

To see more details about the error, run the following command on OnLine Server:

```
/opt/omni/sbin/omnigetmsg 12 <error_number>(HP-UX and Solaris  
systems) or /usr/omni/bin/omnigetmsg 12 <error_number> (other  
UNIX systems).
```

or

Check the `/var/opt/omni/log/informix.log` file (HP-UX and Solaris systems) or `/usr/omni/log/informix.log` file (other UNIX systems).

Using the Data Protector GUI

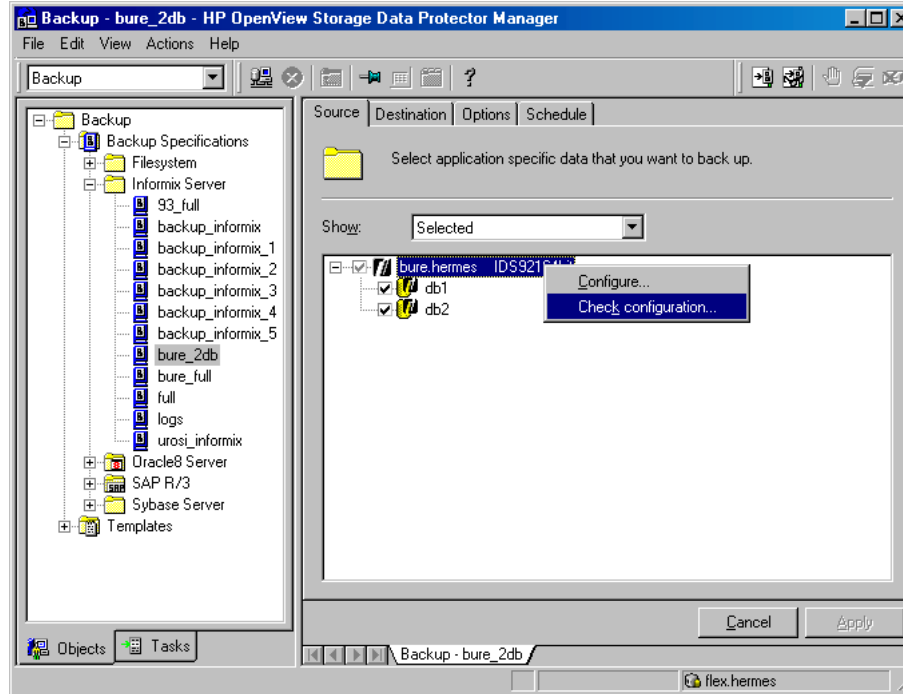
You can also check the configuration of your OnLine Server by performing the following steps in the HP OpenView Storage Data Protector Manager:

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications, and then Informix Server.
3. Go over the configuration procedure described in “Using the Data Protector GUI” on page 354.

Or, if you have already configured a backup specification, click it. The OnLine Server is displayed.
4. Right-click the client and then click Check Configuration.

Figure 5-11

Checking the Informix Configuration Using the Data Protector GUI



5. A message is returned confirming that the integration is properly configured.

What's Next?

Now that you have successfully configured your OnLine Server, you can configure your backup.

Configuring an Informix Backup

To run backups and restores of your Informix dbobjects, you need to configure Data Protector Informix backup specifications.

To configure the backup of Informix dbobjects, perform the following general steps:

Configuration Procedure

1. Configure devices, media and media pools needed for the backup. See the *HP OpenView Storage Data Protector Administrator's Guide* for instructions.
2. Create a Data Protector Informix 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.

Before You Begin

Perform the following preliminary tasks before creating your backup specification:

NOTE

Only general guidelines are given here. Refer to the *INFORMIX-Online Dynamic Server: Backup and Restore Guide* for detailed information about these tasks.

- ✓ Ensure that you have sufficient logical-log space to create a backup
If the total available space in the logical log (all the logical-log files) is less than half of a single log file, OnLine Server does not create a backup.
- ✓ Print or keep a copy of your ONCONFIG file, the sqlhosts file, and the emergency boot file
You need this information when you create a level-0 backup
- ✓ Verify data consistency
- ✓ Synchronize with other administrative tasks

Creating a Data Protector Informix Backup Specification

Informix backup specifications are located in the following directories on the HP-UX or Solaris Cell Manager:

```
/etc/opt/omni/barlists/informix
```

An Informix backup specification is created using the Data Protector GUI. Ensure that you have appropriate privileges. See “Configuring an Informix User in Data Protector” on page 350 for more information.

Integrating Informix and Data Protector Configuring the Integration

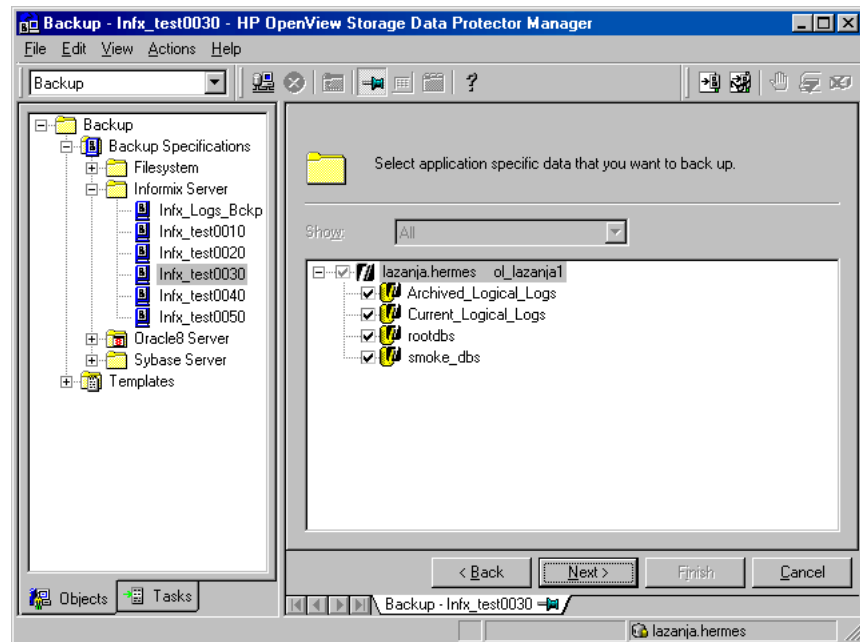
To create a Data Protector Informix backup specification on a client on which no backup specification has yet been configured, proceed from where you left off in “Using the Data Protector GUI” on page 354.

Procedure to Create a Backup Specification

1. In the Results Area, select the dbobjects you want to back up. The dbobjects include dbspaces, archive logical-logs, current logical-logs, and the root dbspace.

Figure 5-12

Selecting dbobjects to Backup



Click Next.

NOTE

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

2. Select devices and media needed for your backup. See online Help for details.

Allocate n primary devices that cover all resource types to be backed up. A typical configuration may include one device for logical logs and one for all other types.

Allocate m devices as secondary devices that cover all resource types in case any of the primary devices fails.

Configure primary devices and after that secondary devices.

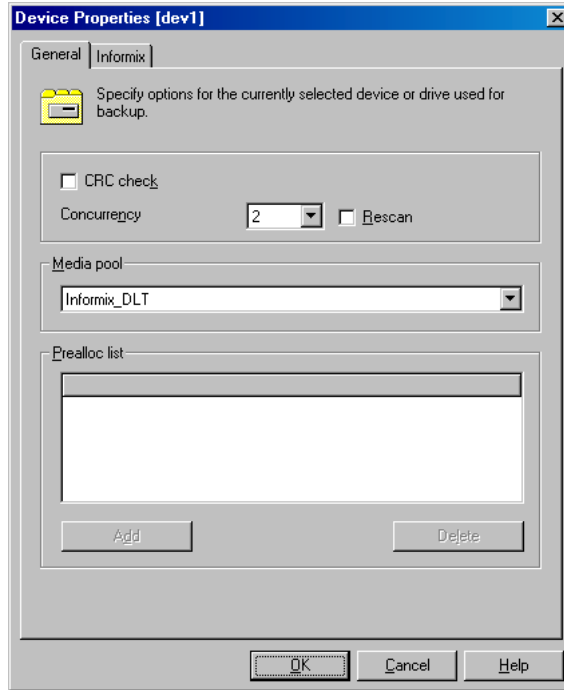
Select a device you want to use and click `Properties`.

The `Device Properties` dialog box is displayed. Specify the number of parallel backup streams in the `Concurrency` tab and the media pool you will use.

NOTE

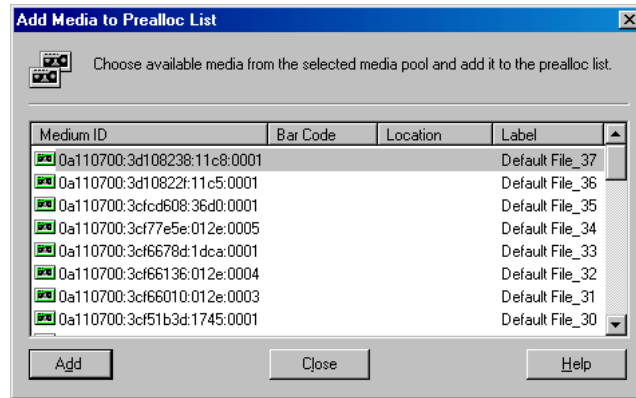
Except for whole-system backups and restores, OnLine Server backs up and restores dbobjects concurrently to achieve better performance than it would if it backed up or restored dbobjects serially. ON-Bar creates a new process for each object up to a limit specified by the `BAR_MAX_BACKUP` configuration parameter.

Figure 5-13 Specifying Device Properties



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

Figure 5-14 Adding Media to a Prealloc List



Select the media and click Add until all the media subsets have been added to the Prealloc List, and then click the Informix tab to set Informix resource types.

The resource type determines the types of dbobjects that will be backed up on this device. For example, if the resource type is set to R, only the root dbspace is backed up to this device.

The valid resource types are shown in Table 5-3.

Table 5-3 Device Resources

Resource Type	Description
B	blobospace
CD	critical dbspace*
L	logical log
MR	master root dbspace
ND	noncritical dbspace
R	root dbspace

Legend:

* The following dbspaces are critical dbspaces (CD):

A root dbspace

A dbspace that contains the physical log

Any dbspace that contains a logical-log file

As an example, select all types, except L. You will back up logical logs on a different medium dedicated only for that purpose.

NOTE

To backup logical logs, the LTAPE parameter in the ONCONFIG file must be set to a value that is not /dev/null or ''. The value will be ignored and the backup will be performed.

Figure 5-15 Informix Resource Types

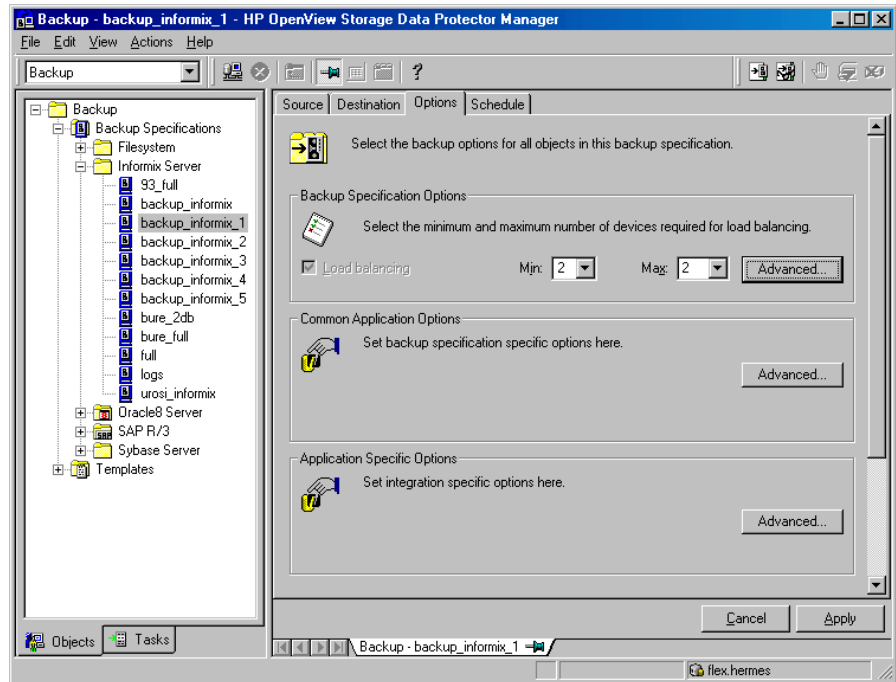


Click OK.

Repeat this step until you specify properties for all media.

3. Click Next to specify backup options.

Figure 5-16 Specifying Backup Options



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.

Specify the number of the primary devices (n) for the minimum and maximum value.

Make sure primary devices are not locked prior to starting Informix backup. You can do so by scheduling backups appropriately. If this is not true, secondary devices with possibly wrong resource type may get used.

Object-Specific Pre-Exec and Post-Exec Commands

Under Application Specific Options, click Advanced, to specify pre-exec and post-exec commands that will be started before ob2onbar is started on the Informix server and after it has finished.

Configuring the Integration

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 started by BSM on any specified client before and after ob2onbar is started and finished on the Informix server.

Under Informix Integration, optionally specify the following:

- Pre-exec

a command that will be started on the OnLine Server before backup. The command is started by the ob2onbar.exe command. The full path for the command must be provided.

TIP

Use the onmode -l OnLine command as a Pre-exec command, to ensure that you always have a log file to back up. This is useful if you specified a logical-log backup, since the backup will fail if there are no logical logs to back up.

-
- Post-exec

a command that will be started on the OnLine Server after backup. The command is started by the ob2onbar.exe command. The full path for the command must be provided.

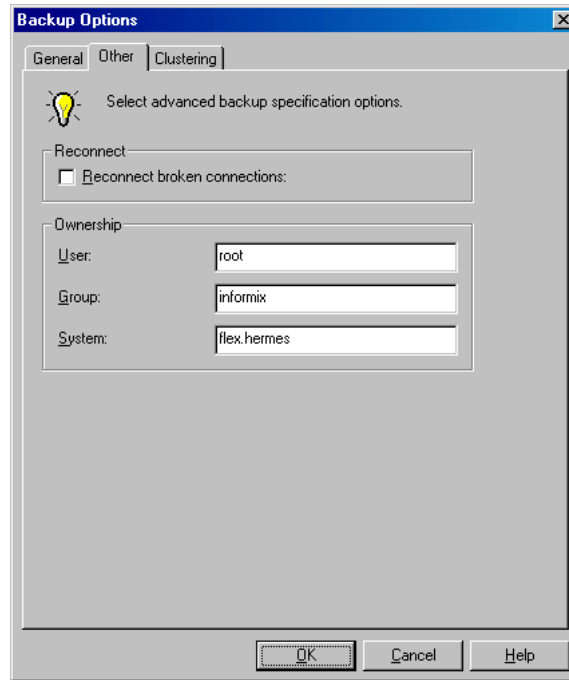
IMPORTANT

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

Changing the Informix User

The Informix user is the person 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 Informix user, click the Other tab in the Backup Options dialog box.

Figure 5-17 Changing the Informix User



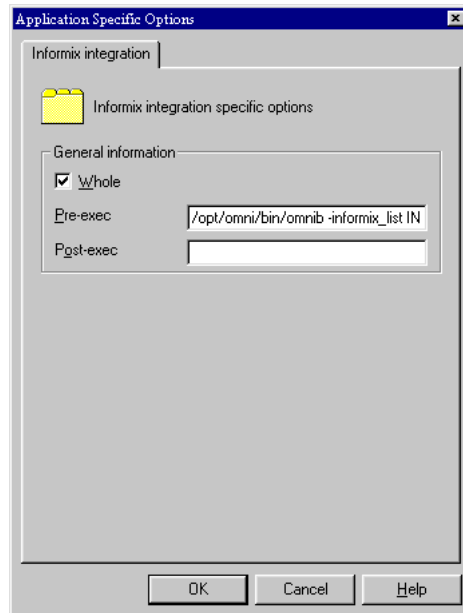
Click OK to return to the main backup options dialog box.

Informix Whole-System Backup

A whole-system backup is a backup of all OnLine Server dbobjects from one onbar command. ON-Bar cannot back up dbobjects concurrently during a whole-system backup, so dbobjects are backed up serially. A whole-system backup is useful for disaster recovery and for restoring from a client other than the one to which the backup was made. If a computer's disk is completely destroyed and needs to be replaced with a new disk, you will need either a full (level-0) backup of every dbspace, blob space, and logical-log file or a full (level-0) whole-system backup to completely restore data on the replacement computer.

As an example, select `Whole`, to make a whole-system Informix backup.

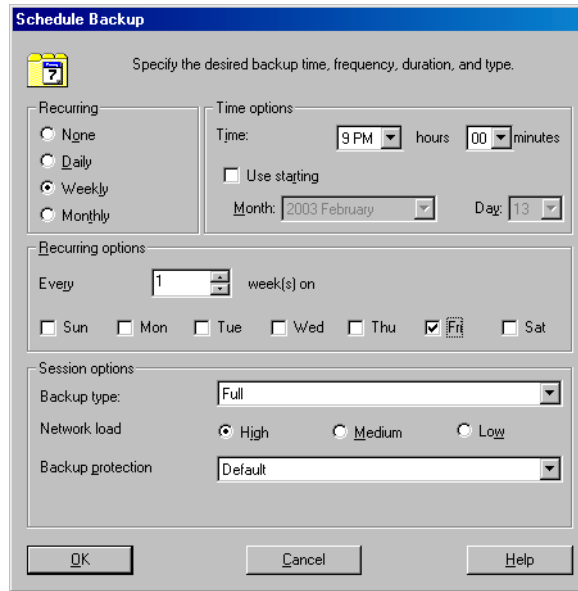
Figure 5-18 Whole-System Backup



4. 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.

For example, to schedule a full backup 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 5-19.

Figure 5-19 Scheduling a Weekly Full Backup

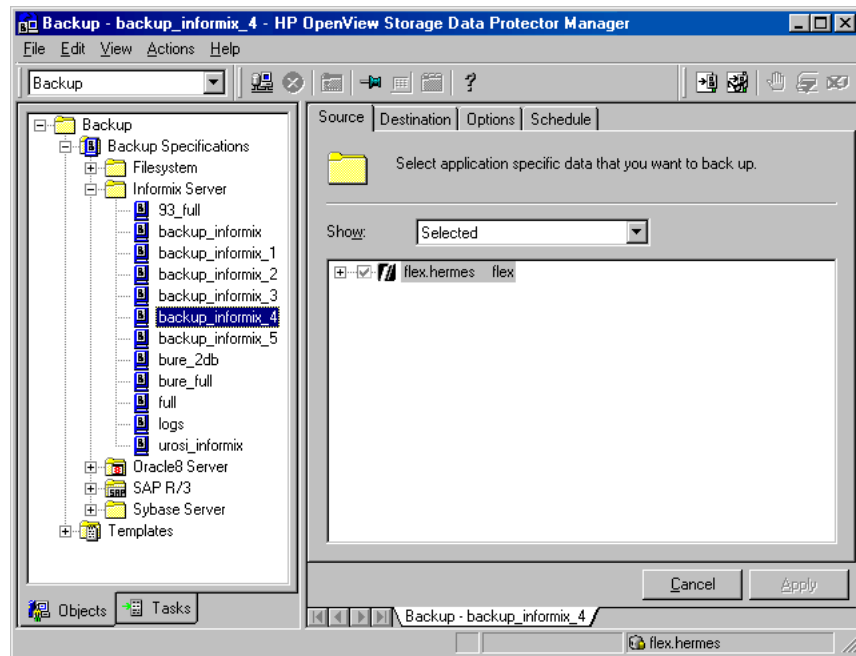


Click Next and then save your backup specification. Click Start Preview, to test your backup specification.

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 as shown in Figure 5-20. Click the appropriate tab and implement the changes you want. You need to save the backup specification afterwards.

Figure 5-20 **Editing a Backup Specification**



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 logical logs.

Test your backup specifications thoroughly before using them for real. Refer to “Testing the Integration” on page 374.

Configuring Informix Enterprise Decision Server (EDS)

Limitation

Currently the Informix Enterprise Decision Server is supported only on HP-UX and Solaris systems. Refer to the latest versions of the Support Matrices at <http://support.openview.hp.com/support.jsp> for up-to-date information.

The Informix Enterprise Decision Server uses onbar-workers to backup data. When a backup is started, the onbar starts onbar_d which then starts onbar-workers (if they are not running yet) using the

<Informix_home_dir>/etc/start_worker.sh script. After onbar-workers are started, the onbar_d passes the information about what needs to be backed up to onbar-workers.

To configure a Data Protector Informix Enterprise Decision Server backup, you need to add the following lines to the *<Informix_home_dir>/etc/start_worker.sh* script after configuring and saving a backup specification:

```
export OB2APPNAME = <INFORMIXSERVERNAME>
```

```
export OB2BARLIST = <saved_backup_specification_name>
```

Testing the Integration

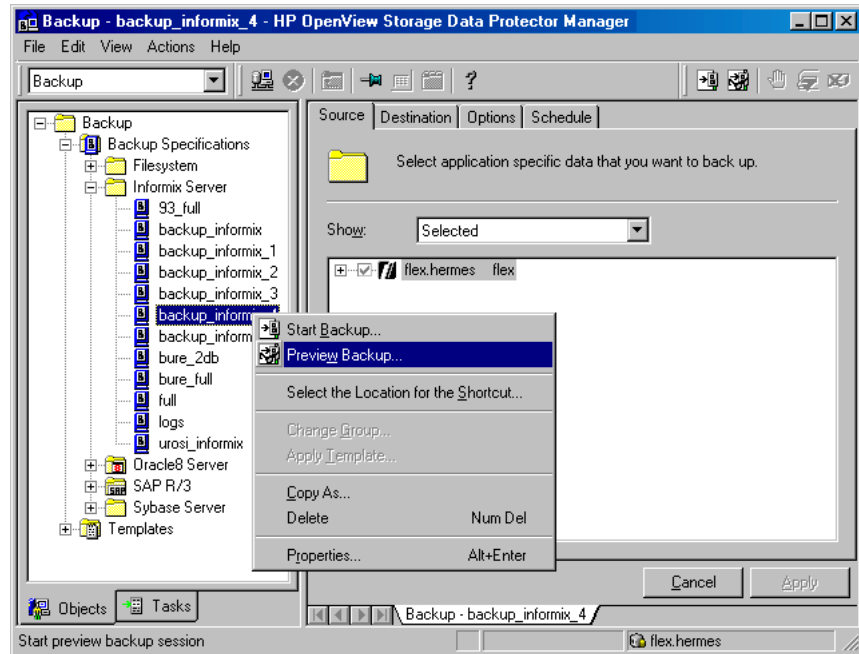
Test your backup specifications thoroughly by previewing them, then running them on file devices that are not NULL devices and then finally on the actual devices you intend to use. You can use either the Data Protector GUI or the Data Protector CLI.

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 `Informix Server` and then right-click the backup specification you want to preview.

Figure 5-21 Testing the Informix Backup Specification



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. See online Help for a description of these options.

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

Using the Data Protector CLI

The omnib Command

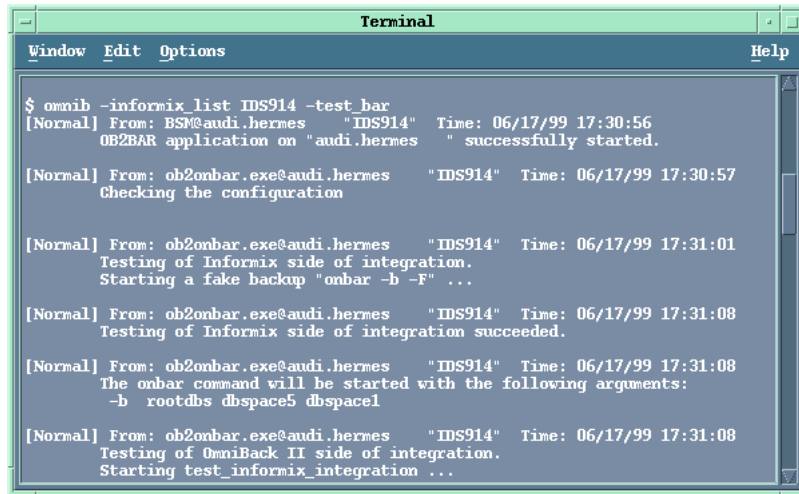
You can also perform the same test using the following Data Protector command in the Data Protector home directory (/opt/omni/bin - HP-UX and Solaris systems or /usr/omni/bin - other UNIX systems):

```
omnib -informix_list <backup_specification_name> -test_bar
```

where <backup_specification_name> is the name of the backup specification to be tested.

In the example, the *IDS914* backup specification was tested.

Figure 5-22 Testing the IDS914 Backup Specification



```
Terminal
Window Edit Options Help
$ omnib -informix_list IDS914 -test_bar
[Normal] From: BSM@audi.hermes "IDS914" Time: 06/17/99 17:30:56
OB2BAR application on "audi.hermes" successfully started.

[Normal] From: ob2onbar.exe@audi.hermes "IDS914" Time: 06/17/99 17:30:57
Checking the configuration

[Normal] From: ob2onbar.exe@audi.hermes "IDS914" Time: 06/17/99 17:31:01
Testing of Informix side of integration.
Starting a fake backup "onbar -b -F" ...

[Normal] From: ob2onbar.exe@audi.hermes "IDS914" Time: 06/17/99 17:31:08
Testing of Informix side of integration succeeded.

[Normal] From: ob2onbar.exe@audi.hermes "IDS914" Time: 06/17/99 17:31:08
The onbar command will be started with the following arguments:
-b rootdbs dbspace5 dbspace1

[Normal] From: ob2onbar.exe@audi.hermes "IDS914" Time: 06/17/99 17:31:08
Testing of OmniBack II side of integration.
Starting test_informix_integration ...
```

What Happens?

The given procedure performs a backup preview that:

1. Starts the Informix onbar command with the `-F` Informix fake option, which tests if the Informix database is correctly configured for backup. This command only tests the Informix part of configuration.
2. Starts the Data Protector
`/opt/omni/bin/utlins/test_informix_integration` or
`test_informix_integration_64bit` (HP-UX and Solaris systems)
or `/usr/omni/bin/utlins/test_informix_integration` or
`test_informix_integration_64bit` (other UNIX systems)
command, which tests:
 - Communication between the OnLine Server and Data Protector
 - The syntax of the Informix backup specification
 - If used devices are correctly specified
 - If the needed media are in devices

The `test_informix_integration` or `test_informix_integration_64bit` command only tests the Data Protector part of the configuration. You can run this command independently as shown in the following example:

**Test Informix
Integration
Example**

```
/opt/omni/bin/utilns/test_informix_integration ODS730  
InformixWhole,
```

where `ODS730` is the name of OnLine Server and `InformixWhole` is the name of the backup specification that you want to test.

Backing Up an Informix Database

In case of system failure, you can make a useful restore of your databases if *and only if* you have been making *regular* backups of the databases *and* logical-logs. This section describes how to configure and run backups of your Informix dbobjects.

To run a backup of Informix dbobjects, use any of the following methods:

Backup Methods

- Schedule the backup of an existing Informix backup specification using the Data Protector Scheduler. Refer to “Scheduling an Existing Backup Specification” on page 381
- Start an interactive backup of an existing Informix backup specification. You can start a backup using the Data Protector GUI or the Data Protector CLI. Refer to “Using the Data Protector GUI” on page 383 or “Using the Data Protector CLI” on page 384
- Start a backup using the Informix `onbar` command. Refer to “Using Informix Commands” on page 386
- Start a backup using the Informix `log_full.sh` script. Refer to “Using the Informix `log_full.sh` Script” on page 388

Backup Types

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

Table 5-4

Informix Backup Types

Backup Type	Description
Full	Full backup
Incr1	First incremental backup. Backs up changes since the last full (level 0) backup.
Incr2	Second incremental backup. Backs up changes since the last incremental (level 1) backup.

Refer to the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for details on these backup types and on the syntax of the `onbar` utility.

What Happens?

The following happens when you start an Informix backup:

1. Data Protector executes the `ob2onbar.exe` command on the OnLine Server. This command checks the configuration of the integration and starts the `onbar` command.
2. During the backup session, the `onbar` utility receives data from OnLine Server, which reads data from disk. The `onbar` utility then sends the data to Data Protector for writing to devices.

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

What OnLine Server Does Not Back Up

OnLine Server backs up all `dbojects` with the following exceptions:

- `Dbspace` pages that are allocated to OnLine Server but are not yet allocated to a `tblspace` extent
- Configuration files, such as the `sqlhosts` file
- Mirror chunks, if the corresponding primary chunks are accessible
- Blobs in `blobspaces` that are stored on optical platters
- Temporary `dbspaces`

What You Need to Back Up

In addition to OnLine `dbojects`, you should back up the files listed in the previous section, if you need them. You *must* back up the following files:

- `ONCONFIG` file, located in the `<INFORMIXDIR>/etc` directory.
- `sqlhosts` file, located in the `<INFORMIXDIR>/etc` directory.
- `oncfg_<INFORMIXSERVER>.<SERVERNUM>` file, located in the `<INFORMIXDIR>/etc` directory.
- emergency boot file, an Informix configuration file that resides in the `<INFORMIXDIR>/etc` directory and is called `ixbar.<server_id>`, where `<server_id>` is the value of the `SERVERNUM` configuration parameter.

IMPORTANT

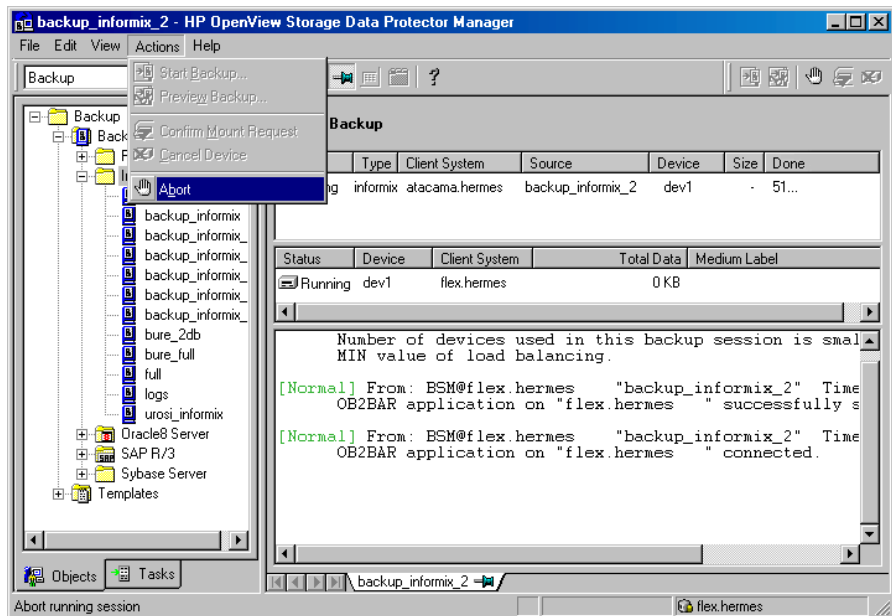
The Informix onbar utility does not back up these files. How often you need to back them up depends on how frequently changes are made to them. Back up the emergency boot file at least daily and always after you back up either a critical dbspace.

Aborting a Running Session

To abort a running Informix backup session, click **Abort** in the **Actions** menu, and then confirm the action.

In the following example, the backup session of the backup specification *InformixLogs* is being aborted.

Figure 5-23 Aborting an Informix Backup Session



IMPORTANT

The `BAR_RETRY` ON-Bar configuration parameter specifies how many times ON-Bar should retry a backup or restore operation if the first attempt fails. To successfully abort a backup or restore operation at the

first attempt, set the `BAR_RETRY` value to 0. Refer to the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for more information on this parameter.

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 Informix backup specification, follow the steps described in “Creating a Data Protector Informix Backup Specification” on page 361.

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 Informix 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 and click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options. See Figure 5-24 on page 382
6. Click OK to return to the Schedule property page.
7. Click Apply to save the changes.

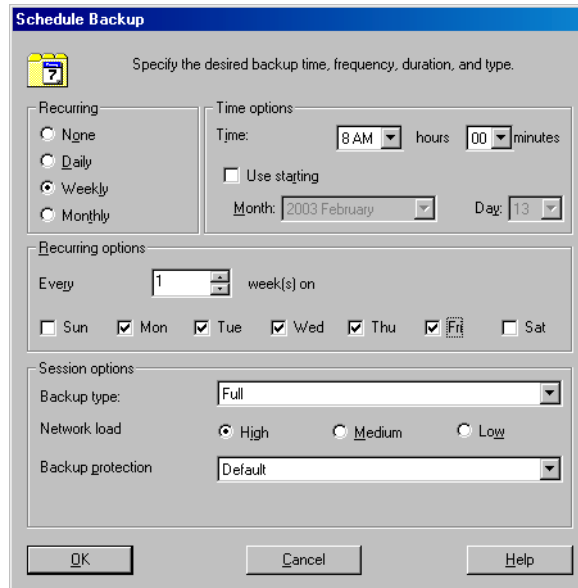
Scheduling Example

To schedule a backup specification called *InformixLogs* such that logical logs are backed up 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, select a date in the calendar and 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. Leave other options as default and click OK.

See Figure 5-24 on page 382.

Figure 5-24 Scheduling the InformixLogs 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.

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

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

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 an Informix dbject, perform the following steps in the HP OpenView Storage Data Protector Manager:

Running an Interactive Backup

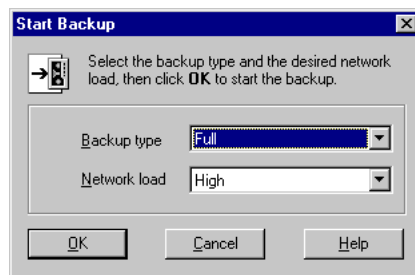
1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup, then Backup Specifications, and then Informix Server.
3. Select the backup specification you want to back up and click Start Backup in the Actions menu.

TIP

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

The Start Backup dialog box is displayed.

Figure 5-25 Starting the Backup of the InformixWhole Backup Specification



Select the backup type {Full | Incr1 | Incr2} and network load {High | Medium | Low}. See online Help for a description of these options.

4. Click OK.

Upon successful completion of the backup session a message confirming the success of the session is displayed.

Using the Data Protector CLI

The omnib Command

You can also start an interactive backup of an Informix dobject using the omnib command located in Data Protector home directory from any client in the Data Protector cell.

NOTE

The Data Protector home directory is /opt/omni/bin/ on HP-UX and Solaris and /usr/omni/bin/ on other UNIX clients.

```
omnib -informix_list <backup_specification_name>  
                [-barmode InformixMode]  
                [List_options]
```

```
InformixMode={full | inf_incr1 | inf_incr2}
```

A Data Protector Informix backup can be any of the following types:

NOTE

The Informix terms level-0, level-1, and level-2 backup are equivalent to Data Protector terms full, incr1, and incr2 backup, respectively.

Table 5-5

Informix Backup Types

Backup Type	Onbar Arguments	Description
Full	-L 0	Full backup
Incr1	-L 1	First incremental backup. Backs up changes since the last full (level 0) backup.

Table 5-5 Informix Backup Types

Backup Type	Onbar Arguments	Description
Incr2	-L 2	Second incremental backup. Backs up changes since the last first incremental (level 1) backup.

List_options can be one 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`.

`-no_monitor`

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

`-test_bar`

Tests both the Informix and the Data Protector parts of the backup specification as described in “Testing the Integration” on page 374.

Backup Examples To start a full backup of the Informix backup specification called *InformixWhole*, execute the following command in the Data Protector home directory:

```
omnib -informix_list InformixWhole -barmode full
```

To start an incremental backup of the Informix backup specification called *InformixIncr*, execute the following command in the Data Protector home directory:

```
omnib -informix_list InformixIncr -barmode inf_incr1
```

Using Informix Commands

This chapter assumes that you are familiar with both OnLine Server and the UNIX operating system.

You can start a backup of an Informix dobject from the client where the database is located using the Informix onbar utility.

See the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for more details on the onbar utility.

Before You Begin Before running a backup using the onbar command, ensure that you are logged in as root and execute the following commands:

- `export ONCONFIG=<onconfig_file>`
Name of OnLine Server ONCONFIG file, for example, ONCONFIG
- `export INFORMIXSQLHOSTS=<full_sqlhosts_file>`
Full pathname of the sqlhosts file, for example,
/applications/informix73/etc/sqlhosts
- `export INFORMIXSERVER=<INFORMIXSERVER>`
Name of OnLine Server, for example, ODS730
- `export INFORMIXDIR=<Informix_home_dir>`
Home directory of OnLine Server, for example,
/applications/informix73/
- `export OB2APPNAME=<INFORMIXSERVER>`
Name of OnLine Server, for example, ODS730
- `export OB2BARLIST=<backup_specification_name>`
Name of the backup specification to be used for the backup, for example, InformixWhole

Note that OB2APPNAME and OB2BARLIST are Data Protector-specific variables.

OnLine Server has to be in online or in quiescent mode to perform a backup. Once you start a backup, do not change the mode until the backup finishes; changing the mode terminates your backup. Only online dbspaces and blobspaces are backed up. To see which dbjects are online, type in the following command:

```
<INFORMIXDIR>/bin/onstat -d
```

where *<INFORMIXDIR>* is the home directory of OnLine Server.

Online Backups

The online mode is convenient if you want your OnLine Server to be accessible while you create the backup. An online backup might contribute to a loss of performance.

Quiescent Backups

The quiescent mode is useful when you want to eliminate partial transactions in a backup. A quiescent backup might not be practical if users need continuous access to OnLine Server databases.

Back Up Your Configuration Files

Keep a copy of your ONCONFIG, sqlhosts, and emergency boot files after you create a full backup. You need this information to restore OnLine Server dbjects.

Onbar Utility Backup Examples

To back up a list of dbspaces, proceed as follows:

```
onbar -b <dbspace_list>
```

To back up dbspaces *dbspace1* and *dbspace3*, type in the following command:

```
onbar -b dbspace1, dbspace3
```

To back up the current logical-log file and switch to the next logical-log file, use the *-c* option:

```
onbar -l -c
```

or

```
onbar -b -l -c
```

if you are using Informix 9.40.

Messages from the Data Protector backup session and messages generated by the Informix Integration Module are logged in the IDB. See “Monitoring an Informix Backup and Restore” on page 403 for additional information.

Using the Informix `log_full.sh` Script

The **`log_full.sh`** script is used to start backing up logical-log files when OnLine Server issues a log-full event alarm on the OnLine Server. See “On-Demand and Continuous Backups” for information on logical-log file backups.

To enable an Informix backup from the `log_full.sh` script, follow these steps:

- Backup Procedure** 1. Add the following line to the Informix ONCONFIG configuration file:

```
ALARMPROGRAM <INFORMIXDIR>/etc/log_full.sh.
```

where `<INFORMIXDIR>` is the home directory of OnLine Server.

2. If you do not have the Data Protector User Interface installed on the OnLine Server, then create an Informix backup specification to back up logical logs only, and edit the `<INFORMIXDIR>/etc/log_full.sh` script.

Add the following at the beginning of the file:

```
export OB2BARLIST=<backup_specification_name>
```

```
export OB2APPNAME=<INFORMIXSERVER>
```

where `<backup_specification_name>` is the name of the Informix backup specification and `<INFORMIXSERVER>` is the name of OnLine Server.

3. If you have the Data Protector User Interface installed on the OnLine Server then create an Informix backup specification to back up logical logs only.

NOTE

The Informix Enterprise Decision Server does not use the `log_full.sh` script.

On-Demand and Continuous Backups

To back up all logical-log files that are full and ready to be backed up, start an *on-demand* backup. An *on-demand* backup backs up all the full logical-log files, then stops at the current logical-log file.

You can also start a *continuous* backup on which OnLine Server backs up each logical-log file as it becomes full. The *continuous* backup process then waits until the next log is full. Use continuous logical-log backups, if you do not want to monitor the logical-log files.

By default, the ALARMPROGRAM configuration parameter is set so that ON-Bar performs continuous backups.

IMPORTANT

If you use *continuous* backups, ensure that a device is always available for the logical log backup process.

Refer to “Troubleshooting Logical Log Backups” on page 413 for information about troubleshooting logical log backups.

Examples

To make an *on-demand* backup of the logical-log files that are full (instead of a *continuous* backup that takes place every time a logical-log file fills), use the *-l* option as shown in the following example:

```
export OB2BARLIST=<backup_specification_name>
export OB2APPNAME=<INFORMIXSERVER>
onbar -l
```

To back up the current logical-log file and switch to the next logical-log file, use the *-c* option, as shown in the following example:

```
export OB2BARLIST=<backup_specification_name>
export OB2APPNAME=<INFORMIXSERVER>
onbar -l -c
```

or

```
onbar -b -l -c
```

if you are using Informix 9.40.

See the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for more information on *on-demand* and *continuous* backups.

Restoring an Informix Database

You can restore an Informix dobject in any of the following ways:

Restore Methods

- Using the Data Protector GUI. Refer to “Using the Data Protector GUI” on page 394.
- Using the `onbar` command on the OnLine Server. Refer to “Using Informix Commands” on page 398.
- Using the Data Protector CLI. Refer to Data Protector man pages for more information.

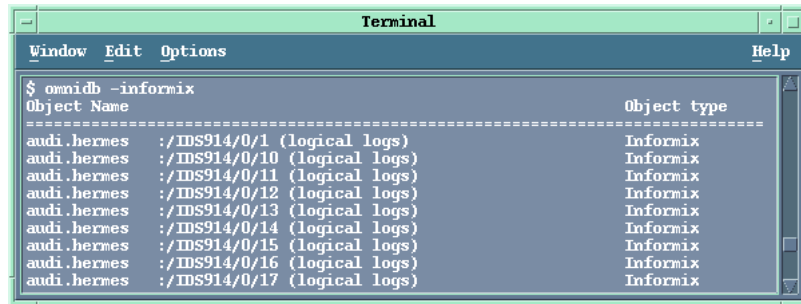
To restore a corrupted database, you need to find the right media and the session ID of the last backup session with a full backup. This and other information can be found using either the Data Protector `omnidb` command or the Data Protector GUI. See “The Data Protector `omnidb` Command” for instructions on using the Data Protector `omnidb` command to find the information needed to restore your data and “Finding Information for Restore” for finding the information using the Data Protector GUI.

The Data Protector `omnidb` Command

To find the information needed to restore your data, execute the following commands in the Data Protector home directory:

- `omnidb -informix`
to get a list of Informix objects.


Figure 5-26 List of Informix Objects



`omnidb -informix "object_name"`

to get details on a specific object, including the Session ID. Figure 5-27 shows how you get the objects pertaining to one of the objects specified in Figure 5-26.

Figure 5-27 Details About a Specific Session

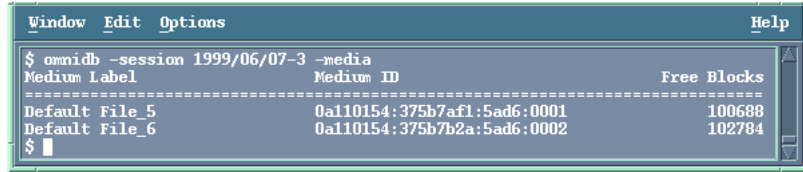


`omnidb -session SessionID -media`

to display media needed for restore. In the example depicted in Figure 5-28, media used for session 1999/06/07-3 are displayed.

Figure 5-28

Finding Media Needed for Restore



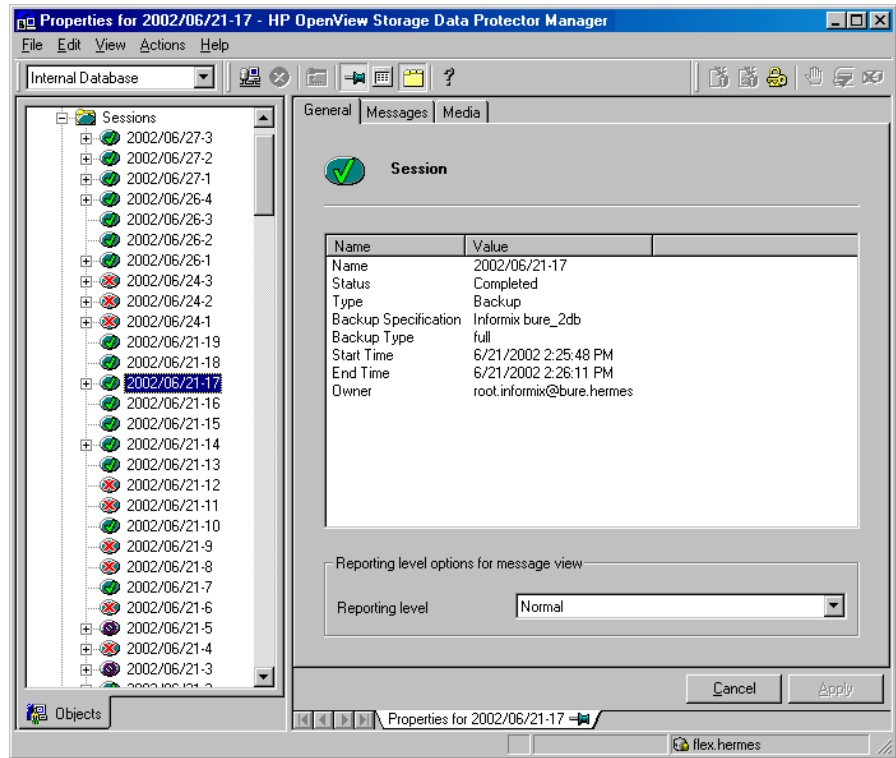
```
Window Edit Options Help
$ omnidb -session 1999/06/07-3 -media
Medium Label          Medium ID              Free Blocks
-----
Default File_5       0a110154:375b7af1:5ad6:0001    100688
Default File_6       0a110154:375b7b2a:5ad6:0002    102784
$
```

See the manual pages for detailed information on the omnidb command.

Finding Information for Restore

You can find information needed for restore in the HP OpenView Storage Data Protector Manager by clicking the Data Protector Internal Database Context and expanding either Sessions or Objects. The sessions are listed by date. Double-click a session to view session details:

Figure 5-29 Checking Session Details



IMPORTANT

The BAR_RETRY ON-Bar configuration parameter specifies how many times ON-Bar should retry a backup or restore operation if the first attempt fails. To successfully abort a backup or restore operation at the first attempt, set the BAR_RETRY value to 0. Refer to the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for more information on this parameter.

Before You Begin Restoring

If you are *not* using the Informix Enterprise Decision Server, shut down OnLine Server before you restore a root dbspace as follows:

Logged on to OnLine Server as user `informix`, type in the following command:

Integrating Informix and Data Protector

Restoring an Informix Database

```
<INFORMIXDIR>\bin\onmode -ky
```

where *<INFORMIXDIR>* is the home directory of OnLine Server.

If you are using the Informix Enterprise Decision Server, bring the Informix server in microkernel mode as follows:

1. Shut down all coservers:

```
xctl onmode -ky
```

2. Bring database server to microkernel mode:

```
xctl -C oninit -m
```

Note that if you intend to restore only non-critical Informix dbspaces (dbspace1, dbspace2, etc.), then OnLine Server can be online.

The following sections describe how to run a restore. Data Protector starts the `onbar` command under the account of the user running the restore. This should be the Informix user.

Using the Data Protector GUI

To run a restore, follow these steps in the HP OpenView Storage Data Protector Manager:

- Restore Procedure**
1. In the Context List, select Restore.
 2. In the Scoping Pane, expand Restore and then Informix Server, to get a list of OnLine Servers from which Informix dbobjects can be restored.
 3. Select the OnLine Server from which you want to restore.

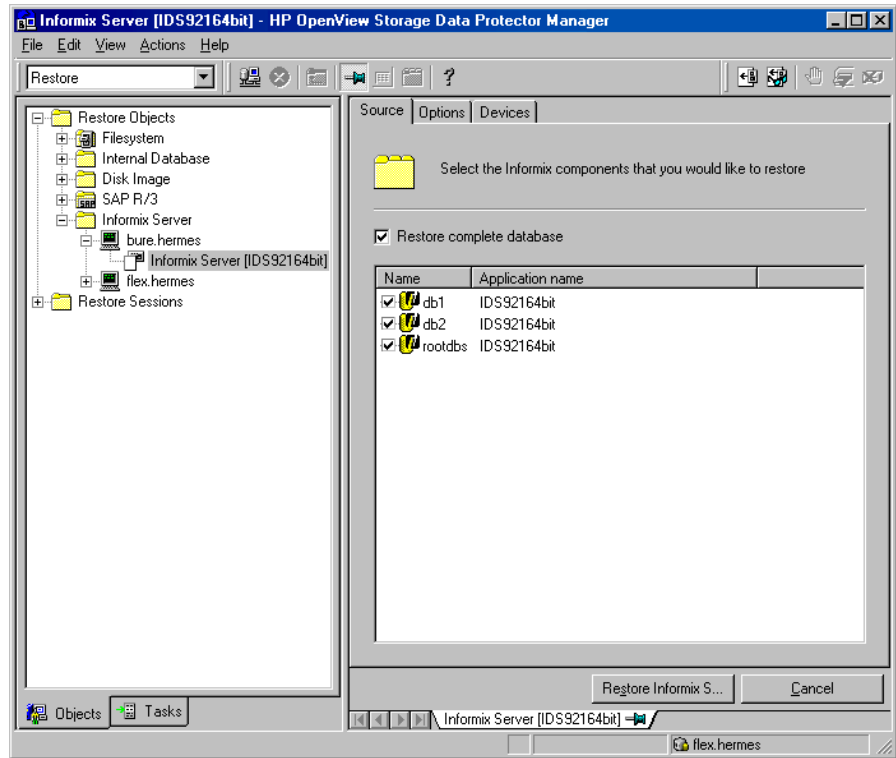
A list of OnLine Servers is displayed in the Results Area. Double-click the server from which you want to restore.

Select the dbobjects you want to restore or Restore complete databases to restore all backed up objects.

NOTE

To make an Informix whole-system restore, you must first select the Restore complete databases option.

Figure 5-30 Selecting Objects for Restore



4. Select **Options**, to specify the options needed for restore. These options are explained in “Restore Options” on page 395:

Restore Options

Backup Specification

Specifies a backup specification to be used to salvage logical-log files that are still on the disk before restoring. Note that this is not necessarily the backup specification used for the backing up.

User name, User group

User name and group. The `onbar` command is started under the account of the specified user.

Restore to client

Specifies the name of the original backup client. To restore to another client, specify the name of the other client.

Restore by log number

Restore all data up to the specific log number. If any logs exist after this one, the onbar utility does not restore them. This option invokes the onbar `-r -n last_log_number` command. Refer to the *INFORMIX OnLine Dynamic Server: Backup and Restore Guide* for details.

Restore by date

Indicates the date of the backup from which the restore is to be performed. This option invokes the onbar `-r -t time` command. Refer to the *INFORMIX OnLine Dynamic Server: Backup and Restore Guide* for details.

NOTE

You can browse your backup dates, regardless of backup type, using the Browse tab. The browse feature works only for backups of the current version of Data Protector. However, you can also enter other restore times, forcing a point in time restore to that particular time.

Restore the latest version

Restores the latest version of a backup.

Whole database restore

Searches the last whole-system backup and restores from that. This option invokes the onbar `-r -w` command. Refer to the *INFORMIX OnLine Dynamic Server: Backup and Restore Guide* for details.

NOTE

This option should only be used after a whole database backup. Data Protector does not automatically detect if you have a whole database backup.

5. Select the `Devices` tab to specify the devices from which you want to restore.

6. If the Informix Full or Whole restore are to be performed and the Informix server to be restored is in online mode, shut down the Informix server by issuing the following command on the Informix server that is to be restored:

```
onmode -ky
```

7. To start your restore session, proceed as follows:

Click the Start Restore button

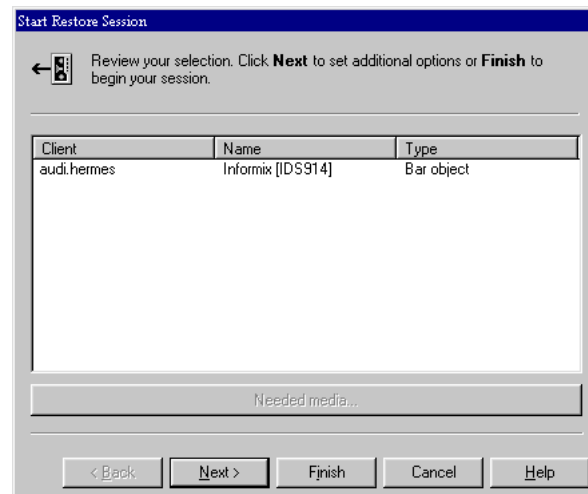
or

In the Actions menu, click Start Restore.

The Start Restore Session dialog box is displayed.

Figure 5-31

Starting the Informix Restore Session



Click Next, to select the Report Level and the Network Load of the restore session and then Finish to start the restore session.

8. Observe the session flow messages in the Data Protector Monitor.
The “Session successfully completed” message is displayed at the end of a successful session.

9. If the Informix Full or Whole restore was performed, put the Informix server back in online mode by issuing the following command on the Informix server that was restored, when the restore has finished:

```
onmode -m
```

What Happens?

The following happens when you start a restore using the Data Protector User Interface:

1. Data Protector executes the `ob2onbar.exe` command on the OnLine Server. This command starts the `onbar` restore command, with the specified options.
2. The `onbar` command contacts OnLine Server, which contacts Data Protector via XBSA and initiates a backup session to salvage logical logs.
3. During this backup session, OnLine Server reads data from the disk and sends it to the `onbar` utility, which sends the data to Data Protector for writing to the device.
4. The `onbar` command contacts OnLine Server, which contacts Data Protector via XBSA and initiates a restore session for the data selected for restore.
5. During this restore session, Data Protector reads the data from the device and sends the data to the `onbar` utility, which in turn sends the data to OnLine Server for writing to disk.
6. OnLine Server switches to quiescent mode.

Using Informix Commands

Before You Begin

Before you restore an Informix database instance using the Informix `onbar` command, ensure that you are logged in as `root`, and execute the following commands:

- `export ONCONFIG=<onconfig_file>`

Name of OnLine Server ONCONFIG file, for example, `ONCONFIG`

- `export INFORMIXSQLHOSTS=<sqlhosts_entry>`

Full pathname of the `sqlhosts` file, for example,
`/applications/informix73/etc/sqlhosts`

- `export INFORMIXSERVER=<INFORMIXSERVER>`

- Name of OnLine Server, for example, ODS730
 - `export INFORMIXDIR=<Informix_home_dir>`
Home directory of OnLine Server, for example,
`/applications/informix73/`
 - `export OB2APPNAME=<INFORMIXSERVER>`
Name of OnLine Server, for example, ODS730
 - `export OB2BARLIST=<backup_specification_name>`
Name of the backup specification used for salvaging logical logs and
not the one used for the backup, for example, Logs
- Note that OB2APPNAME and OB2BARLIST are Data Protector-specific variables.

Examples

The following are some examples of using the `onbar` command syntax for running restore.

Restoring Dbspaces and Logical Logs (Informix Full Restore)

If the Informix server to be restored is in online mode, shut down the Informix server by issuing the following command on the Informix server that is to be restored:

```
onmode -ky
```

To restore dbspaces and blobspaces as well as appropriate logical logs, use the `-r` option:

```
onbar -r
```

When the restore has finished, put the Informix server back in online mode by issuing the following command on the Informix server that was restored:

```
onmode -m
```

Restoring Dbspaces and Logical Logs (Informix Whole Restore)

If the Informix server to be restored is in online mode, shut down the Informix server by issuing the following command on the Informix server that is to be restored:

```
onmode -ky
```

To restore dbspaces and blobspaces as well as appropriate logical logs, use the `-r -w` options:

```
onbar -r -w
```

When the restore has finished, put the Informix server back in online mode by issuing the following command on the Informix server that was restored:

```
onmode -m
```

Restoring Dbspaces and Blobspaces Only

To restore dbspaces and blobspaces and not the logical log, use the `-r` and `-p` options:

```
onbar -r -p
```

Restoring a Particular Dbspace or Blobspace

To restore a specific dbspace, for example `dbspace_1`, use the following syntax:

```
onbar -r dbspace_1
```

Salvaging Logical Log files

If there has been a disk failure, salvage the logical-log files that are still on the disk with the following command before restoring your data from a backup:

```
onbar -l -s
```

See the *Informix-OnLine Dynamic Server: Backup and Restore Guide* for more details on the `onbar` command.

Messages from the Data Protector backup session and messages generated by the Informix Integration Module are logged in the IDB. Refer to “Monitoring an Informix Backup and Restore” on page 403 for additional information on monitoring restore sessions.

To Another OnLine Server

To restore to an OnLine Server other than the one from which the backup was made, proceed as follows:

1. Install and configure the Informix Integration Module to the other client.
2. Create an Informix user on the client to which you intend to restore.

Restoring to Another OnLine Server

3. Create an Informix database with the same Informix instance name and number as the original database by using the Informix `onmonitor` utility. Before going to the next step, ensure that OnLine Server is running.
4. Configure the Informix integration with the same OnLine Server name on the target client as was on the original client. Refer to Section “Configuring an OnLine Server” on page 353 for instructions.
5. Shut down the Informix database created in step 3.
6. Copy the main Informix configuration files (`ONCONFIG`, `sqlhosts`, emergency boot file, `oncfg_<INFORMIXSERVER>.<SERVERNUM>`) to the other client.
7. Modify the main client name in the Informix configuration files. This is necessary because the client name is changed when you restore to the other client.
8. Start a whole-system restore of the Informix objects from the Data Protector User Interface.

Using Another Device

Data Protector supports restore using a different device than the original one, which was used at backup time.

Restoring Using the Data Protector GUI

If you are performing a restore using the Data Protector GUI, refer to the “Restoring Under Another Device” section of the *HP OpenView Storage Data Protector Administrator’s Guide* for more information on how perform a restore using another device.

Restoring Using the Data Protector CLI or Informix Commands

If you are performing a restore using the Data Protector CLI or Informix commands, specify the new device in the `/etc/opt/omni/cell/restoreddev` 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.

Example

Suppose you have Informix 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 disaster recovery. Also see 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. Recover the operating system.
2. Install, configure, and initialize the database/application so that data on Data Protector media can be loaded back to the system. Consult database/application vendor documentation for a detailed procedure and steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in this chapter and 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 an Informix Backup and Restore

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

To monitor a currently running session, proceed as follows in the HP OpenView Storage Data Protector Manager:

Monitoring Procedure

1. In the `Context List`, select `Monitor`.

The progress and status of current sessions appear in the `Results Area`. You can sort the sessions by clicking the column headings in the `Results Area`.

If no sessions appear in the `Monitor` view, there are no sessions running. See the next section for instructions on how to view the finished sessions.

2. Double-click the running session you want to monitor.

At the end of the session a message is displayed indicating the success or failure of the session.

All actions and messages are logged to both Data Protector and Informix log files. Error messages from the last backup are logged in the `/var/opt/omni/log/informix.current.log` (HP-UX and Solaris systems) or `/usr/omni/log/informix.current.log` (other UNIX systems) file. Mount prompt requests are displayed on the Data Protector monitor.

When the onbar utility encounters an error or a condition that warrants a warning, it writes a message to the Informix ON-Bar message file. The full pathname of this file is specified in the `BAR_ACT_LOG` configuration parameter. For more information on this file, refer to the *INFORMIX-OnLine Dynamic Server: Backup and Restore Guide*.

Configuring the Integration as Cluster-Aware

Installation and Configuration

The Data Protector Informix integration can be configured in the MC Service Guard cluster. This means that either the Data Protector Cell Manager can be configured in a cluster, or Data Protector client can be configured in the cluster. Refer to *HP OpenView Storage Data Protector Concepts Guide* for more information on supported configurations.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* and “Configuring the Integration” on page 348 for information on how install and configure the Data Protector Informix MC Service Guard integration.

When configuring the Data Protector Informix integration, configure it only on one of the cluster nodes per one Informix server. Use the virtual hostname when configuring the integration. However, if the Cell Manager is outside the cluster, you need copy/append the Data Protector Informix configuration files to all other nodes *after the integration has been configured on one node*.

Copy the following directory(ies) to the same position on all other nodes:

```
/etc/opt/omni/informix/<informix_server_name>/
```

Append the following file to the same file on the same position on all other nodes:

```
/etc/opt/omni/informix/SERVERLIST
```

For information on the Data Protector Cell Manager package configuration (if you want to install and configure a Data Protector Cell Manager in the MC/SG cluster), refer also to the *HP OpenView Storage Data Protector Administrator's Guide*.

Backup and Restore

When creating a Data Protector Informix MC/SG cluster backup specification, always select the virtual hostname in the cluster and not a particular node.

Refer to “Backing Up an Informix Database” on page 378 for information on how to create a Data Protector Informix backup specification and to *HP OpenView Storage Data Protector Administrator’s Guide* for information on MC/SG cluster backing up specifics.

Refer to “Restoring an Informix Database” on page 390 for information on how to restore a Sybase database.

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.

Configuration Problems

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

1. Make a test backup of any filesystem on the problematic client.

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

2. Ensure that OnLine Server is up and running:

- a. Log on to OnLine Server as user `informix`

- b. Type in the following command:

```
<INFORMIXDIR>/bin/onstat -d
```

where `<INFORMIXDIR>` is the home directory of OnLine Server.

If OnLine Server is up and running, the `-- On-Line --` message is displayed.

If not, then start OnLine Server using the following command:

```
<INFORMIXDIR>/bin/oninit
```

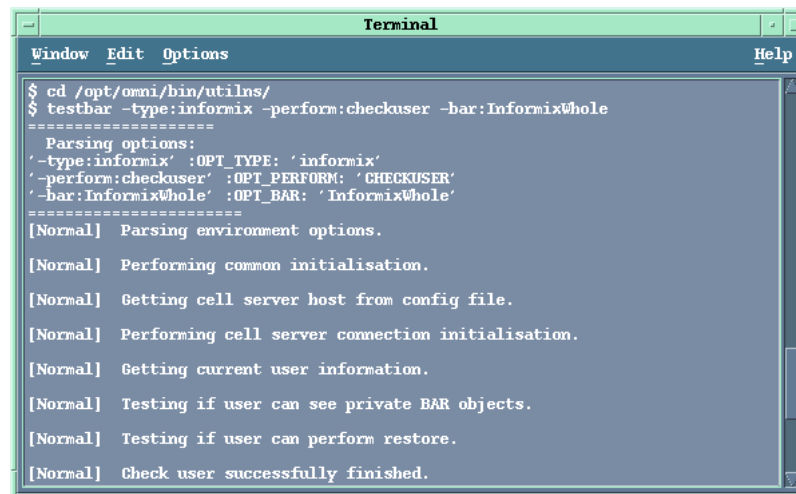
where <INFORMIXDIR> is the home directory of OnLine Server.

3. If you have any non-default Informix settings, then ensure that they are registered in the
 /etc/opt/omni/informix/<INFORMIXSERVER>/ .profile (HP-UX and Solaris systems) or
 /usr/omni/config/informix/<INFORMIXSERVER>/ .profile (other UNIX systems) file, where <INFORMIXSERVER> is the name of OnLine Server.
4. Examine system errors reported in /usr/omni/log/debug.log on the OnLine Server.
5. Test if the Informix user has the appropriate privileges in Data Protector. Log in as the Informix user, for example, as user informix, change to the /opt/omni/bin/utilns/ (HP-UX and Solaris systems) or /usr/omni/bin/utilns/(other UNIX systems) directory and run the following command on the OnLine Server:

```
testbar -type:informix -perform:checkuser  
-bar:InformixWhole
```

Figure 5-32

Checking the Informix User



In the preceding example, the user has all the appropriate rights for the backup specification called *InFormixWhole*.

If a user *ana* on OnLine 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.
```

Refer to “Configuring an Informix User in Data Protector” on page 350 for information about the right privileges.

Backup Problems

If you have problems backing up Informix dobjects, proceed as follows:

1. Make a test backup of any filesystem on the problematic client.

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

2. Ensure that OnLine Server is up and running:

- a. Log on to OnLine Server as user *informix*
- b. Type in the following command:

```
<INFORMIXDIR>/bin/onstat -d
```

where *<INFORMIXDIR>* is the home directory of OnLine Server.

If OnLine Server is up and running, the -- On-Line -- message is displayed.

If not, then start OnLine Server using the following command:

```
<INFORMIXDIR>/bin/oninit
```

where *<INFORMIXDIR>* is the home directory of OnLine Server.

3. Verify the configuration of your OnLine Server using the following command:

where *<INFORMIXSERVER>* is the name of OnLine Server.

```
util informix.exe -CHECKCONF <INFORMIXSERVER>,
```

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

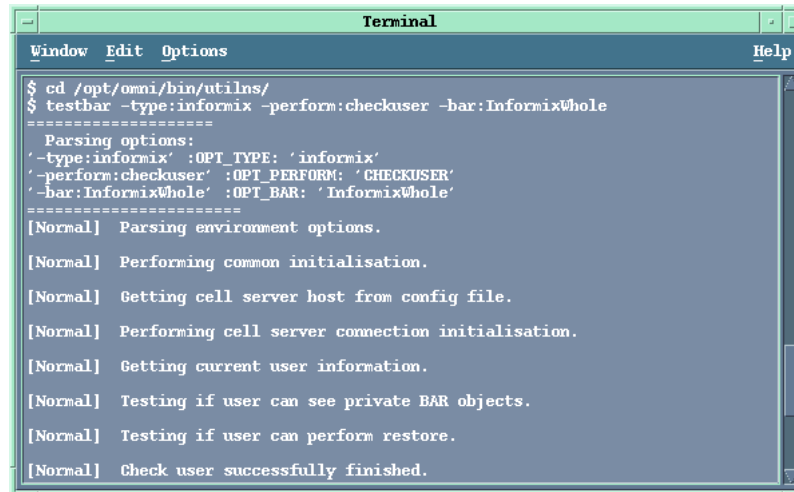
To get the error description, start the command,
`/opt/omni/lbin/omnigetmsg 12 <error_number>` (HP-UX and Solaris systems) or `/usr/omni/bin/omnigetmsg 12 <error_number>` (other UNIX systems).

4. Test if the Informix user has the right privileges in Data Protector. Log in as the Informix user, for example, as user `informix`, and run the following command on the OnLine Server:

```
/opt/omni/bin/utilns/testbar -type:informix
-perform:checkuser -bar:InformixWhole (HP-UX and Solaris systems)

/usr/omni/bin/utilns/testbar -type:informix
-perform:checkuser -bar:InformixWhole (other UNIX systems)
```

Figure 5-33 Checking the Informix User



In the preceding example, the user has all the appropriate rights for the backup specification named `InformixWhole`.

If a user `andrea` on OnLine 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.
```

Refer to “Configuring an Informix User in Data Protector” on page 350 for information about the right privileges.

5. Verify that the Data Protector Cell Manager is correctly set on the OnLine Server:

```
cat /etc/opt/omni/cell/cell_server (HP-UX and Solaris systems) or
```

```
cat /usr/omni/config/cell/cell_server (other UNIX systems) or
```

This command returns the Data Protector Cell Manager. In Figure 5-34, the configured Data Protector Cell Manager is called *toplarna.hermes*.

Figure 5-34

Verifying the Cell Manager



6. Verify that the `onbar` (or `onbar_d` for Informix 7.3x or later) command has the switch ownership (`s`) bit set and that it is owned by the Informix user, for example, `root:informix`.
7. Test the Data Protector Informix configuration as per instructions in “Testing the Integration” on page 374.

Example

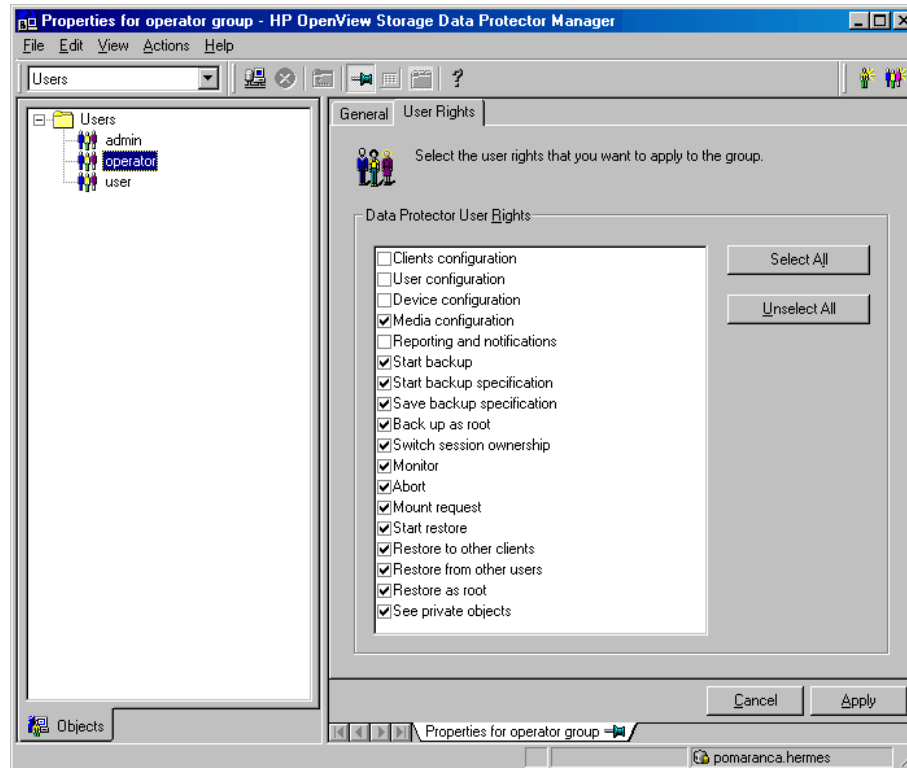
Run the following command to test the configuration of the backup specification called *InformixWhole*:

```
opt/omni/bin/omnib -informix_list InformixWhole -test_bar
```

- If the Informix part of the test fails, then proceed as follows:
Make a test run using the `onbar -F -b` option. If the test fails, then this is probably not a Data Protector problem. Refer to Informix manuals for further instructions.
- If the Data Protector part of the test fails then create an Informix backup specification to back up to a null or file device. If the backup succeeds, then proceed as follows:

- a. Verify that the owner of the backup specification is the Informix user, and if they are in the Data Protector operator or admin group.
- b. Ensure that the See private objects user right of the Data Protector operator group, which allows users to browse private objects, is selected:
 1. In the Context List, select Users.
 2. In the Results Area, right-click Operator and click Properties.

Figure 5-35 **Selecting the See Private Objects User Right**



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

- c. Create an Informix 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 an OnLine Server. See “Using Informix Commands” on page 386 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.

8. Test Data Protector data transfer using the testbar utility. Log in as the Informix user on the OnLine Server, change to the `/opt/omni/bin/utilns/` (HP-UX and Solaris systems) or `/usr/omni/bin/utilns/` (other UNIX systems) and proceed as follows:

```
testbar
-type: Informix
-appname: <INFORMIXSERVER>
-bar: <backup_specification_name>
-perform: backup
```

where `<INFORMIXSERVER>` is the name of OnLine 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` (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the OnLine Server.
9. If you have any non-default Informix settings, then ensure that they are registered in the `/etc/opt/omni/informix/<INFORMIXSERVER>/ .profile` (HP-UX

and Solaris systems) or

`/usr/omni/config/informix/<INFORMIXSERVER>/ .profile` (other UNIX systems) file where `<INFORMIXSERVER>` is the name of OnLine Server.

10. Start the backup directly from the OnLine Server. Refer to “Using Informix Commands” on page 386.

Troubleshooting the Informix Side

This chapter is not meant to teach you about OnLine Server. Following the given procedure might help you solve some Informix-related problems:

1. Check the following Informix files for error descriptions:

`bar_act.log`

`bar_dbg.log`

`online.log`

The locations of these files are specified in the Informix ONCONFIG file.

2. Start a backup, not using Data Protector:

Set the `BAR_BSALIB_PATH` shell variable to
`<INFORMIXDIR>/lib/ibsad001.sl`

Use the `onbar` command to start the backup.

Troubleshooting Logical Log Backups

Problem

Backup of Informix logical logs fails

Description

After the continuous backup of logical logs is done, the Backup Session Manager waits for a specified timeout for the next logical log to be backed up. If there is no new connection in the specified timeout, the Backup Session Manager completes the session and goes down. If Informix sends a request for the backup of the next logical log, Data Protector first checks if the Backup Session Manager and other processes are up and running. If the Backup Session Manager is up, then a request is sent to the Backup Session Manager to create a new backup object.

And if between these last two events the Backup Session Manager goes down because it didn't receive a new request, you receive a system error and a new session is not started.

Resolution Organize the backup of logical logs with more than two backup specifications.

Detailed Description The backup of logical logs is started when a logical log is full. At that time, Informix starts a script specified by `ALARMPROGRAM` configuration parameter in the `ONCONFIG` file. This script then starts the backup using the specified backup specification.

When the next logical log is full, then it restarts the `ALARMPROGRAM` script. This script will now start the backup using a different backup specification than the previous one and this way the new session will be started and the problem cannot appear.

Use at least 3 backup specifications, because during the backup of one logical log it can happen that Informix calls the `ALARMPROGRAM` script more than once.

About the Backup Specifications Backup specifications must be the same. You can use the same device in all the backup specifications.

IMPORTANT The backup specification names must be different.

The following is an example of an alarm script for 4 backup specifications named `BARLIST1`, `BARLIST2`, `BARLIST3`, and `BARLIST4`. The script automates logical log backups using event alarms from the database server. To install this script, add the following line to the `ONCONFIG` file:

```
ALARMPROGRAM <INFORMIXDIR>/etc/log_full.sh, where  
<INFORMIXDIR> is OnLine Server home directory.
```

NOTE

The Informix Enterprise Decision Server does not use the `ALARMPROGRAM` script for continuous logical log backup. `Onbar_d` starts the backup automatically (if the `LOG_BACKUP_MODE` parameter in the `ONCONFIG` file is set to `CONT`) and passes a request to workers which perform the backup.

Figure 5-36 Example of an Alarm Script

```
PROG=`basename $0`
Barlist=`cat /etc/opt/omni/informix/IDS914/barlist`
export OB2BARLIST=$Barlist
export OB2APPNAME=IDS914
USER_LIST=informix
BACKUP_CMD="/opt/omni/bin/omnib -informix_list $Barlist"
EXIT_STATUS=0

EVENT_SEVERITY=$1
EVENT_CLASS=$2
EVENT_MSG="$3"
EVENT_ADD_TEXT="$4"
EVENT_FILE="$5"

case "$EVENT_CLASS" in
    23)
        # onbar assumes no operator is present,
        # so all messages are written to the activity
        # log and there shouldn't be any output, but
        # send everything to /dev/null just in case
        $BACKUP_CMD 2>&1 >> /dev/null
        EXIT_STATUS=$?
        ;;

# One program is shared by all event alarms.  If this ever gets expanded
# handle more than just archive events, uncomment the following:
    *)
        #
        EXIT_STATUS=1
        ;;
esac

case "$Barlist" in
    "BARLIST1")
        echo BARLIST2 > /etc/opt/omni/informix/IDS914/barlist
        ;;
    "BARLIST2")
        echo BARLIST3 > /etc/opt/omni/informix/IDS914/barlist
        ;;
    "BARLIST3")
        echo BARLIST4 > /etc/opt/omni/informix/IDS914/barlist
        ;;
    "BARLIST4")
        echo BARLIST1 > /etc/opt/omni/informix/IDS914/barlist
        ;;
    *)
        echo BARLIST2 > /etc/opt/omni/informix/IDS914/barlist
        ;;
esac

exit $EXIT_STATUS
```


Restore Problems

If you have problems restoring Informix dbobjects, proceed as follows:

1. Examine system errors reported in the
`/var/opt/omni/log/debug.log` (HP-UX and Solaris systems) or
`/usr/omni/log/debug.log` (other UNIX systems) file on the OnLine Server.
2. Make a Data Protector test backup and restore of any filesystem on the problematic client.

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

3. Verify that the Data Protector Cell Manager is correctly set on the OnLine Server:

```
cat /etc/opt/omni/cell/cell_server (HP-UX and Solaris systems) or
```

```
cat /usr/omni/config/cell/cell_server (other UNIX systems).
```

This command returns the Data Protector Cell Manager. In the following example, the configured Data Protector Cell Manager is called `toplarna.hermes`.

Figure 5-37

Verifying the Cell Manager

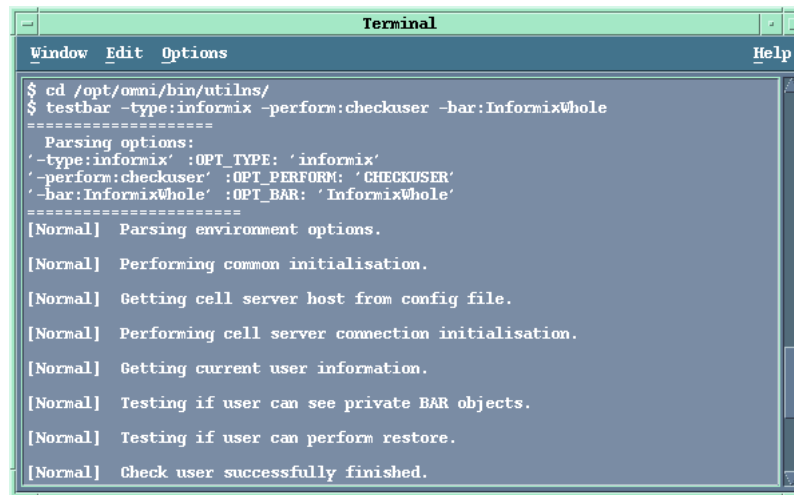


4. Verify that the user specified for the restore session is the Informix user, and that they are in the Data Protector operator or admin group.
5. Ensure that the See private objects user right of the Data Protector operator group is selected
6. Verify that the onbar (or onbar_d for Informix 7.3x or later) command has the switch ownership (s) bit set and that it is owned by the Informix user, for example, root:informix.

7. Test if the Informix user has the right privileges in Data Protector. Log in as the Informix user, for example, as user `informix`, and run the following command in the `/opt/omni/bin/utilns/` (HP-UX and Solaris systems) or in the `/usr/omni/bin/utilns/` (other UNIX systems) directory on the OnLine Server:

```
testbar -type:informix -perform:checkuser  
-bar:InformixWhole
```

Figure 5-38 Checking the Informix User



```
Terminal  
Window Edit Options Help  
$ cd /opt/omni/bin/utilns/  
$ testbar -type:informix -perform:checkuser -bar:InformixWhole  
=====  
Parsing options:  
'-type:informix' :OPT_TYPE: 'informix'  
'-perform:checkuser' :OPT_PERFORM: 'CHECKUSER'  
'-bar:InformixWhole' :OPT_BAR: 'InformixWhole'  
=====  
[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.
```

In the preceding example, the user has all the appropriate rights for the backup specification named `InformixWhole`.

If a user `andrea` on OnLine 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.
```

Refer to “Configuring an Informix User in Data Protector” on page 350 for information about the right privileges.

8. Ensure that the backup specification used for salvaging logical logs is properly configured. Note that this is *not* the same backup specification used to back up your data.

9. Test Data Protector data transfer using the testbar utility. Log in as the Informix user on the OnLine Server, change to the `/opt/omni/bin/utilns/` (HP-UX and Solaris systems) or `/usr/omni/bin/utilns/` (other UNIX systems) directory and proceed as follows:

```
testbar
-type:Informix
-appname:<INFORMIXSERVER>
-bar:<backup_specification_name>
-perform:restore
```

where `<INFORMIXSERVER>` is the name of OnLine Server and `<backup_specification_name>` the name of the Data Protector backup specification.

If the test is not successful, 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` (HP-UX and Solaris systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the OnLine Server.

Restoring to Another Client

Description

If you backed up your data to one client and exported the media and then imported them to another client in a different cell, the Data Protector session IDs of backup sessions may be changed in the IDB. However, the session IDs are not automatically changed in the Informix emergency boot file (`ixbar.<server_id>`, where `<server_id>` is the value of the `SERVENUM` configuration parameter).

Therefore, the restore of such objects may fail.

Action

Edit the emergency boot file to reflect the changed Data Protector session IDs. List the changed session IDs during the import procedure.

Information about backed up objects is stored in the emergency boot file in the format shown in Table 5-6.

Table 5-6 Emergency Boot File Format

```
ODS730  rootdbs  R  1  7  0  9  1999008018  1999-08-18  18:10:25  1
```

Information that makes up the Data Protector session ID is in columns 7 and 9. Column 9 represents the date and column 7 the unique session number.

For example, the session ID denoted in Table 5-6 is 1999/08/18-9. Note that the delimiter in the date field is “-” in the emergency boot file and “/” in the Data Protector session ID.

Also note that the value of the SERVERNUM configuration parameter is given in column 4.

Troubleshooting the Informix Side

Following the given procedure might help you solve some Informix-related problems:

1. Check the following Informix files for error descriptions:

```
bar_act.log  
bar_dbg.log  
online.log
```

The locations of these files are specified in the Informix ONCONFIG file.

2. Verify that the dbspaces you want to restore are offline in order to run a cold restore:
 - a. Log on to your OnLine Server as UNIX user `informix`
 - b. Type in the following command:

```
<INFORMIXDIR>/bin/onstat -d
```

where `<INFORMIXDIR>` is the home directory of OnLine Server.

3. Ensure that Informix configuration files (ONCONFIG, sqlhosts, emergency boot file, `oncfg_<INFORMIXSERVER>.<SERVERNUM>`) are not corrupted. If they are corrupted, restore them manually.

In This Chapter

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

This chapter is organized into the following sections:

- “Overview” on page 423
- “Prerequisites and Limitations” on page 425
- “Integration Concept” on page 427
- “Network Data Management Protocol (NDMP)” on page 429
- “Installing the NDMP Server Integration” on page 434
- “Configuring the Integration” on page 435
- “Network Appliance Configuration” on page 446
- “Backing Up the NDMP Server Data” on page 449
- “Restoring the NDMP Server Data” on page 454
- “Media Management” on page 458
- “The NDMP Related omnirc File Variables” on page 459
- “Troubleshooting” on page 460

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.
 - ✓ *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instruction on how to configure and run backups.

General NDMP Limitations

The following is a list of general NDMP Server integration limitations:

- Only filesystem backups are possible.
- A limited set of standard Data Protector media management functions is supported.
- Only static backup specifications are supported.
- It is not possible to have an NDMP backup session and a normal Data Protector session on the same medium.
- It is not possible to use the ACS or DAS Media Agents with NDMP. ACS and DAS libraries cannot be used with the NDMP client (Data Protector NDMP Media Agent client). Locally attached SCSI devices, however, are supported by the NDMP client.
- Maximum device concurrency is 1.
- Only the devices supported by Data Protector and the NDMP Server are supported. Please refer to the *HP OpenView Storage Data Protector Software Release Notes* for the device support matrix.

Prerequisites and Limitations

- 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.
- Localization for the NDMP Server specific messages is not possible.
- It is not possible to deselect a subtree/file of the selected tree to be restored.

**NetApp NAS
Device Limitation**

Direct access restore is supported only on the NDMP Server ONTAPP v6.1.x and higher.

**Celerra NAS
Device Limitation**

If you select directory restore on the Celerra NAS device using the direct access restore, only the selected directory without its contents is restored.


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. Refer to Figure 6-1 on page 428.

Figure 6-1 The NDMP Environment Configuration



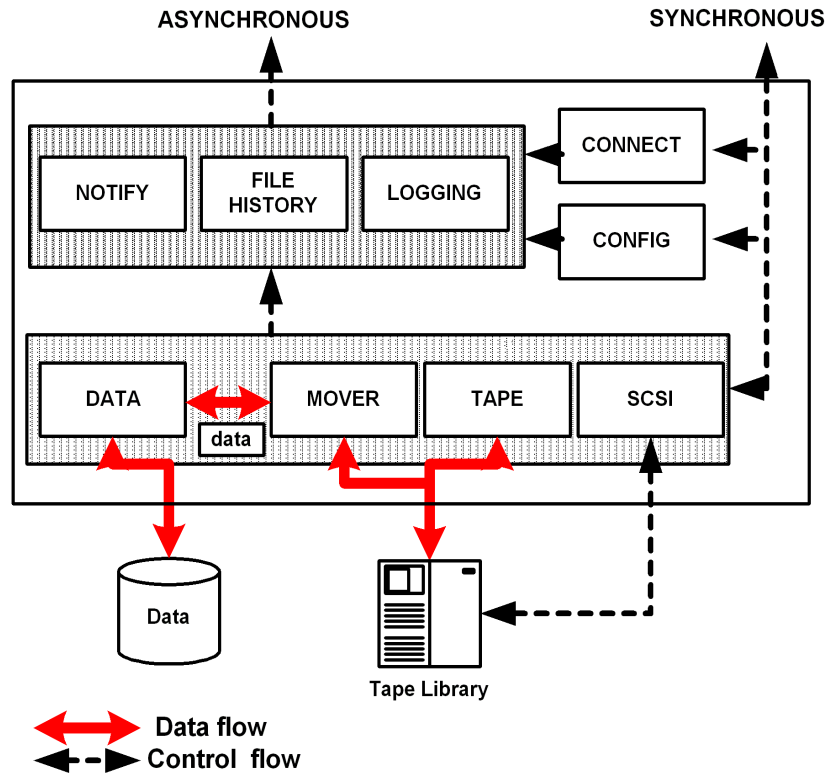
The important consequence of such a configuration is that the actual writing and reading of backed up data is not done by Data Protector, but by the NDMP Server itself in its own format. Since Data Protector cannot recognize this data format, some limitations apply for management of the NDMP media. For more information, refer to “Prerequisites and Limitations” on page 425.

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 Data Protector client because of the processing of catalog information.

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 6-2 shows the NDMP interfaces:

Figure 6-2 **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

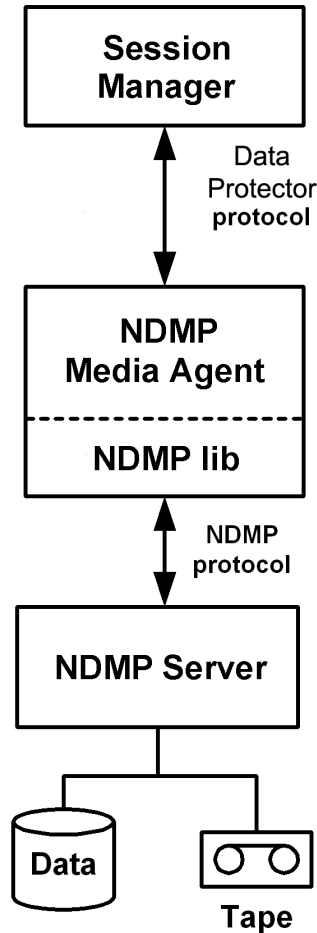
Using Data Protector, you can configure the NDMP Server not to create the file history information. Refer to “Direct Access Restore” on page 455 and “The NDMP Related omnirc File Variables” on page 459.

- LOGGING 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.

Refer to Figure 6-3 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 6-3 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. The NDMP Media Agent's functionality is limited to one provided by the NDMP protocol.

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.

Installing the NDMP Server Integration

- Prerequisites** Before installing the integration software, ensure that you have the Data Protector A.05.10 environment running.
- Installation** Install the NDMP Media Agent using remote installation. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on remote installation.
- You may install this software component on a system in the cell functioning as a Data Protector NDMP client (Data Protector NDMP Media Agent client). This software component is responsible for connection to the NDMP Server.
- Verifying the Installation** Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- What’s Next?** You have installed the Data Protector NDMP Server integration software. At this point, you are ready to proceed to the configuration procedure described in the next section.

Configuring the Integration

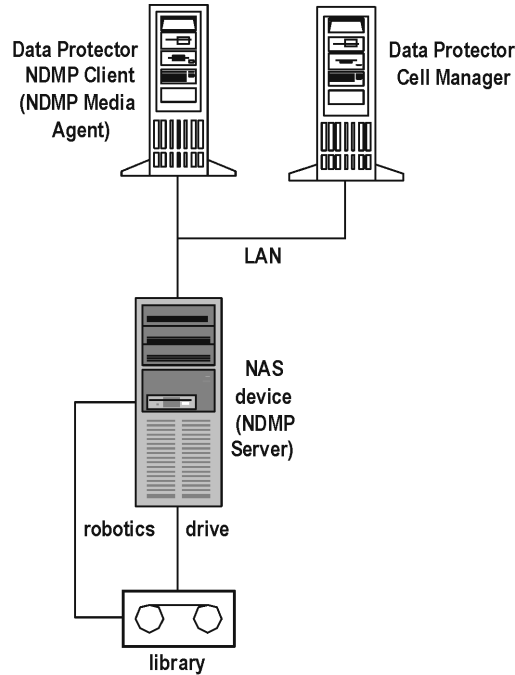
- Prerequisite** The NDMP Server must always have a tape drive, connected. The library robotics can be connected either to the NDMP Server or to a Data Protector NDMP client (Data Protector NDMP Media Agent client) depending on the chosen configuration. Refer to “Supported Configurations” on page 436.
- Supported Library Devices** Library devices can be controlled in two ways:
- The library robotics is attached to the NDMP server.
In this case, the library is entirely controlled through the NDMP Server. All library drives must be connected to the same NDMP Server and cannot be shared with other systems. Backup operations will work normally.
For details, refer to Figure 6-4 on page 436.
 - The library robotics is attached to a Data Protector NDMP client.
In this case, the library is controlled directly by Data Protector. The library drives may be connected to different systems, or shared between systems.
Refer to Figure 6-5 on page 437 and Figure 6-6 on page 438.
- 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.

Supported Configurations

Data Protector supports three NDMP Server integration configurations:


Figure 6-4

Supported Configurations-I



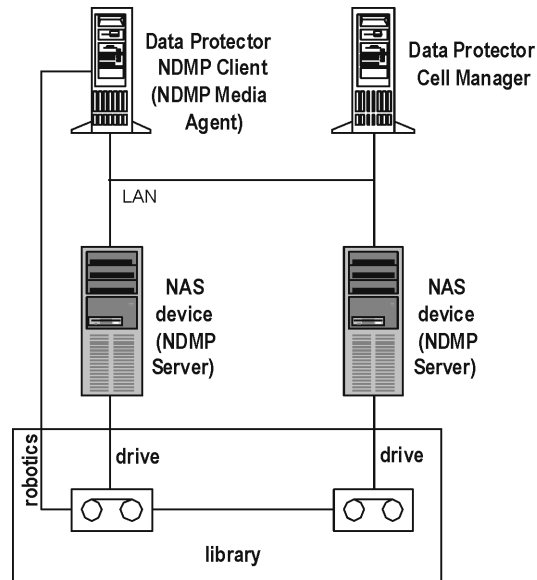
There is only one NDMP Server. The library drive is controlled through the NDMP Server TAPE interface and the library robotics is controlled through the NDMP Server SCSI interface. All drives in the library must be connected to the same NDMP Server in a direct SCSI connection.

Figure 6-5 Supported Configurations-II



Everything is as in the previous case, except that the library robotics is controlled directly from the NDMP client (Data Protector NDMP Media Agent client). This means that the support of different library devices depends on Data Protector.

Figure 6-6 Supported Configurations-III



There are two NDMP Servers in this configuration. The library drives are controlled through the NDMP Server TAPE interfaces, while the library robotics is controlled directly from the NDMP client (Data Protector NDMP Media Agent client) through locally attached SCSI devices. A part of the tape drives can be connected to the NDMP Server(s) in a direct SCSI connection, and the other part of the drives can be connected to other systems including the NDMP client. The support of different library devices depends 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.

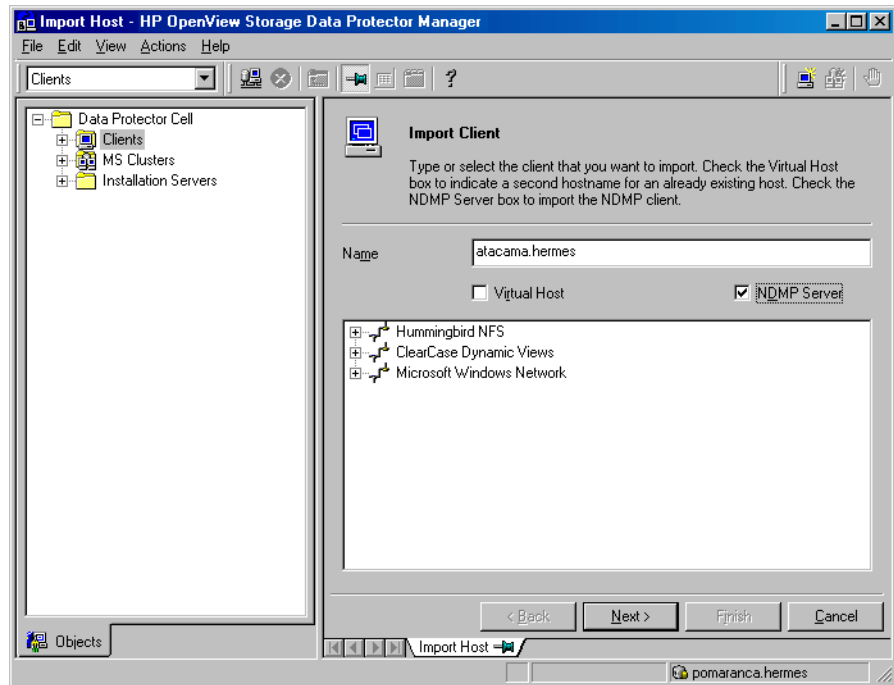
Importing the NDMP Server Host

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 NDMP Server.

Figure 6-7

Importing the NDMP Server Host



Click Next.

5. In the Import NDMP Host window, enter the specific import parameters:

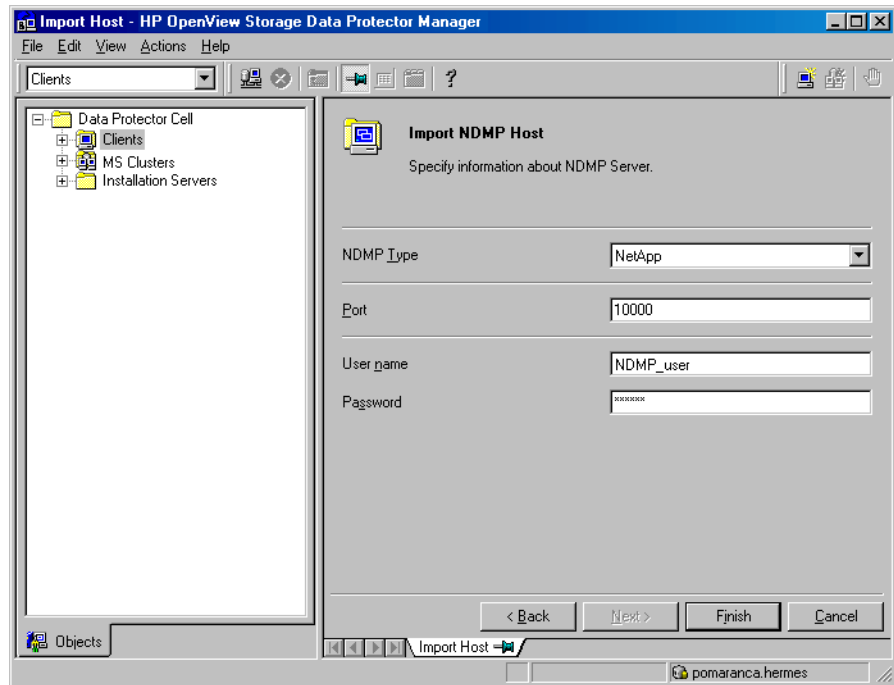
Integrating the NDMP Server and Data Protector Configuring the Integration

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.

Figure 6-8 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 a Data Protector NDMP client.

If the tape library unit (TLU) robotics is attached to a Data Protector NDMP client (*Data Protector NDMP Media Agent client*), it is possible to share drives in such a TLU. The TLU has to be configured for the Data Protector NDMP client with the TLU robotics attached. Drives can be shared among several NDMP Servers or among NDMP Servers and *Data Protector Media Agent clients*.

If the TLU robotics is attached to the NDMP Server, it is not possible to share drives in such a TLU and the TLU has to be configured for the NDMP Server.

Tape Library Units (TLU) Configuration Procedure

TLU Connected to a Data Protector NDMP Client To configure a TLU with robotics attached to a Data Protector NDMP 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 443.

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.

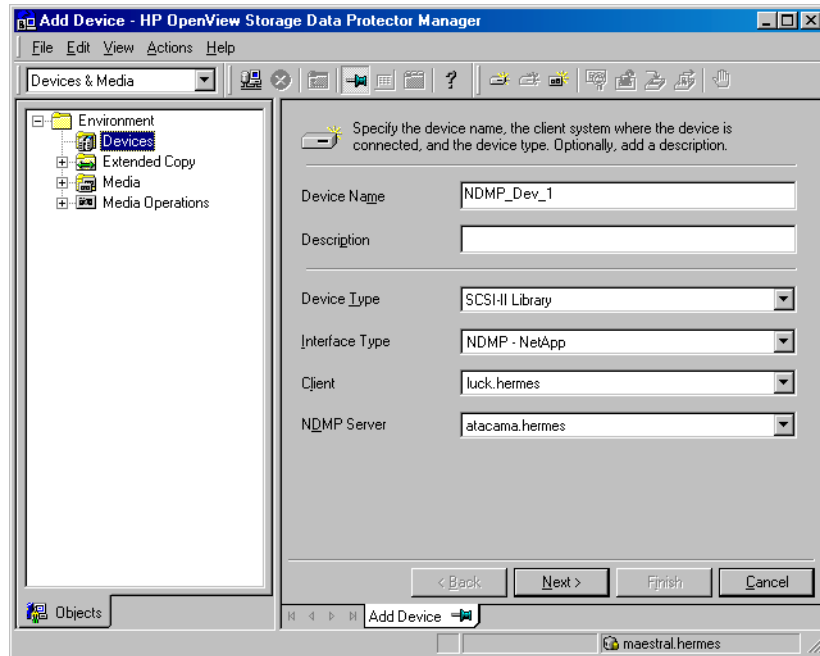
In the `Device Type` drop-down list, select `SCSI II Library`.

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 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.

Refer to Figure 6-9 for details.

Figure 6-9 Library Configuration



- Click Next.
4. Following the wizard, enter the required information about library SCSI ID and drive handling. Refer to “Network Appliance Configuration” on page 446 and “EMC Celerra Configuration” on page 448 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.
- If the library’s robotics is controlled by the NDMP Server, click Next.

If the library's robotics is controlled by the NDMP client (Data Protector NDMP Media Agent client), 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. Refer to "Network Appliance Configuration" on page 446 and "EMC Celerra Configuration" on page 448 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.
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.

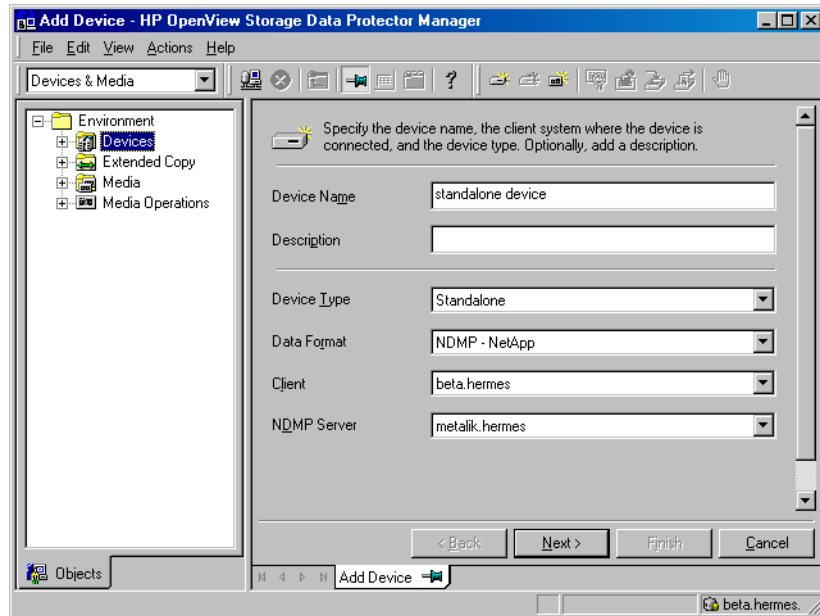
In the `Device Type` drop-down list, select `Standalone`.

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 6-10 Standalone Device Configuration



- Click Next.
- Follow the wizard and enter the information about the standalone device's SCSI address. Refer to “Network Appliance Configuration” on page 446 and “EMC Celerra Configuration” on page 448 for more information. Click Next.
 - 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.

- 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 6-1

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 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.

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.

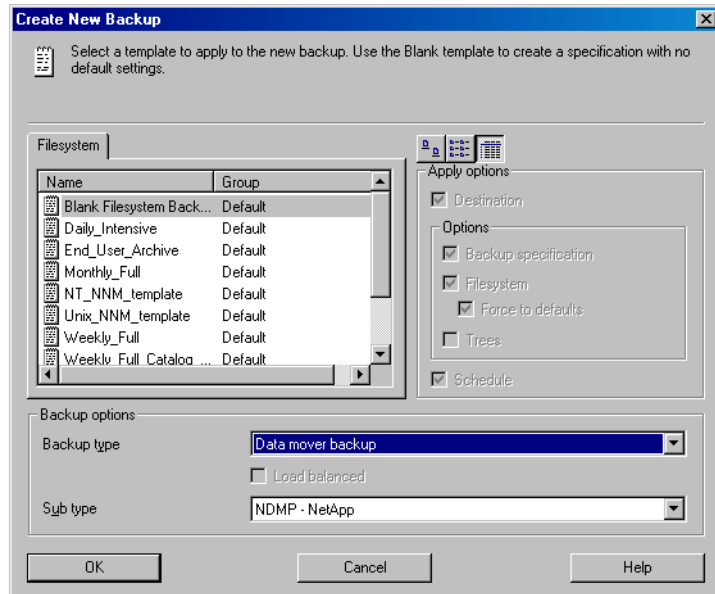
To back up a filesystem, perform the following steps:

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.

In the Backup type drop-down list, select Data mover backup. In the Sub type drop-down list, select NDMP-NetApp or NDMP-Celerra.

Refer to Figure 6-11 on page 450.

Figure 6-11 **Creating a New Backup**



NOTE

Load balancing is not supported, so the Load balanced check box is disabled.

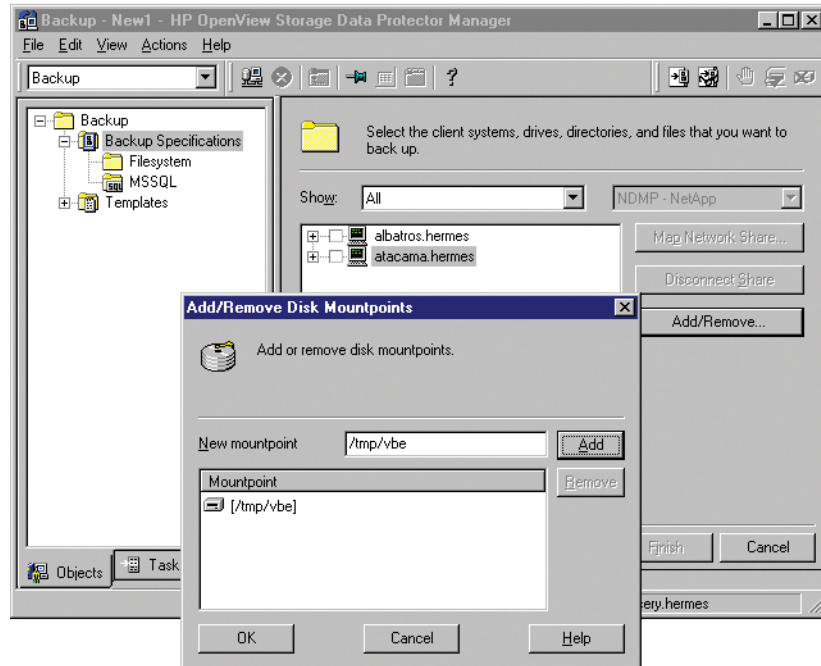
Click OK to start the Backup wizard.

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's mount point manually.

6. Specify the new directory (with the full path). After you have selected what you want to back up, click Add and then OK. Refer to Figure 6-12 on page 451.

Figure 6-12 Add/Remove Mountpoint of the NDMP Server



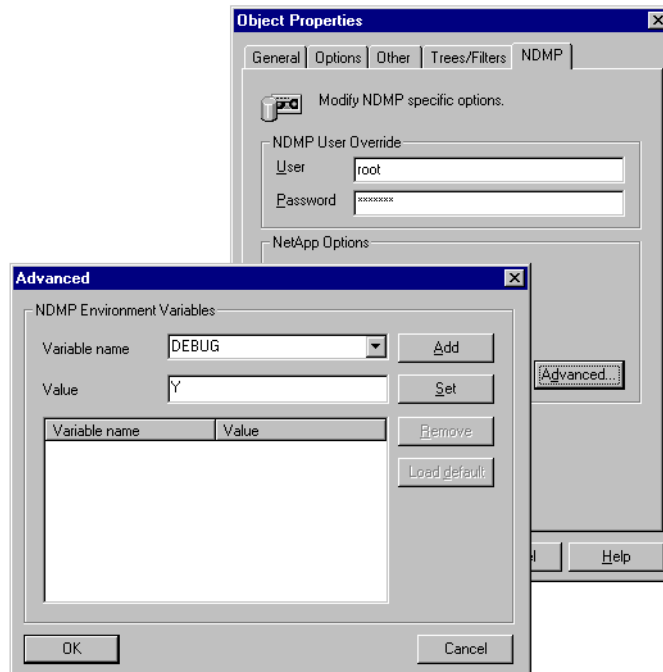
- Click Next to display the next page of the Backup wizard.
7. Select a device you want to use for the backup, and then click Next to proceed.
 8. Specify the Filesystem Options and Backup Specification Options. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on these options. Click Next.
 9. 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. See 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 options` dialog box. Refer to Figure 6-13.

Figure 6-13 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.

NDMP Environment Variables

The following tables represent the supported user defined NDMP environment variables for the NetApp and Celerra NAS device:

Table 6-2

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 6-3

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

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/file of the selected tree to be restored. This means that it is not possible to exclude subdirectories or files from a restore tree.
- Direct access restore on NetApp NAS device requires ONTAPP 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). Refer to “The NDMP Related omnirc File Variables” on page 459 for more information on how to switch the NDMP Server file history on or off.

Restore Procedure

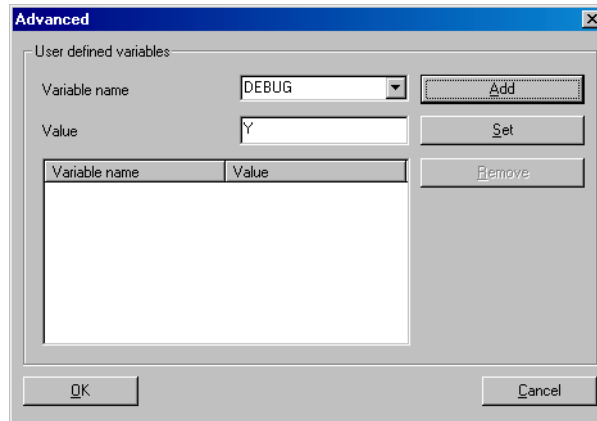
To restore a filesystem backed up from a NAS device through NDMP, perform the following steps:

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. Click the Advanced button to set the advanced options. Refer to Figure 6-14 on page 455.

Figure 6-14 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 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.

Integrating the NDMP Server and Data Protector Restoring the NDMP Server Data

Use of direct access restore functionality is defined through the `OB2NDMPDIRECT omnirc` file variable. By setting this variable you can enable or disable direct access restore. Refer to “The NDMP Related omnirc File Variables” on page 459 for details.

Prerequisite

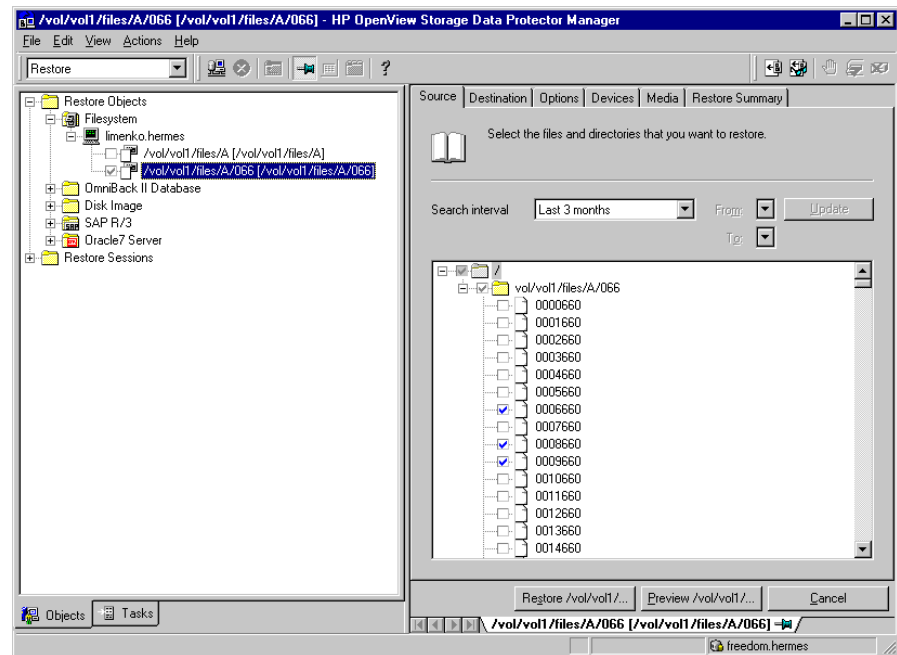
Direct access restore is possible only if the file history is switched on the NDMP Server during the backup session. Refer to “The NDMP Related omnirc File Variables” on page 459 for more information on how to switch the NDMP Server file history on or off.

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 more files in the Results Pane of the Data Protector Restore context.

NOTE

Direct access restore is supported for files only. Directories cannot be restored in direct access mode.

Figure 6-15 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.

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.

The NDMP Related omnirc File Variables

There are two NDMP related omnirc file variables, OB2NDMPFH and OB2NDMPDIRECT. For more information on the location and usage of the omnirc file refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

When the NDMP related variable OB2NDMPFH is set to Y, the NDMP Server file history is switched on. When it is set to N, the NDMP Server file history is switched off. This Data Protector variable overrides the file history settings on the NDMP Server every time a backup session is started.

The NDMP Server file history setting affects the following two aspects of the integration:

- 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.
- If it is switched off, you cannot browse and restore single files.

The two NDMP related omnirc file variables are:

- **OB2NDMPFH**

Default value: Y

If this variable is set to Y, the NDMP Server will create the file history information.

If this variable is set to N, the NDMP Server will not create the file history information.

- **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.

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, refer to “Network Data Management Protocol (NDMP)” on page 429.

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.

7**Integrating Network Node
Manager and Data Protector**

In This Chapter

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

It is organized into the following sections:

“Overview” on page 465

“Prerequisites and Limitations” on page 466

“Integration Concept” on page 467

“Installing the NNM Integration” on page 469

“Configuring an NNM Backup” on page 470

“Backing Up an NNM Database” on page 474

“Restoring NNM” on page 477

“Monitoring an NNM Backup and Restore” on page 479

“Troubleshooting” on page 480

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.

- 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. See 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.

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 Cell Manager. 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.
 - ✓ *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.

Limitations

See the *HP OpenView Storage Data Protector Software Release Notes* for a list of general Data Protector limitations. This section describes limitations 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.

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`.

Installing the NNM Integration

Prerequisites

There are no prerequisites to installing the Data Protector NNM Integration Module. The Data Protector Disk Agent must be installed on the NNM systems you want to back up but that installation can be done concurrently with the NNM Integration Module installation.

Installation

Install the Data Protector NNM integration software on your NNM systems either locally, from the CD-ROM, or remotely, using the Data Protector GUI.

You must install these Data Protector software components on the NNM system:

- HP OpenView NNM Backup Integration
- Disk Agent

Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

You have now installed the Data Protector NNM integration software on the NNM system. The NNM system is also a Data Protector client (because you have also installed the Data Protector Disk Agent). At this point, you are ready to proceed to creating a backup specification as described in the section, “Creating a Backup Specification” on page 471.

Configuring an NNM Backup

To configure an NNM backup, perform the following steps:

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

See 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, refer to “Creating a Backup Specification” 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. In the next step of the wizard, select the device to be used for backup. Click `Next`.
 7. Follow the wizard to define options, devices, 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 Data Protector online Help and the *HP OpenView Storage Data Protector Administrator's Guide* for backup options common to all objects. See “NNM Backup Options” on page 473 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 `/etc/opt/omni/datalists/<Backup_Specification_Name>` file on HP-UX and Solaris Cell Manager systems.

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.

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 475.
- Start an interactive backup of the NNM backup specification. See “Starting an Interactive Backup” on page 476.

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

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.

Backing Up an NNM Database

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 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 see 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.

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 /var/opt/OV/tmp/ovpause.lock 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 /var/opt/OV/tmp/ovpause.lock 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.

Backup/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.

Troubleshooting

- ✓ 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 install, configure, and use the Data Protector Lotus Notes Integration. It explains the concepts and methods that you need to understand to backup and restore Lotus Domino R5 Server.

The chapter is organized into the following sections:

“Overview” on page 485

“Prerequisites and Limitations” on page 488

“Integration Concept” on page 489

“Installing the Lotus Notes Integration” on page 492

“Configuring the Integration” on page 494

“Backing Up Lotus Domino R5 Server” on page 506

“Restoring Lotus Domino R5 Server Data” on page 517

“Monitoring a Lotus Domino Server Backup and Restore” on page 523

“Troubleshooting” on page 526

Overview

The Data Protector integration with Lotus Domino R5 Server allows you to perform offline as well as online backups. In order to enable a recovery from an online backup, the respective Lotus Domino R5 Server has to be set to use transactional logging. This way the transactions are stored to the transaction log directory and can be used to apply or undo database transactions during database recovery.

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 minimal impact on database performance.

The integration also provides you with features such as library support, parallel backups, and media management for backup and restore.

Data Protector backs up all types of databases (NSF, NTF, and BOX). Full and incremental backups are possible on offline as well as online databases. You are able to back up specific database or databases, or the whole server (all databases under the Lotus Domino R5 Server).

Database restore is possible even if Lotus Domino R5 Server is running, therefore the restore of specific database will not have an impact on other databases currently in use. There is also the possibility to perform a recovery to a specific point in time on a given database or on all databases under a specific server.

Lotus Notes Integration Agent

Data Protector Lotus Notes Integration Agent helps to protect and manage Domino Server data by making it easy to perform the following actions:

- Online backup of the whole Lotus Domino R5 Server or specific database
- Backup of archived transaction logs when archive logging is in effect
- Backup of the currently filling transaction log file if the Lotus Domino Server 5.0.4 or later is installed
- Centralized, online, full and incremental backup of Lotus Domino R5 databases
- Maintain multiple versions of Lotus Domino databases backups

Overview

- Automate scheduled backups
- Restore without performing a recovery
- Restore of backup versions of a Lotus Domino R5 database and apply changes made since the backup from the transaction log
- Restore Lotus Domino R5 databases to a specific point in time or to the latest possible consistent state
- Recover to same or different Lotus Domino R5 Server
- Restore database to other Lotus Domino R5 Server location than originally backed up from
- Automatic restoration of archived transaction logs in case of recovery

The Lotus Notes Integration Agent provides online and offline backups of Notes databases and transaction logs. The Lotus Notes Integration Agent supports two types of backup:

1. Full backup

Performs a full backup of specified Lotus Domino R5 Server databases. In case archive transaction logging is enabled, the full backup of all archived transaction logs is taken, including the current filling transaction log, which is not yet marked as 'ready to be archived'.

2. Incremental backup

If the data changed from the last backup is more than specified in the `Amount of log options`, performs a full backup of specified databases; otherwise, the specified database is skipped. In case when archive transaction logging is enabled, the full backup of all archived transaction logs is also taken.

Advantages

Using Data Protector together with Lotus Domino R5 Server offers several advantages over using Lotus Domino R5 Server alone:

- Central Management for all backup operations:
The administrator can manage backup operations from a central point.

- **Media Management:**

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

- **Scheduling:**

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

- **Device Support:**

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

- **Reporting:**

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

- **Monitoring:**

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

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

Prerequisites and Limitations

This is a list of prerequisites and limitations for the Data Protector Lotus Notes integration:

- You need a special license to use the Data Protector Lotus Notes integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for detailed information about Data Protector licensing.
- Before you begin, ensure that you have correctly installed and configured Lotus Domino R5 Server and Data Protector systems. 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, limitations, 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.
- It is assumed that you are familiar with Lotus Domino R5 Server administration and the basic Data Protector functionality.
- Only Lotus Notes R5 database is supported. You cannot backup Lotus Notes database R4 format, because Lotus Domino R5 API does not provide backup and restore functions for R4 database formats. For complete details on upgrading to Lotus Domino Release 5, please refer to *Lotus Domino R5 Administration Guide*.

Integration Concept

Data Protector Lotus Notes integration provides efficient online backup, restore and recovery of Lotus Domino R5 Server. It uses new Lotus C API to allow third party applications to perform online backups and restores.

The central component of the Data Protector Lotus Notes Integration is the Data Protector `ldbar.exe` executable, which is installed on the Lotus Domino R5 Server system and which controls the activities between Lotus Domino R5 Server and Data Protector backup and restore processes.

From the perspective of the Lotus Domino R5 Server, Data Protector is seen as a media management software. On the other hand, the Lotus Domino R5 Server is a Data Protector client from the Data Protector Cell Manager's point of view.

Backup Flow

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

The Data Protector Backup Session Manager reads the backup specification and starts the `ldbar.exe` command on the Lotus Domino R5 Server system.

The `ldbar.exe` reads data from the Lotus Domino R5 Server and passes it to the Data Protector Media Agent.

Lotus Domino Server databases are backed up in parallel depending on the sum of all concurrencies for individual device defined in the backup specification.

Backup session messages are sent to the Backup Session Manager, which then writes the messages and information regarding the respective session to the Data Protector database.

The two types of backup supported by the Data Protector Lotus Notes integration, are **Full** and **Incremental**.

Full backup includes all backup objects specified in the backup specification regardless of whether they have changed since the last backup. The incremental backup performs a full backup of specified databases if the data changed from the last full backup is bigger than specified in the `Amount of Log` options; otherwise the specified databases are not backed up.

Integrating Lotus Domino R5 Server and Data Protector
Integration Concept

There is only one level of incremental backup. It references the previous full or incremental backup, whichever was performed last.

Restore Flow

Using the Data Protector User Interface, you define which objects and objects versions to restore. The Restore Session Manager is invoked, which then starts the `ldbar.exe` with specific restore parameters. The `ldbar.exe` passes the information about the objects and backup versions on to the Lotus C API. Media Agents are started by `ldbar.exe`, and data flows from the media to target Lotus Domino R5 Server. Refer to Figure 8-1.

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

Figure 8-1 Data Protector Lotus Notes Integration Concept

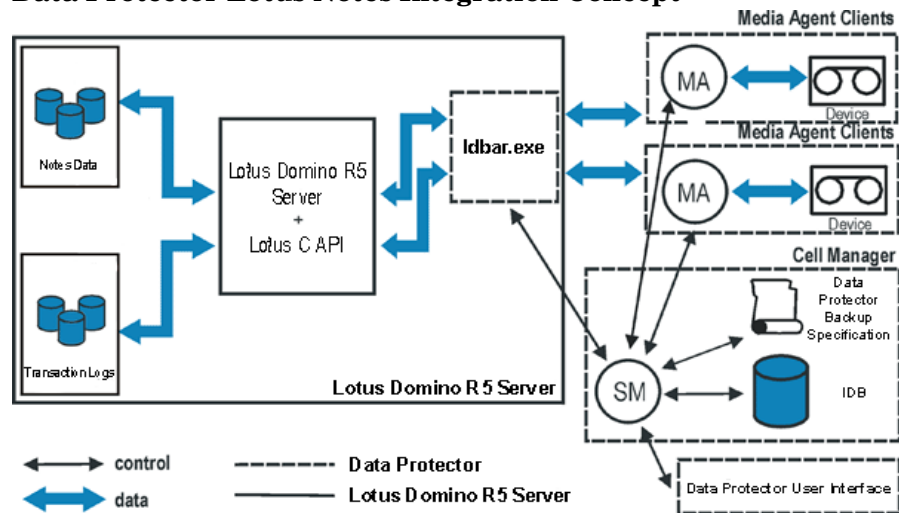


Table 8-1

Legend:

SM	Data Protector Session Manager that is Data Protector Backup Session Manager during a backup and Data Protector Restore Session Manager during a restore.
MA	Data Protector Media Agent

Table 8-1

Legend:

Lotus C API	The Lotus defined interface that enables data transfer between Data Protector and Lotus Domino R5 Server.
Notes Data	The Notes database is the basic component of a Notes application. It is a repository where users create, update, store, and track documents in various formats.
Transaction Logs	Domino supports transaction logging and recovery by capturing database changes and writing them to the transaction log.

Installing the Lotus Notes Integration

Install Data Protector Lotus Notes integration software on your Lotus Domino R5 Server either locally, from the CD-ROM, or remotely, by using the Data Protector User Interface.

Once installed and configured, the system also becomes a Data Protector client.

Install the following Data Protector software components:

- Lotus Notes Integration
- Disk Agent

Data Protector requires a Disk Agent to be installed on backup servers (clients with filesystem data to be backed up). It is important to install the Disk Agent for two reasons:

- ✓ It is recommended that you configure and run a test filesystem backup of the Lotus Domino R5 Server using Data Protector. By doing this, you will check whether the Lotus Domino R5 Server system and the Data Protector Cell Manager can communicate properly.
- ✓ It is also recommended to run a filesystem backup of important Domino R5 Server data that cannot be backed up using Lotus Notes Integration agent. These are so called non-database files, which need to be backed up to provide a complete data protection solution for a Domino R5 server. Following are examples of Domino non-database data:

- notes.ini
- desktop.dsk
- all *.id files

To backup this data and ensure complete data protection use Data Protector filesystem backup solutions.

- Media Agent (if you have devices connected to the system)
- User Interface

NOTE

If you are performing a local installation of the Data Protector software on the Lotus Domino R5 Server system, you have to select the Lotus Notes Integration software component during the setup procedure. Other required components are selected by default.

Verifying the Installation

Once the installation has completed, you can verify it. For the procedure, refer to “Verifying Data Protector Installation” in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

What’s Next?

If you have installed the Data Protector Lotus Notes integration software on the Lotus Domino R5 Server system, then you are ready to proceed to the configuration procedure described in the following section.

Configuring the Integration

The configuration is a set of procedures needed after the installation of the Data Protector Lotus Notes integration software. It consists of the following:

1. “Configuring the Lotus Domino R5 Server” on page 494
2. “Configuring Data Protector Lotus Notes Integration” on page 497

Configuring the Lotus Domino R5 Server

The configuration is performed using the ‘Lotus Domino Administrator’ on the Lotus Domino R5 Server computer. The same can be achieved with ‘Web Administrator’ or by editing the `notes.ini` file directly.

To configure Lotus Domino R5 Server, you have to enable transaction logging with archive transaction log style. No transaction logging is set as the default mode for Lotus Domino R5 Server. There are also two transaction logging styles.

In case you select circular logging, the transaction log files are automatically overwritten, when the disk space available for transaction log files is reached. If turned on, this option reduces disk storage space requirements, but does not allow you to perform incremental backups or use the database recovery feature.

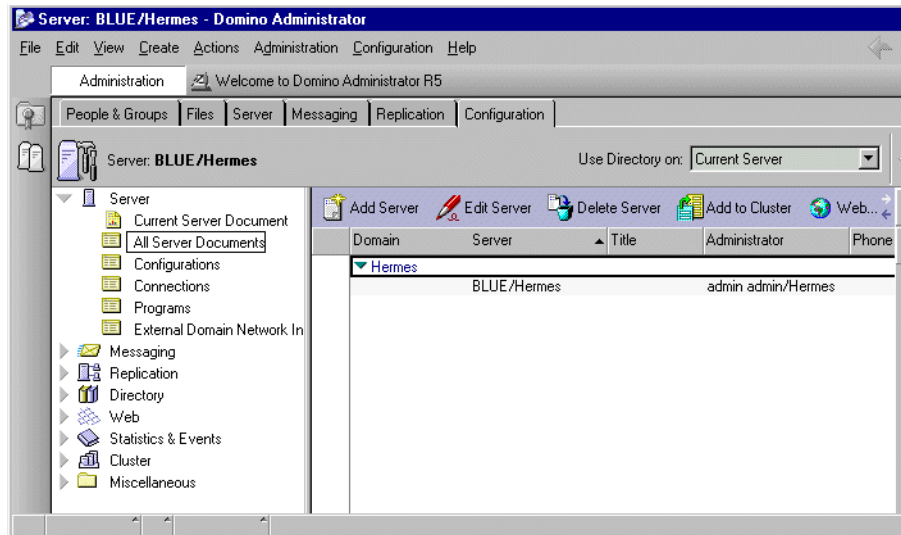
IMPORTANT

To perform incremental backups and archive log files, the transaction logging has to be set to archive logging style.

Proceed as follows to enable transaction logging and turn the circular logging style off.

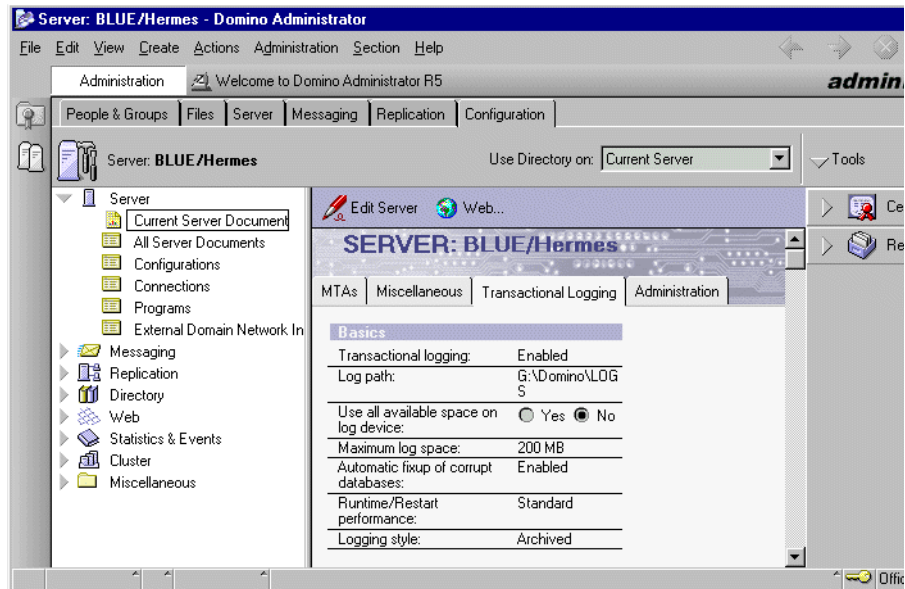
1. Start the Lotus Domino Administrator.
2. Log on to the Domino R5 Server, select Configuration tab and expand Server. Move to All Server Documents and select the Domino R5 Server you want to edit.

Figure 8-2 **Browsing Lotus Domino R5 Server**



3. Select the Transactional Logging tab, and set appropriate values. Save settings.

Figure 8-3 **Enabling Archive Transactional Logging Style**



4. Restart the server for changes to take effect.

For complete details on transactional logging please refer to *Domino R5 Administration Help*.

Transaction Logging (Circular versus Archive)

Domino supports transaction logging and recovery by capturing database changes and writing them to the transaction log. If a system or media failure occurs, you can use the transaction log and a third-party backup utility to recover your databases.

Transaction logging simplifies your daily backup procedure. You can use a third-party backup utility (like Data Protector Lotus Notes Integration Agent) to perform daily full backups of the transaction logs, instead of full database backups.

Transaction logging only works with databases in Domino Release 5 format but not with databases that use formats from earlier releases. After you enable transaction logging, all databases in Release 5 format are automatically logged.

As mentioned above there are two different logging styles when transaction logging is enabled:

- Circular logging

This is the default mode when transaction logging is enabled. The Domino R5 Server continuously reuses the same log file, which is defined at a designated size, thus overwriting old transactions once the transaction log is filled. You are limited to restoring only the transactions stored in the transaction log. Archiving of transaction logs is not possible if circular transaction logging is used.

- Archive logging

Transaction logging is an integral part of recovering from system and media failures. As mentioned before, this is the only way to perform a backup of transaction logs. This reduces the time needed to perform a restore in case of media or system failure.

The Domino R5 Server does not reuse the log extents until they are backed up. The system uses the transaction logs to apply or undo database transactions not flushed to disk for databases that were open during the system failure.

A media failure may cause a database to be damaged or lost. To recover, use the Data Protector Lotus Notes Integration Agent to restore database backups and archived transaction log files.

When transaction logging is enabled, you may see multiple S0000000.TXN files in the \log directory (you can, optionally, specify different log directory). A transaction log is a binary file where transactions are written. The maximum size of each log extent (.txn file) is 64 MB. Default log space used for log extents is 192 MB. Maximum is 4 GB. Domino formats at least 3 and up to 64 log files, depending on the maximum log space you allocate.

Configuring Data Protector Lotus Notes Integration

It is assumed that the installation of the Data Protector software components on the Lotus Domino R5 Server was successful.

It is recommended that you configure and run a Data Protector filesystem backup of the Domino R5 Server system. A filesystem backup can be performed only if you have installed the Disk Agent on the Domino R5 Server system.

In case of failures, this type of backup is much easier to troubleshoot than the integration of the Domino R5 Server with Data Protector.

Configuring Data Protector Lotus Notes Integration means preparing the environment for starting backup. The environment parameters such as the Domino Server name, path to the `notes.ini` file, path to Lotus home directory, path to Domino data directory and path to Domino executables are saved in the Data Protector Lotus configuration files on the Cell Manager. The configuration has to be done for each Domino R5 Server.

Data Protector Lotus Notes Configuration Files

Data Protector stores Lotus Notes integration parameters in two files on the Cell Manager. These files are created during the configuration of the Lotus Domino R5 Server with Data Protector. These files are:

- Global configuration file

This file is used to define the names of all configured Lotus Domino R5 Servers. It is stored on the following location:

```
/etc/opt/omni/integ/config/Lotus/<client_name>%_OB2_LOTUS
```

- Server specific configuration file

This file is used to define the absolute pathname to the `notes.ini` file, Lotus home directory, Domino data directory, and Domino executables for every configured Lotus Domino client. It is stored on the following location:

```
/etc/opt/omni/integ/config/Lotus/<client_name>%<srv_name>
```

Syntax

The syntax of the global configuration file is as follows:

IMPORTANT

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

```
SRV_LIST=( ' <SRV_NAME1>' [ , ' <SRV_NAME2>' , ' <SRV_NAME3>' ... ] );
```

Example

This is an example of the global configuration file:

```
SRV_LIST=( ' RED' , ' BLUE' );
```

Syntax

The syntax of the server specific configuration file is as follows:

IMPORTANT

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

```
INI_FILE='<full path to notes.ini file>';  
LOTUS_HOME='<full path to Lotus home directory>';  
LOTUS_DATA='<full path to Domino data directory>';  
LOTUS_EXEC='<full path to Domino executables>';
```

Example

This is an example of the server specific configuration file:

```
INI_FILE='/opt/lotus/notesdata/notes.ini';  
LOTUS_HOME='/opt/lotus/';  
LOTUS_DATA='/opt/lotus/lotusdata';  
LOTUS_EXEC='/opt/lotus/notes/latest/hppa';
```

Creating a Link to Lotus C API Library

It is necessary to create a link to the Lotus C API library in order to run Lotus Notes integration agent. To create the link, follow the steps below:

1. Connect to the respective Lotus Domino R5 Server. You must be logged in as a root user.
2. Change to the `/opt/omni/lib` directory:

```
cd /opt/omni/lib on HP-UX systems  
cd /usr/omni/lib on other UNIX systems
```
3. If you have Lotus Domino R5 Server installed on an HP-UX system, create a soft link of the `libnotes.sl` library pointing to the `<DOMINO_EXEC>/libnotes.sl` library by executing the following command:

```
ln -s <DOMINO_EXEC>/libnotes.sl libnotes.sl
```
4. If you have Lotus Domino R5 Server installed on an AIX system, create a soft link of the `libnotes_r.a` library pointing to the `<DOMINO_EXEC>/libnotes_r.a` library by executing the following command:

Integrating Lotus Domino R5 Server and Data Protector

Configuring the Integration

Example of Creating Soft Links

```
ln -s <DOMINO_EXEC>/libnotes_r.a libnotes_r.a
```

If you are running a Lotus Domino R5 Server on an HP-UX system, execute the following command:

```
ln -s /opt/lotus/notes/latest/hppa/libnotes.sl libnotes.sl
```

If you are running a Lotus Domino R5 Server on an AIX system, execute the following command:

```
ln -s /opt/lotus/notes/latest/ibmpow/libnotes_r.a  
libnotes_r.a
```

Configuring a Lotus Domino R5 Server User in Data Protector

In order to start a Lotus Domino backup session, you need an operating system logon for the system on which the Lotus Domino R5 Server is running.

By default, the username is notes and the group is notes.

This user is allowed to backup or restore a Lotus Domino R5 Server database. In order to start a backup of a Lotus Domino R5 Server database using Data Protector, this user has to become the owner of the Data Protector backup specification.

IMPORTANT

Additionally, the operating system root user on the Lotus Domino R5 Server also has to be added to either the Data Protector admin or operator user group.

After the two users are added to either the Data Protector admin or operator user group, Data Protector sessions can be started under the user account with all the privileges required to perform a Lotus Domino R5 Server database backup with Data Protector.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for detailed information on Data Protector user rights and how to add a user to a user group.

Configuring the Lotus Notes Integration Using the Data Protector GUI

The configuration is performed at the same time that a new backup specification is created. Refer to “Creating a Backup Specification” on page 508 for instructions on how to create a backup specification.

Checking the Configuration

You can check the configuration after it has been configured or it can also be checked if you have already created and saved a backup specification for backing up a particular Domino R5 Server:

1. In HP OpenView Storage Data Protector Manager, switch to the Backup context. In the Scoping Pane, expand Backup, Backup Specification, then Lotus Notes.
2. In the Results Area, right-click on the backup specification.
3. In the Source property page, right-click on the name of the client system, then click Check Configuration.

If the configuration is successful, you will receive a message confirming that the integration was properly configured. If not, you will receive a message explaining reasons for the unsuccessful configuration.

Configuring the Lotus Notes Integration Using the Command Line

Execute the following command to configure the Lotus Notes Integration using the Data Protector CLI:

Syntax

On HP-UX systems:

```
/opt/omni/sbin/util_notes.exe -CONFIG -SERVER:<SRV_NAME>  
-INI:<path to the notes.ini file> -HOMEDIR:<path to Lotus  
home directory> -DATADIR:<path to Domino data directory>  
-EXECDIR:<path to Domino executables directory>
```

On other UNIX systems:

```
/usr/omni/bin/util_notes.exe -CONFIG -SERVER:<SRV_NAME>  
-INI:<path to the notes.ini file> -HOMEDIR:<path to Lotus  
home directory> -DATADIR:<path to Domino data directory>  
-EXECDIR:<path to Domino executables directory>
```

The variables are defined as follows:

- <path to the notes.ini file>

Integrating Lotus Domino R5 Server and Data Protector Configuring the Integration

Full path to the Lotus Domino R5 Server `notes.ini` file.

- `<SRV_NAME>`

The Lotus Domino R5 Server name.

- `<path to Lotus home directory>`

Full path to the Lotus Domino R5 Server home directory.

- `<path to Domino data directory>`

Full path to the Lotus Domino data directory.

- `<path to Domino executables directory>`

Full path to the Lotus Domino executables.

Example

In the example below, the Lotus Domino R5 Server name is `BLUE` and `notes.ini` is located in the `/opt/lotus/notesdata/notes.ini` directory.

```
/opt/omni/lbin/util_notes.exe -CONFIG -SERVER:BLUE
-INI:/opt/lotus/notesdata/notes.ini -HOMEDIR:/opt/lotus
-DATADIR:/opt/lotus/notesdata
-EXECDIR:/opt/lotus/notes/latest/hppa
```

Checking the Configuration

To check the configuration, you can run the following command on the Lotus Domino R5 Server system:

On HP-UX systems:

```
/opt/omni/lbin/util_notes.exe -CHKCONF -SERVER:BLUE
```

On other UNIX systems:

```
/usr/omni/bin/util_notes.exe -CHKCONF -SERVER:BLUE
```

Data Protector will check the path to the specified directories and files.

In case of an error, the error number is displayed in the form `*RETVAL*<Error_number>` otherwise the `*RETVAL*0` is displayed.

Testing the Integration

Once you have created and saved a backup specification, you should test it before running a backup. The test verifies both parts of the integration, the Lotus Domino Server side and the Data Protector side. The configuration is tested as well.

The procedure consists of checking the Lotus Domino Server and the Data Protector part of the integration to ensure that communication between Domino Server and Data Protector is established, such that the data transfer works properly.

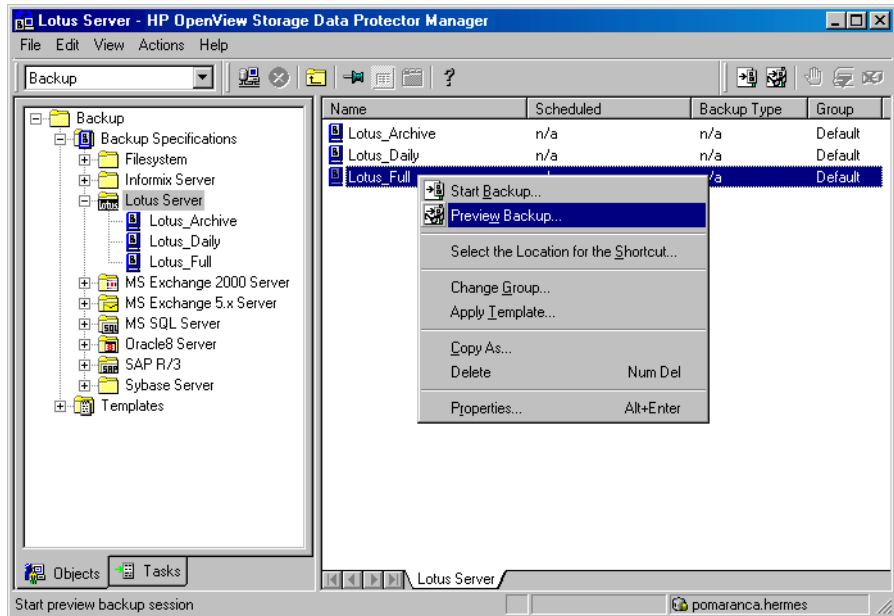
Test your backup specifications thoroughly by previewing them, then running them on file devices and then finally on the actual devices you intend to use. To test your backup specifications, you can use either the Data Protector GUI or the Data Protector CLI.

Testing Using the Data Protector GUI

Follow the procedure below to test the backup of a Lotus Domino R5 Server backup specification:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then double-click Backup Specifications. Expand Lotus Server and right-click the backup specification you want to preview.
3. Click Preview Backup.

Figure 8-4 Previewing a Backup



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

Testing Using the Command Line

A test can be executed from the command line on the Lotus Domino R5 Server system or on any Data Protector client system within the same Data Protector cell, provided that the system has the Data Protector User Interface installed.

Execute the following command:

On HP-UX systems:

```
/opt/omni/bin/omnib -lotus_list <backup_specification_name>  
-test_bar
```

On other UNIX systems:


```
/usr/omni/bin/omnib -lotus_list <backup_specification_name>  
-test_bar
```

What Happens?

The given procedure performs a backup preview that tests:

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

The `testbar.exe` command only tests the Data Protector part of the configuration.

Backing Up Lotus Domino R5 Server

Before You Begin Before performing a backup of Lotus Domino R5 Server, make sure that archive transaction logging is enabled.

What to Backup? Lotus Notes integration agent provides functions for backing up and restoring Notes data in the Domino R5 Server data directory.

The Lotus Domino R5 Server databases consist of the following files:

- NSF (Notes Storage Facility) files.
These files are databases.
- NTF (Notes Template Facility) files.
These files are templates for creating new NSF databases.
- BOX files.
Mail router uses these files.
- Transaction log files, named SXXXXXXX.TXN, where XXXXXXXX is a 7 digit number that is automatically incremented for every new transaction file (the maximum size of this file is 64 MB).

NOTE

When the online backup and database are under transaction log, Lotus Notes Integration Agent (Agent) backs up transactions made during backup. Agent saves the changed information to a newly created file `<db_name>.CI`. This file is also backed up and after the backup is complete, it is deleted. During restore, the Agent applies all the changed information from `.CI` file to restored database. Following this, the recovery is performed.

Because the log files log the changes made to the database, the backup functions use the log files to perform archive backups. Lotus Domino R5 Server automatically recycles archived transaction logs after backup.

IMPORTANT

It is important that archived transaction log files are backed up often enough, so log files do not exceed the amount of disk space meant for log files.

NOTE

You cannot backup Notes database R4 format, since only Lotus Domino Server R5 format is supported. For complete details on upgrading from Domino Release 4 to Domino Release 5, please refer to *Domino 5 Administration Help*.

Domino Templates (NTF)

NTF files (Domino templates) are templates for creating new NSF databases which, in contrast with NSF files, never change. To speed up a full Lotus Domino Server backup, it is recommended to create a separate backup specification for NTF files and back them up. When you create a backup specification for a full Lotus Domino Server backup, add NTF files to the private exclusion list to exclude them from the full backup. For more information on exclusion lists, refer to the *HP OpenView Storage Data Protector Administrator's Guide*. Note that NTF files are not backed up during the incremental backup because they do not change.

The Lotus Notes Integration Agent is an agent that backs up and restores Lotus Domino databases, Domino templates, and transaction logs. Besides Notes databases and archived transaction log files (Notes data) there are a lot of non-database files which needs to be backed up to provide a complete data protection solution for a Domino server. Domino non-database data include:

- `notes.ini`
- `desktop.dsk`
- All `*.id` files

To backup this data and ensure complete data protection use Data Protector file system backup solutions. Refer to "*HP OpenView Storage Data Protector Administrator's Guide*" for detailed info.

Configuring a Lotus Domino R5 Server Backup

To configure an Lotus Domino R5 Server backup, perform the following steps:

1. Configure the devices you plan to use for the backup. See the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instructions.
2. Configure media pools and media for the backup. See the *HP OpenView Storage Data Protector Administrator's Guide* or online Help for instructions.
3. Create a Lotus Domino R5 Server backup specification, specifying what to back up, which devices to use, and how to back it up.

Refer to the following section for the procedure for creating a backup specification.

Creating a Backup Specification

The Data Protector backup specification contains information on how, when and which Lotus Domino R5 Server backup objects to back up.

NOTE

All Lotus Domino R5 Server backup specifications are collected under "Lotus Server" in the Scoping Pane.

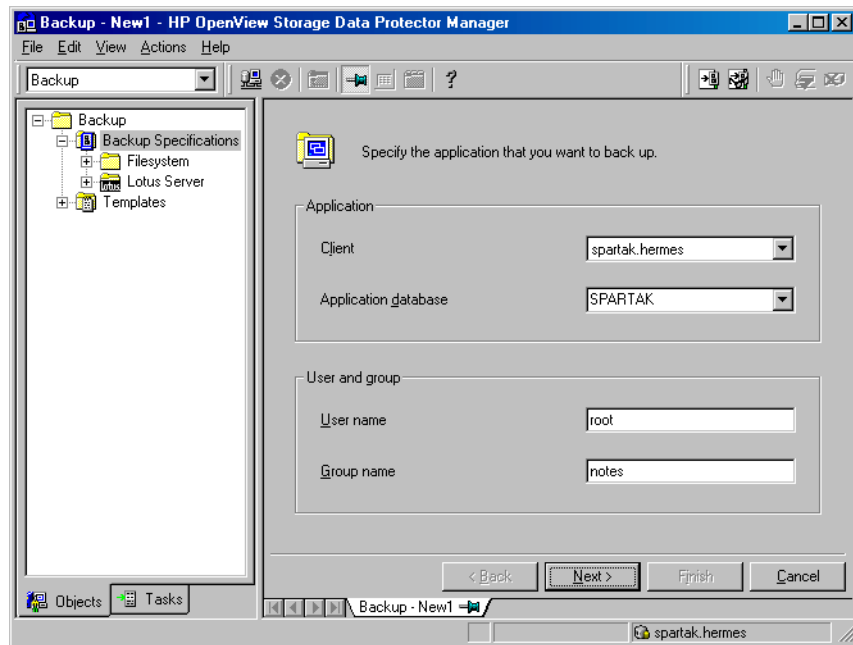
To create a backup specification, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then double-click Backup Specifications.
3. In the Results Area, right-click Lotus Server and then click Add Backup. The Create New Backup dialog box is displayed.
4. Select the Blank Lotus Notes Backup template, and then click OK to start a backup wizard.
5. In the first page of the wizard, specify a client name.

In the Application database drop-down window, select the name of the Lotus Domino R5 Server that you want to back up. If the Lotus Domino R5 Server has not yet been configured, type the name of the Domino R5 Server manually in the drop-down window.

You must also enter the name of the user who is the owner of the backup and the name of the user group. This was configured during the installation of the application. Click Next to proceed.

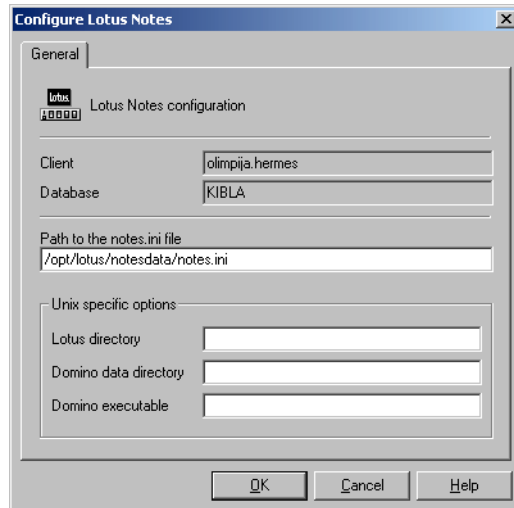
Figure 8-5 Specifying a Client Name and Selecting an Application Database



6. If you are configuring the Lotus Notes integration for the specified Domino Server for the first time, a window `Configure Lotus Notes` is displayed.

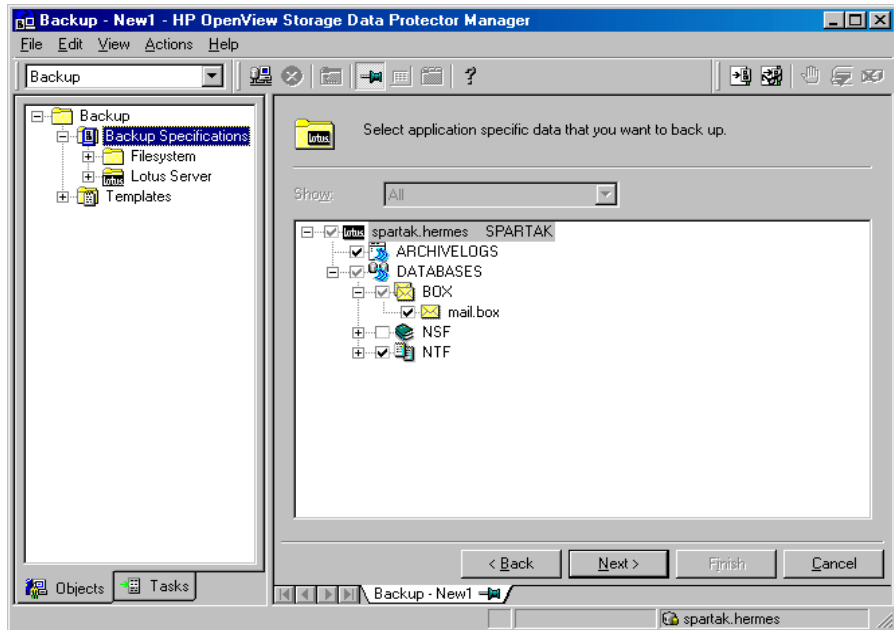
Specify the full pathname to the `notes.ini` file located on the Lotus Domino client system. You must also specify the full pathname to the Lotus home directory, Domino data directory and Lotus Domino executables directory. Click Next to continue.

Figure 8-6 Specifying the Pathname to the Notes.ini File



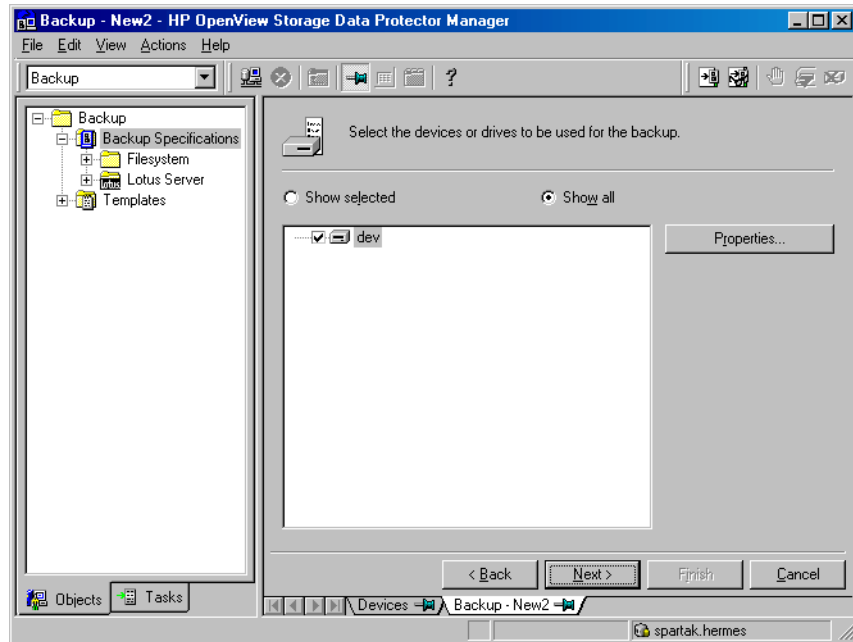
7. Select the Lotus Domino R5 Server objects you want to back up. Click Next.

Figure 8-7 **Selecting Objects for a Backup**



8. Select the device you want to use for the backup. Click *Properties* to set the device concurrency, media pool, and preallocation policy. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information about these options.

Figure 8-8 **Selecting Backup Devices**



9. Click Next to proceed to select backup options and to schedule your backup.

Refer to online Help or the *HP OpenView Storage Data Protector Administrator's Guide* for details about options common to all Data Protector backup specifications.

Refer to “Lotus Domino R5 Server Specific Backup Options” on page 512 for details about Lotus Domino R5 Server specific options.

10. Once you have defined all backup options and optionally scheduled the backup, name and save the newly created backup specification.

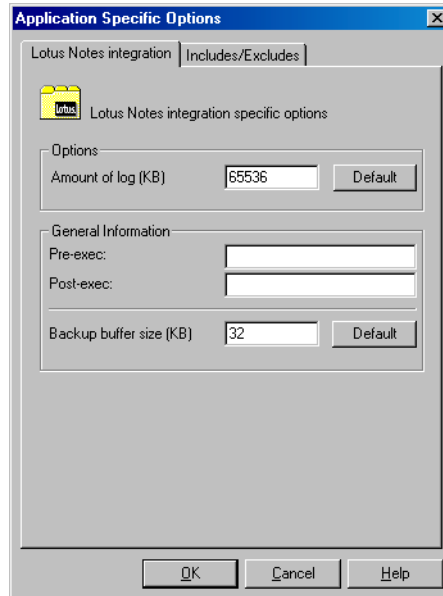
Once saved, the backup specification can be tested by clicking Start Preview, or started by clicking Start Backup.

Lotus Domino R5 Server Specific Backup Options

This section describes backup options specific to the Data Protector Lotus Notes integration.

You can access these options in the Options property page of a backup specification. Click the Advanced button next to the Application Specific Options. Refer to “Application Specific Options” on page 513.

Figure 8-9 Application Specific Options



The following options can be selected from this window:

Amount of log This is the size of the amount of log needed for the database recovery. In case that the database has less amount of log than specified, the incremental backup skips the database. If the database exceeds the specified amount of log, the full backup of the database is performed.

Pre-exec Specifies a command with arguments or a script that will be started on the Lotus Domino R5 Server client before the backup starts. The command/script is started by Data Protector `ldbar.exe` and has to reside in the `/opt/omni/bin` on HP-UX and `/usr/omni/bin` directory on other UNIX systems. Only the filename must be provided in the backup specification.

Post-exec Specifies a command with arguments or a script that will be started on the Lotus Domino R5 Server client after the backup. The command/script is started by Data Protector `ldbar.exe` and has to reside in the `/opt/omni/bin` on HP-UX and `/usr/omni/bin` directory on other UNIX systems. Only the filename must be provided in the backup specification.

Backup buffer size This is the size of the Lotus Notes integration buffer which is used to transfer data to Data Protector.

The application specific options are applied to all backup objects that have been selected in the backup specification.

Running an Online Backup

To run an online backup of a Lotus Domino R5 Server object, use any of the following methods:

- Schedule the backup of an existing Lotus Domino R5 Server backup specification using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI.

Scheduling a Backup

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

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

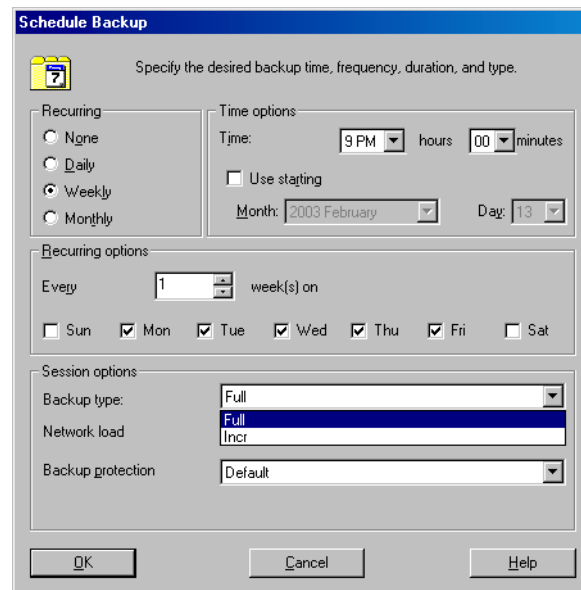
A backup schedule can be tailored according to your business needs. If you have to keep the databases online continuously, then you should back it up frequently, including the backup of the archived transaction logs, which are required in case you need a recovery to a particular point in time.

To schedule a Lotus Domino R5 Server backup specification, perform the following steps:

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

- 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. See Figure 8-10 on page 515.
 6. Click OK to return to the Schedule property page.
 7. Click Apply to save the changes.

Figure 8-10 Scheduling Backups



Starting an Interactive Backup

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

To start an interactive backup of a Lotus Domino R5 Server backup object, perform the following steps:

Backing Up Lotus Domino R5 Server

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.

2. In the Scoping Pane, expand the Backup, and then the Backup Specifications items.

Expand Lotus Server. A list of backup specifications appears.

3. Right-click the backup specification you want to back up, and then select Start Backup from the pop-up menu.

The Start Backup dialog box appears.

Select the backup type and network load.

Refer to online Help for a description of network load.

4. Click OK.

Messages appear in the Results Area as the backup session proceeds. Upon successful completion of the backup session, the Session completed successfully message is displayed.

Restoring Lotus Domino R5 Server Data

You can restore Lotus Domino R5 Server objects using the Data Protector GUI.

Databases are restored directly to the host with the installed Domino R5 Server using `ldbar.exe`. Through the Lotus Notes Integration Agent you are able to bring database offline, put the databases online, and put perform database recoveries after restores. In the case that you perform a recovery, transaction logs are also restored if needed. This step (restoring archive logs) is performed automatically during recovery process.

Database restore can be done while the server is online. You can specify the restore location. So you can restore the database to the same location as it was backed up from (so in case the database is corrupted or deleted you can replace it) or you can restore the database to other location than original and you keep the original database intact.

After restore of Domino database, database is not active. If you access it, it will be automatically brought online. But in this case the recovery is not performed. In most cases you would like to get the last possible consistent state of the databases or to do recovers to a specific point in time. In this situation you must use the recover option.

IMPORTANT

In case when restore location resides the database with the same file name as restored one, then this database is taken offline and deleted.

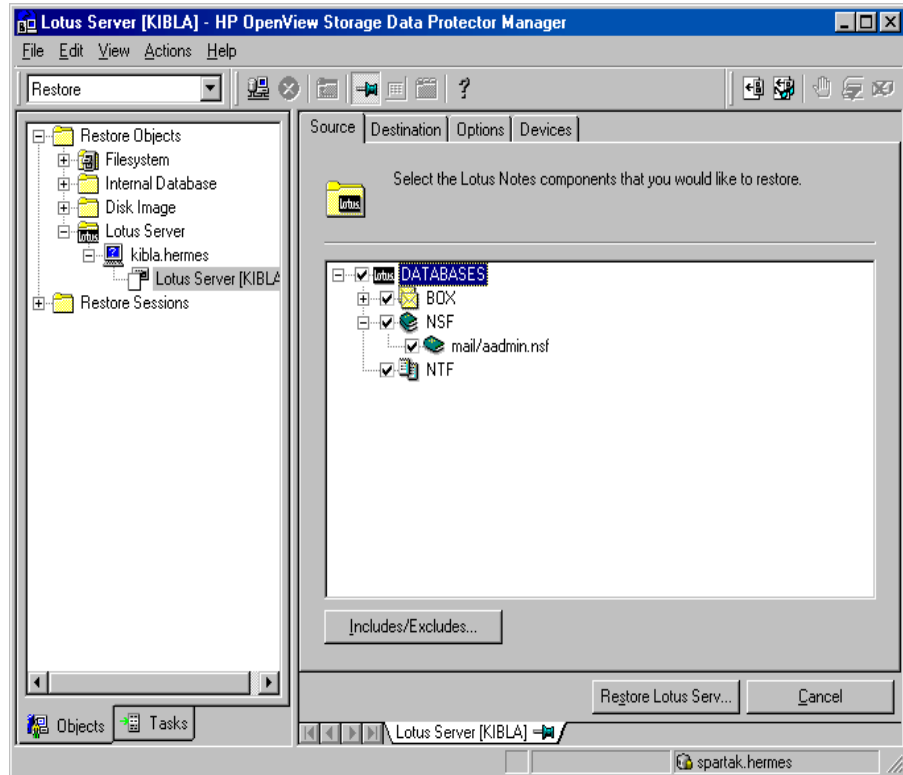
Restore Procedure

Use the following procedure to restore the Lotus Domino R5 Server objects:

1. In HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand Lotus Server and then the name of the client system from which you want to restore.

3. Browse for and select the backed up Lotus Domino R5 Server objects you want to restore.

Figure 8-11 Restore Objects



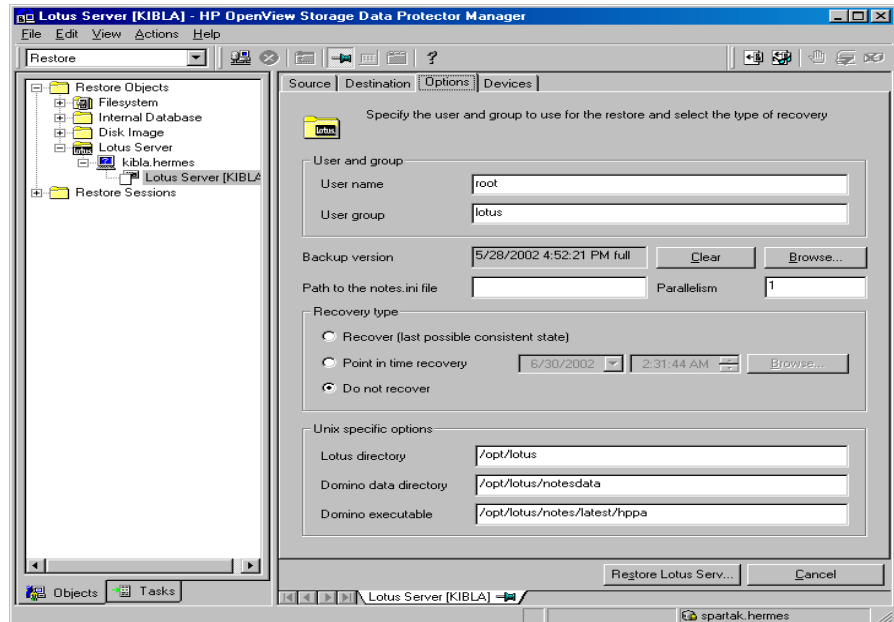
Backup version can be selected in the Options property page. Click **Browse** to select a different version of backup.

NOTE

In the source property page all the backed up databases are listed. If you are restoring multiple databases from a specific backup session, you must be sure that database selected for restore were backed up in the selected backup session. If this criteria is not met, a warning 'object not found in the database' appears at the restore time. Restoring from different backup sessions demands separate restore sessions. The only

exception is when the backup session is not specified. In such cases, Lotus Notes agent finds the latest backup version of each database for restore.

Figure 8-12 Selecting Restore Options



4. Select the restore options from the Options property page. Refer to “Restore Options” on page 521.

The devices and media for restore are automatically selected.

Note that you can change the device used for the restore. Therefore, you have the possibility of using a different device for a restore than the one that was used for the backup. 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.

5. Click Restore Lotus Server. Review your selection, and then click Finish to start a restore session.

The restore session messages are displayed in the Results Area.

Restore Options

The following destination and restore options are specific to the Data Protector Lotus Notes integration.

Destination Options

Destination options are the following:

- Target client

By default, the target Data Protector Lotus Domino R5 Server client is the Lotus Domino R5 Server from which the application data was backed up. If the target client is a Windows system, then UNIX specific options are disabled. If the target client is a UNIX system, then UNIX specific options are enabled and you must manually enter the Lotus home directory, Lotus Domino data directory, and the Lotus Domino executables directory. The new target Lotus Domino R5 Server must be a part of the Data Protector cell and have the Lotus Notes Integration software component installed.

Nevertheless, the databases can be restored to a Lotus Domino R5 Server other than the one the backup was made from. The new target Lotus Domino R5 Server must be a part of the Data Protector cell and have the Lotus Notes Integration software component installed.

- Restore Location

- ✓ Restore to original location

This is the default option. You can restore the databases to the same directory from which it was backed up (it can be on the original client system or on some other client system which you selected).

- ✓ Restore to new location

This option enables you to restore your data to another directory. When defining the restore location you can specify the relative directory to Domino data directory where you want to restore your data.

Example

Domino data directory resides on `/opt/lotus/notesdata/BLUE`.

In case you want to restore database to `/opt/lotus/notesdata/BLUE/restore_dir/` directory, you just have to specify `restore_dir` directory. The restored database filenames are the same as they were at the backup time.

Restore Options

You can specify the following restore options:

- User and group

Enter the Lotus user name and group, for example, "notes", "notes".

- Backup version

By default restore is done from the last full backup of a database. Click **Browse** button to define backup version other than the last one.

- Path to the `notes.ini` file

Specify the full path to the Lotus Domino R5 Server `notes.ini` file.

- Parallelism

Specify how many Lotus Notes agents will start the restore. By default, this value is set to 1.

- Recovery type

- ✓ Perform recovery (last possible consistent state)

This is the default option. Select this option to restore the database to the last possible consistent state. This also includes restore of archived transaction logs if needed during recovery.

- ✓ Point-in-time recovery

You can specify a point in time to which the database state should be restored. Click **Browse** to specify the desired date and time for the point-in-time recovery. In this case, only transactions written before the specified date and time are applied to the database.

- ✓ Do not perform recovery

If this option is selected, the restore of the specified database is performed from the last backup or from the specified backup version. This option only restores databases without trying to recover the database. Any transactions made during or after backup are not reflected in the restored database.

Restoring Lotus Domino R5 Server Data

- UNIX specific options

These options are only enabled if the target system is a UNIX system. If you want to perform the restore, the following UNIX restore options must be specified:

- ✓ Lotus directory

Specify the full path to the Lotus Domino R5 Server home directory.

- ✓ Domino data directory

Specify the full path to the Lotus Domino data directory.

- ✓ Domino executables

Specify the full path to the Lotus Domino executables.

Monitoring a Lotus Domino Server Backup and Restore

The Data Protector GUI allows you to monitor the current or previous backup and restore sessions. Note that you have to have the appropriate privileges to view previous sessions.

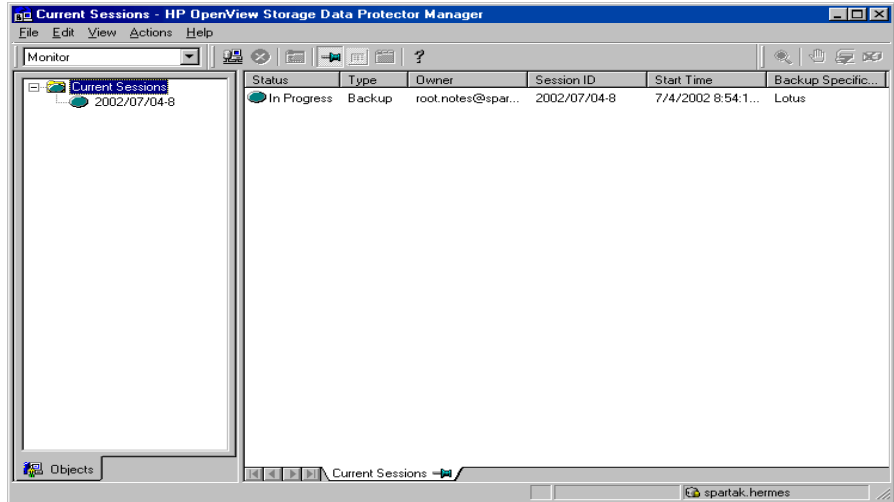
Monitoring is automatically activated when you start a restore or backup.

Monitoring Current Sessions

The procedure below describes how to monitor a current session using in the Data Protector GUI:

1. In HP OpenView Storage Data Protector Manager, switch to Monitor context.
2. The sessions that are currently in progress are displayed in the Results Pane. If there are no sessions currently in progress, the Results Pane is empty.
3. Double-click on the sessions you want to monitor.

Figure 8-13 **Monitoring Current Sessions**




Viewing Previous Sessions

The procedure below describes how to monitor previous sessions:

1. In the HP OpenView Storage Data Protector Manager, switch to the Internal Database context. A list of items is displayed in the Scoping Pane.
2. Expand Sessions. A list of sessions appears in the Scoping Pane.
3. Select the session you want to view.

See the *HP OpenView Storage Data Protector Administrator's Guide* for details.

Figure 8-14 Viewing Previous Sessions



Troubleshooting

This section is divided into the following subsections:

- General troubleshooting
- Checking prerequisites related to the Lotus Domino R5 Server side of the integration
- Configuration problems
- Backup problems
- Restore problems
- Recovery problems

General Troubleshooting

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, problems and workarounds, as well as the 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.

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

Checking Prerequisites Related to the Lotus Domino R5 Server Side of the Integration

For more detailed information about how to perform any of the following procedures, refer to the Lotus Domino R5 Server documentation.

1. Check the environment variables:

Prior to any Lotus C API call from Data Protector Lotus integration agent, the Lotus C API has to be initialized. To successfully initialize it, the following environment variables must be set:

```
LOTUS=/opt/lotus
```

```
NOTES_DATA_DIR=/local/notesdata
```

```
Notes_ExecDirectory=<Lotus_home>/notes/latest/ibmpow
```

```
PATH=$PATH:$LOTUS:$NOTES_DATA_DIR:$Notes_ExecDirectory:/opt/lotus/bin
```

```
PATH=$PATH:$Notes_ExecDirectory/res/C
```

These variables are usually exported by Lotus Notes integration agent or utility prior to Lotus Notes C API initialization. If you experience problems with Lotus C API initialization, please try to export these variables manually or put them in `.omnirc` file. For more information on how to use `.omnirc` file, refer to *HP OpenView Storage Data Protector Administrator's Guide*.

2. Check if the soft link to `libnotes` library exists:

If the soft link to Lotus Domino `libnotes` library does not exist, you can experience the following error while running Lotus Notes utility from the command line:

```
#./util_notes.exe
```

```
/usr/lib/dld.sl: Can't find path for shared library:  
libnotes.sl
```

```
/usr/lib/dld.sl: No such file or directory
```

```
Abort (coredump)
```

Check if the soft link from
<Data_Protector_home>/lib/libnotes.sl directory to Lotus
Domino library libnotes.sl exists. The name of the link must be the
same as libnotes.sl library name in the Lotus Domino Exec
directory.

NOTE

The soft link must be checked on all Lotus Domino clients that you have
in your cell.

Example

On HP-UX systems, the following link must exist in the /opt/omni/lib
directory:

```
libnotes.sl -> /opt/lotus/notes/latest/hppa/libnotes.sl
```

On AIX systems, the following link must exist in the /usr/omni/lib
directory:

```
libnotes_r.a -> /opt/lotus/notes/latesst/ibmpow/libnotes_r.a
```

After setting the soft link, check if the soft link works:

- a. If you are using an HP-UX system, run:

```
/opt/omni/lbin/util_notes.exe -app
```

If the following output is displayed, the libnotes library is not linked
properly. Please check the link again.

```
/usr/lib/dld.sl: Can't find path for shared library:  
libnotes.sl
```

```
/usr/lib/dld.sl: No such file or directory
```

If the soft link is set correctly, the *RETVAL*0 message is displayed.

- b. If you are using an AIX system, run:

```
/usr/omni/bin/util_notes.exe -app
```

If the following output is displayed, the libnotes library is not linked
properly. Please check the link again.

```
exec():0509-036 Cannot load program ./util_notes.exe  
because of the following errors:
```

```
0509-022 Cannot load library libnotes_r.a
```

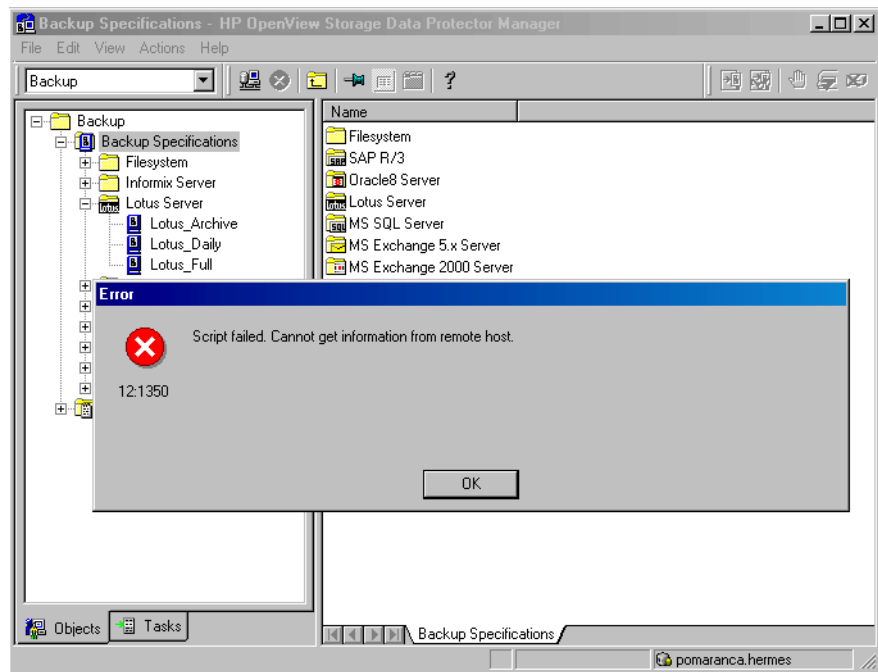

0509-026 System error: A file or directory in the path does not exist.

If the soft link is set correctly, the *RETVAL*0 message is displayed.

3. Script failed error:

You can get the following error message while configuring or starting a backup using the Data Protector GUI

Figure 8-15 Script Failed Error



To solve this problem, refer to the procedure described in “Checking Prerequisites Related to the Lotus Domino R5 Server Side of the Integration” on page 527.

Configuration Problems

IMPORTANT

If you have encountered any errors up to this point when performing procedures described in the previous section, refer to Lotus Domino R5 Server support. The respective tests have to be done before you even start checking the Data Protector Lotus Domino R5 Server configuration.

1. **Verify that the Data Protector software has been installed properly.**

See “Installing the Lotus Notes Integration” on page 492 for details.

2. **Check the .omnirc environment settings:**

Examine the environment settings in the .omnirc file, which is located in the /opt/omni directory.

Backup Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. **Check your Lotus Domino R5 Server configuration:**

To check the Lotus Domino R5 Server configuration, login as a Lotus Domino group dba user to the Lotus Domino R5 Server system and start the following command on the Lotus Domino R5 Server system:

```
/opt/omni/lbin/util_notes.exe -CHKCONF -SERVER:<SRV_NAME>
```

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 <set_number> <Error_number>
```

*RETVAL*0 indicates a successful configuration.

2. **Perform a filesystem backup of the Lotus Domino R/5 Server system:**

Perform a filesystem backup of the Lotus Domino R/5 Server system so that you can eliminate the chance of any potential communication problems between the Lotus Domino R/5 and the Data Protector Cell Manager system.

Do not start troubleshooting an online database backup unless you have successfully completed a filesystem backup of the Lotus Domino R/5 Server system.

Refer to online Help or the *HP OpenView Storage Data Protector Administrator's Guide* for details about how to do a filesystem backup.

If the Lotus Domino R5 Server part of the filesystem backup fails, examine the system errors reported in `/var/opt/omni/log/debug.log`, which is located on the Data Protector Lotus client system. Try to restart the Lotus Domino R5 Server and observe the server messages.

If the Data Protector part of the filesystem backup fails, examine the system errors reported in `/var/opt/omni/log/debug.log`, which is located on the Data Protector Cell Manager system.

If the filesystem backup succeeds, the problem is probably insufficient memory, disk space or other OS resource of the client running the Data Protector User Interface.

3. **Verify Data Protector internal data transfer using the `testbar` utility.**

Before you run the `testbar` utility, verify that the Cell Manager name is correctly defined on the Lotus Domino R5 Server system. Check the `/etc/opt/omni/cell/cell_server` (HP-UX systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system. Then run the following command:

On HP-UX systems:

```
/opt/omni/bin/utilns/testbar -type:Lotus -apname:<SRV_NAME>  
-bar:<backup_specification_name> -perform:backup
```

On other UNIX systems:

```
/usr/omni/bin/utilns/testbar -type:Lotus -apname:<SRV_NAME>  
-bar:<backup_specification_name> -perform:backup
```

Switch to the Data Protector Manager and examine the errors reported by the testbar utility by clicking the Details button in the Data Protector Monitor context.

Create an Lotus Domino R5 Server backup specification to back up to the null device or file. If the backup succeeds, the problem may be related to the backup devices. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for instructions on troubleshooting devices.

4. **Start a backup session using ldbar.exe:**

You can start a test backup session using the Data Protector command line interface, where the backup options have to be specified as the ldbar.exe command line options.

The command has to be started from the /opt/omni/bin directory on a Data Protector Lotus client system. Run the following command:

```
ldbar -perform:backup -db:<DB_NAME> -server:<SRV_NAME>  
-ini:<path to notes.ini file>  
-bar:<backup_specification_name> -homedir:<path to Lotus  
home> -datadir:<path to Domino data> -execdir:<path to  
Domino executables>
```

For other ldbar.exe parameters, please refer to command help by executing:

```
ldbar.exe -help
```

NOTE

The -bar option is mandatory since the ldbar.exe reads the device options from the backup specification as opposed to other options in the respective backup specification, which are ignored. The command line options are used instead.

5. **If the Lotus Domino R5 Server freezes during backup:**

During the backup session, it can happen that Lotus Domino R5 Server freezes with the following error:

```
Fatal Error signal = 0x0000000b PID/TID = xxxx/1  
Freezing all server threads ...
```

This can happen in the following cases:

- a. The Lotus Notes C API initialization failed and caused the server to freeze. Perform the following steps:

i. Logon to the Lotus Domino client system as a root user.

ii. Kill all the `ldbar.exe` processes:

```
ps -ef | grep ldbar.exe | grep -v 'grep' | awk  
{ 'print $2' } | xargs kill -9
```

iii. If the Lotus Domino R5 Server is running, restart it. Before restarting make sure that none of the Domino processes are still running.

iv. Logon as a `notes` user and run the following check to see if the server recovered:

```
/opt/omni/bin/util_notes.exe -box -ini:<path to  
notes.ini file>
```

If everything is working properly, the `*RETVAL*0` message is displayed.

NOTE

This is caused by corrupted shared memory and semaphores that the program does not clean up. Even if you do not have any troubles after the crash, it is good practice to clean up before restarting any process.

- b. If the Lotus Domino R5 Server is not online and the Lotus Domino daemon `logasio` is not running, then while the Lotus Notes integration agent is initializing Lotus C API, the `logasio` daemon automatically starts. Since the environment for `notes` user is not set because the `.profile` is not executed, the `logasio` server could fail to start. Perform the following steps:

i. Logon to the Lotus Domino client system as a root user.

ii. Kill all the `logasio` processes:

```
ps -ef | grep logasio | grep -v 'grep' | awk { 'print  
$2' } | xargs kill -9
```

iii. If the Lotus Domino R5 Server is running, restart it. Before restarting make sure that none of the Domino processes are still running.

- iv. Logon as a notes user and run the following check to see if the server recovered:

```
/opt/omni/bin/util_notes.exe -box -ini:<path to notes.ini file>
```

If everything is working properly, the *RETVL*0 message is displayed.

6. Check errors during the backup session

Observe the messages reported during the backup session. In cases when the error is related to the Lotus Domino R5 Server, the following type of error may be displayed:

```
Lotus ERROR [error #]: <Error description>
```

Examine the error description and take appropriate actions.

Example of Lotus Error Message

```
[Critical] From: OB2BAR@ice.hermes "BLUE" Time: 1.10.01  
16:26:48
```

```
Lotus ERROR [3748]: Attempt to backup a database that is  
currently being backed up.
```

Restore Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Perform a filesystem restore:

Check whether the filesystem restore of the problematic client works. It is much easier to troubleshoot a filesystem restore.

2. Check whether the Data Protector Lotus Notes agent `ldbar.exe` is installed on the system.

3. Check the restore to another client:

To restore to another system, ensure that the Lotus Domino R5 Server is installed and that it has the same non-database files as the Lotus Domino R5 Server whose backup is to be restored. Those files must be first restored from a file system backup.

4. Examine system errors:

If the Lotus Domino R5 Server restore fails, examine the system errors reported in `/var/opt/omni/log/debug.log`, which is located on the Data Protector Lotus client system.

5. Test a restore session using `ldbar.exe`:

A restore session can be tested using the Data Protector `ldbar.exe` command. The command has to be started from the `/opt/omni/bin` directory on the Data Protector Lotus Domino R5 Server system:

```
ldbar.exe -perform:restore -db:<DB_NAME>  
-server:<SRV_NAME> -ini:<path to notes.ini file>
```

For other `ldbar.exe` parameters, please refer to command help by executing:

```
ldbar.exe -help
```

6. Check errors during the restore session:

Observe the messages reported during the restore session. In case the error is related to the Lotus Domino R5 Server, the following type of error can be displayed:

```
Lotus ERROR [error #]: <Error description>
```

Examine the error description and take appropriate actions.

Example of Lotus Error Message

```
[Minor] From: OB2BAR@ice.hermes "BLUE" Time: 30.9.01  
21:56:24
```

```
Lotus ERROR [5098]: The database is in use and cannot be  
taken offline.
```

Recovery Problems

At this stage, you should have performed all the verification steps described in the previous sections. After this, proceed as follows:

1. Recovery of restored Lotus Domino NSF database failed:

You can experience the following error at the recovery process:

```
[Critical] From: OB2BAR@ice.hermes "BLUE" Time: 19.10.01  
17:24:23
```

```
Lotus ERROR [5114]:Recovery Manager: Recovery only  
supported for Backup Files.
```

This error indicates that at least one database from the restore list was accessed, for example: by Lotus Domino Server, any user or any process, before the recovery ended.

Proceed as follows:

- a. Restart the Lotus Domino Server and perform the restore again.
- b. Restore the failed database to a location other than the one it was backed up from.

2. Check the recovery time parameter setting:

The recovery time parameter must be set as follows:

yyyy/mm/dd.hh:mm:ss

The recovery time parameter must be in the above mentioned format, otherwise the recovery time can be misunderstood. It is very important that time format is accurate and that you use a 24 hour format.

Example

2001/01/25.18:15:00

3. Recovery failed with a Lotus ERROR [520] error:

The following error message indicates that the recovery has failed.

```
[Critical] From: OB2BAR@ice.hermes "BLUE" Time: 1.10.01  
9:04:23
```

Lotus ERROR [520]:

This can happen, if you have restored several databases but some of them were not under transaction logging at the backup time. Therefore no database was in the list for recovery. This is the case when you are recovering NTF database types or an NSF database that is not recoverable.

To resolve the problem, try to restore only one database, for which you are sure that it is recoverable, and observe messages. It might be that one database in the restore list is corrupt (was corrupted at backup time) and therefore the Lotus C API recovery call fails.

NOTE

No description is listed for error number 520. This is because the internal error text of Lotus C API. There are several error codes that are internal type and have no description listed.

Before You Call Support

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

1. Provide details about your hardware and software configuration, including official patches you use, the Lotus Domino R5 Server version, etc...
2. Provide a detailed description about the action performed that failed. If you have backup problems, attach the backup specification.
3. Provide the information from the `<Data_Protector_home>\debug.log`. Describe what happened after the failure.
4. Copy the session output to a file.

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 original units.

See also **backup system** and **original unit**.

archived redo log (*Oracle specific term*)

Also called offline redo log. If the Oracle8/9 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. 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 object

Any data selected for backup, such as a disk, a file, a directory, a database, or a part of it. During the backup session, Data Protector reads the objects, transfers the data (through the network), and writes them to the media residing in the devices.

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

Glossary

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

See **media set**.

backup set (*Oracle specific term*)

Backup for (one or more) Oracle8/9 files, where the files are multiplexed together. The reason for multiplexing is to give performance benefits. Files in backup sets have to be extracted using a restore command. There are two types of backup sets: data file backup set and archive log backup set.

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 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 replica units 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 storage version. *See also* **application system** and **replica unit**.

backup types

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

backup view

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

BC (*EMC Symmetrix specific term*)

Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

See also **BCV**.

Glossary

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 **CA** (*HP StorageWorks Disk Array XP specific term*), **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

BC (*HP StorageWorks Virtual Array specific term*)

Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication. 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 sets should be connected to the backup system. See also **HP StorageWorks Virtual Array LUN**.

BC Process (*EMC Symmetrix specific term*)

A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. See also **BCV**.

BC VA (*HP StorageWorks Virtual Array specific term*)

Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

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

Glossary

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
- 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

Glossary

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 backup sessions, restore sessions, and backed up data. Depending on the selected log level, it also contains file names and file versions. This part of the database is always local to the cell.

See also **MMDB**.

CDF file (*UNIX specific term*)

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

cell

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

Cell Manager

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

centralized licensing

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

Glossary

allocate licenses to specific cells to suit your needs.

See also MoM.

Centralized Media Management Database (CMMDB)

See CMMDB.

channel (*Oracle specific term*)

An Oracle8/9 Recovery Manager resource allocation. Every allocated channel starts a new Oracle8/9 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 Oracle8/9 is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging (*MS Exchange and Lotus Domino Server specific term*)
Microsoft Exchange database and Lotus Domino Server database mode in which transaction log files are automatically overwritten as soon as the data they contain is committed to the database.

client backup

A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or client system

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

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

Glossary

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

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.

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.

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

Glossary

users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account

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

data stream

Sequence of data transferred over the communication channel.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle8/9 Server.

database parallelism

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

database server

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

Dbobject (*Informix 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 occupying approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the `<Data_Protector_home>\db40` directory. 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 2 GB.

DCBF

The Detail Catalog Binary Files (DCBF) are a part of the IDB. The files in store information about file versions and attributes occupying approximately 80% of the IDB. By default, DCBF consist of one DC directory with a maximum size of 2 GB. You can create more DC directories.

Glossary

delta backup

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

See also **backup types**

device

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

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (*EMC Symmetrix specific term*)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to

the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information. Data Protector can back up DHCP server data as part of the Windows configuration.

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.

Glossary

direct backup A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems. *See also* **XCOPY engine**.

directory junction (*Windows specific term*)

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

Directory Store (DS) (*MS 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.

Glossary

disk group (*Veritas Volume Manager specific term*)

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

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

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. The DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating

Glossary

system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. Active DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

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

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)

See **Symmetrix Agent (SYMA)**

EMC Symmetrix Application Programming Interface (SYMAPI)

(EMC Symmetrix specific term)

See **Symmetrix Application Programming Interface (SYMAPI)**

EMC Symmetrix CLI Database File

(EMC Symmetrix specific term)

See **Symmetrix CLI Database File**

EMC Symmetrix Command-Line Interface (SYMCLI) *(EMC Symmetrix specific term)*

See **Symmetrix Command-Line Interface (SYMCLI)**

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

Glossary

Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

EVA Agent (*HP StorageWorks Enterprise Virtual Array specific term*)
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array snapshot integration on the application system and the backup system. It communicates with the HSV Element Manager to control the HP StorageWorks Enterprise Virtual Array.

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 II Exchanger.
See also **library**.

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

Extensible Storage Engine (ESE) (*MS Exchange specific term*)
A database technology used as a storage system for information exchange by Microsoft Exchange 2000 Server.

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

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

Glossary

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 Replication Service (FRS)

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

file version

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

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**.

Glossary

full database backup

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

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/options` directory on HP-UX and Solaris systems and in the `<Data_Protector_home>\config\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 (X11/Motif and Windows) graphical user interface, provided by Data Protector for easy access to all configuration and administration tasks.

hard recovery (*MS Exchange specific term*)

Recovery of data on the level of the database engine (Extensible Storage Engine 98).

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/Holidays` on the UNIX Cell Manager and `<Data_Protector_home>\Config\Holidays` on the Windows Cell Manager.

host backup

See client backup with disk discovery.

Glossary

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 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 (*HP StorageWorks Disk Array XP specific term*)

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that are mirrored using Continuous Access XP (CA) and Business Copy XP (BC) configurations. See also **BC** (*HP StorageWorks Disk*

Array XP specific term) and **CA** (*HP StorageWorks Disk Array XP specific term*).

HP StorageWorks Virtual Array

LUN (*HP StorageWorks Virtual Array specific term*)

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that are replicated using the HP StorageWorks Business Copy VA configuration. See also **BC** (*HP StorageWorks Virtual Array specific term*).

HP VPO

See **OVO**.

HSV Element Manager (HP

StorageWorks Enterprise Virtual Array specific term)

The HSV Element Manager is used by the Data Protector HP StorageWorks Enterprise Virtual Array integration to provide the features that enable virtualization technology and the management interface for the HP StorageWorks Enterprise Virtual Array environment.

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,

Glossary

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 (*MS Exchange specific term*)

A backup of changes since the last full or incremental backup. Only transaction logs are backed up.

See also **backup types**.

incremental mailbox backup

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

incremental1 mailbox backup

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

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

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

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

incremental restore (*EMC Symmetrix specific term*)

A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV

Glossary

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 (MDB) *(MS Exchange specific term)*

This is the default message store provider for the Microsoft Exchange Server. The information store consists of the following stores:

- Public information store (MS Exchange 5.5 Server) or Public folder store (MS Exchange 2000 Server)
- Private information store (MS Exchange 5.5 Server) or Mailbox store (MS Exchange 2000 Server)
- 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 a Microsoft Exchange Server organization, even if multiple 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 Microsoft Exchange Server.

See also **Directory Store (DS)**.

Initialization Parameter File *(Oracle specific term)*

An Oracle8/9 file that contains information on how to initialize a database and instance.

initializing *See* **formatting**.

Installation Server

A computer system that holds a repository of the Data Protector

Glossary

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 where data replicated during the ZDB disk backup or ZDB disk/tape backup is restored at high speed using split mirror or snapshot technology. The restore takes place within the disk array and there is no restore from the standard backup media involved. Full recovery of a database application may require further steps, such as applying the log files, to be performed afterwards. Instant recovery restores the user-selected replica storage version to the original storage.

See also zero downtime backup (ZDB), ZDB disk backup, ZDB tape backup, ZDB disk/tape backup and replica storage pool.

integrated security (*MS SQL specific term*)

Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be

used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

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 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.

Glossary

LBO (*EMC Symmetrix specific term*)

A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library

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

lights-out operation or **unattended operation**

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

LISTENER.ORA (*Oracle specific term*)

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

load balancing

By default, Data Protector automatically balances the usage of devices selected

for backup, so that they are used evenly. Load balancing optimizes the usage by balancing the number and the size of the objects backed up 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 load balancing is not selected, you select which device will be used for each object in your 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.

Glossary

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.

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`.

logging level

The logging level determines the amount of details on files and directories written to the IDB during backup. 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 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

Glossary

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 (Oracle8/9) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle8/9) 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 (*MS Exchange specific term*)

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

Mailbox Store (*MS Exchange 2000 Server specific term*)

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

Main Control Unit (MCU) (*HP StorageWorks Disk Array XP specific term*)

An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP*

Glossary

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, the Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, the Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. The Media Agent also manages the robotics control of a library.

MAPI (*MS 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.

Glossary

media pool

A set of media of the same type (such as DDS) used and tracked as a group.

Media are formatted and assigned to a media pool.

media set

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

media type

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

media usage policy

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

merging

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

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 (DMPAPI), 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.

Glossary

Microsoft Volume Shadow Copy service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.

See also shadow copy, shadow copy provider, writer.

mirror (*ZDB specific term*)

See replica unit.

mirror rotation (*HP StorageWorks Disk Array XP specific term*)

See replica storage 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 mountpoints 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.

Glossary

MU number (*HP StorageWorks Disk Array XP specific term*)

A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.

See also **first level mirror**.

multi-drive server

A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat

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

An object can be one of the following:

- for Windows clients, an object is a logical disk (such as d:);

- for UNIX clients, an object is a mounted filesystem or a mount point;
- for Novell Netware clients, an object is a volume.

The scope of the data can be further reduced by selecting files or directories. Additionally, an object can be a database entity.

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.

offline recovery

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI-II 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.

Glossary

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).

online backup

A backup that is performed while the application (or database) is available for use. Application-specific interfaces allow backup products, like Data Protector, to back up logical units of the database while retaining access for the application. In simple configurations (non ZDB), the application remains in a backup mode for the entire duration of the backup. In contrast to that, for ZDB configurations, the backup mode lasts only for the duration of the split/snapshot operation. After that, the application can resume to the standard mode. Depending on the configuration, resource requirements vary significantly.

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**.

Glossary

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.

original unit (*ZDB specific term*)

A logical unit that is used as a source for data replication using snapshot or split mirror technologies. Depending on the vendor and technology used, an original unit denotes P-VOL on HP StorageWorks Disk Array XP, parent LUN on HP StorageWorks Virtual Array, logical drive on HP StorageWorks Modular SAN Array 1000, or virtual disk on HP StorageWorks Enterprise Virtual Array. Data in an original unit is replicated to data in a replica unit. Original units are on systems interpreted as physical drives

(Windows) or physical volumes (UNIX).

See also **replica unit**, **original storage**, and **replica storage version**.

original storage (*ZDB specific term*)

A set of original units that contain the backup objects selected in one Data Protector backup specification. Data in an original storage is replicated to data in a replica storage version by replicating the set of original units. An original storage is typically used by the application system.

See also **original unit**, **replica unit**, and **replica storage version**.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.

See also **merging**.

OVO

HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were

Glossary

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.

package (*MC/ServiceGuard and Veritas Cluster specific term*)

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

pair status (*HP StorageWorks Disk Array XP specific term*)

A mirrored pair of disks can have

various status values depending on the action performed on it. The three most important status values are:

- **COPY** - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- **PAIR** - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- **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

Glossary

data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism

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

physical device

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

post-exec

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

See also **pre-exec**.

pre- and post-exec commands

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

prealloc list

A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec

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

See also **post-exec**.

Primary Volume (P-VOL)

(HP StorageWorks Disk Array XP specific term)

Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.

See also **Secondary Volume (S-VOL)**.

Private Information Store

(MS Exchange 5.5 Server specific term)

A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection

See **data protection** and also **catalog protection**.

Glossary

public folder store (*MS Exchange 2000 Server specific term*)

The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID

Redundant Array of Inexpensive Disks.

RAID Manager Library (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup

See disk image backup.

RCU (*HP StorageWorks specific term*)

The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS

Relational Database Management System.

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

Glossary

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 Oracle8/9 tables and views that are used by Recovery Manager to store information about Oracle8/9 databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle8/9 databases. The recovery catalog contains information about:

- The physical schema of the Oracle8/9 target database
- Data file and archivelog 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,

and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) (*Oracle specific term*)

An Oracle8/9 command-line interface that directs an Oracle8/9 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 (*HP StorageWorks Disk Array XP specific term*)

The Remote Control Unit (RCU) acts as a slave of an MCU in a CA

Glossary

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 unit (*ZDB specific term*)

A logical unit that is used as a target for data replication using snapshot or split mirror technologies. Depending on the vendor and technology used, a replica unit denotes S-VOL on HP StorageWorks Disk Array XP, child (BC) LUN on HP StorageWorks Virtual Array, logical drive on HP StorageWorks Modular SAN Array

1000, or virtual disk on HP StorageWorks Enterprise Virtual Array. Data in an original unit is replicated to data in a replica unit. Replica units are on systems interpreted as physical drives (Windows) or physical volumes (UNIX). A replica unit is also referred to as snapshot or mirror. *See also* **original unit**, **original storage**, and **replica storage version**.

replica storage version (*ZDB specific term*)

A set of replica units, created or reused during one ZDB backup session, which contain replica copies of the backup objects selected in one Data Protector backup specification. Data in an original storage is replicated to data in a replica storage version. A replica storage version is typically used by the backup system. *See also* **original unit**, **replica unit**, and **original storage**.

replica storage pool (*ZDB specific term*)

A number or group of replica storage versions produced during ZDB sessions to be used for the purpose of replica storage rotation, instant recovery, and split mirror restore. The replica storage versions in the replica storage pool are all created using the same backup specification. The size of a replica storage pool is defined for each backup specification as the maximum number

Glossary

of replica storage versions that are to be kept on a disk array before the oldest replica storage version for the backup specification is reused.

See also **replica storage rotation**.

replica storage rotation (*ZDB specific term*)

A ZDB process that denotes either a reuse of the oldest replica storage version in the replica storage pool whenever the size of the replica storage pool is reached or, if the size of the replica storage pool is not reached, a creation of a new replica storage version in the replica storage pool.

See also **replica storage pool**.

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.

SA Agent (*HP StorageWorks Modular SAN Array 1000 specific term*)

A Data Protector software module that executes all tasks required for the HP StorageWorks Modular SAN Array 1000 snapshot integration on the application system and the backup system. It communicates with the HP StorageWorks Modular SAN Array 1000 Business Copy Manager to control the HP StorageWorks Modular SAN Array 1000.

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

This environment variable is set by Data Protector during actual backup sessions (not during preview). It identifies a session and is recorded in the database.

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 (e.g. 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.

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) is a part of the IDB that stores session messages generated during backup and restore sessions. One binary file is created per session. The files are grouped by year and month.

snapshot (*ZDB specific term*)

See **replica unit**.

snapshot backup (*ZDB specific term*)

A ZDB term encompassing ZDB disk backup, ZDB tape backup and ZDB disk/tape backup utilizing snapshot technology.

See also **zero downtime backup (ZDB)**.

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 medium

When copying media, the source medium is the medium that contains backed up data and is being copied.

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 backup (*EMC Symmetrix specific term*)

See **ZDB tape backup**.

Glossary

split mirror backup (*HP StorageWorks Disk Array XP specific term*)
See **ZDB tape backup**, **ZDB disk/tape backup** and **ZDB disk backup**.

split mirror restore (*HP StorageWorks Disk Array XP specific term*)
A process where data backed up using the ZDB tape backup or ZDB disk/tape backup process is restored from tape media to the replica storage version selected by the replica rotation process or by the user. The replica storage version is then synchronized to the original storage. Split mirror restore is limited to filesystem restore.
See also **ZDB tape backup**, **ZDB disk/tape backup**, and **replica storage rotation**.

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

Glossary

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.

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

(MS Exchange 2000 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.

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

Symmetrix Application Programming Interface (SYMAPI) *(EMC Symmetrix specific term)*

A linkable library of functions that can interface with EMC Symmetrix units attached to the Data Protector clients. Provided by EMC.

Symmetrix CLI Database File
(EMC Symmetrix specific term)

The EMC Symmetrix database file that stores EMC Symmetrix configuration data on each system with a configured EMC Symmetrix ICDA and installed SYMCLI.

Symmetrix Command-Line Interface (SYMCLI) *(EMC Symmetrix specific term)*

An application written using the Symmetrix Application Programming Interface (SYMAPI) that retrieves data from an EMC Symmetrix unit using special low-level SCSI commands. The SYMCLI allows you to run commands on the client to obtain configuration, status, and performance data from the EMC Symmetrix units attached to clients that are running in an open systems environment.

System Backup to Tape *(Oracle specific term)*

An Oracle interface that handles the actions required to load, label, and

unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*

The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system 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

Glossary

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.

tapeless backup (*HP StorageWorks Disk Array XP specific term*)

See **ZDB disk backup**.

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 medium

When copying media, the target medium is the medium to which data is copied.

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.

Glossary

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 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 (*MS Exchange and Lotus Domino Server specific term*)

Files in which changes made to a database are recorded.

transaction logs (*Data Protector specific term*)

Keeps track of IDB changes. The

Glossary

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.

TSANDS.CFG file (*Novell NetWare specific term*)

A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM/TSA directory on the server where TSANDS.NLM is loaded.

unattended operation

See lights-out operation.

user account

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

user disk quotas

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

user group

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

user profile (*Windows specific term*)

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

user rights

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

Glossary

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' 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 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 Enterprise Virtual Array 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 **original unit** and **replica unit**.

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.

Glossary

volume mountpoint (*Windows specific term*)

An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service

See **Microsoft Volume Shadow Copy service**.

VPO

See **OVO**.

VSS

See **Microsoft Volume Shadow Copy service**.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN

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

Web reporting

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

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 database repository about a computer's configuration.

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

writer

(*MS VSS specific term*)

A process that initiates change of data on the original volume. Writers are typically applications or system services

Glossary

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 disk backup (*ZDB specific term*)

The basic concept of ZDB disk backup is to create a copy of data from the

original storage at a specific point-in-time, and keep this copy of data in the disk array in the replica storage version selected from or created in the replica storage pool. Data in the replica storage version is not moved to standard backup media. The data backed up utilizing the ZDB disk backup functionality can be either restored by utilizing the instant recovery process or used for data mining and similar purposes.

See also **zero downtime backup (ZDB), ZDB tape backup, ZDB disk/tape backup, instant recovery, and replica storage pool**.

ZDB disk/tape backup (*ZDB specific term*)

The basic concept of ZDB disk/tape backup is to create a copy of data from the original storage at a specific point-in-time, and keep this copy of data in the replica storage version. The copy of data in the replica storage version is additionally used for a backup to a standard backup medium, typically a tape. The data backed up using the ZDB disk/tape backup can be restored using the instant recovery or the standard Data Protector restore procedure. It can also be used for data mining and similar purposes.

See also **zero downtime backup (ZDB), ZDB disk backup, ZDB tape backup, instant recovery, and replica storage pool**.

Glossary

ZDB part of the IDB (*ZDB specific term*)

A part of the IDB, storing ZDB related information such as original and replica storage versions, security information and other. The ZDB part of the IDB is used for ZDB, instant recovery, and split mirror restore.

See also **zero downtime backup (ZDB)**.

ZDB tape backup (*ZDB specific term*)

The basic concept of ZDB tape backup is to create a copy of data from the original storage at a specific point-in-time, and use this copy of data in the replica storage version for a backup to a standard backup medium, typically a tape. After the backup is complete, the data in the replica storage version may be overwritten. Instant recovery is not possible from such a backup, the data must be restored following the standard Data Protector restore procedure.

See also **zero downtime backup (ZDB)**, **ZDB disk backup**, **instant recovery**, **ZDB disk/tape backup**, and **replica storage pool**.

zero downtime backup (ZDB)

A backup process utilizing data replication technologies (the split mirror and snapshot technologies) to minimize the backup window for the application system; typically to few minutes. With this technique, application database downtime (offline backup) or backup

mode (online backup) is limited to the very short time it takes to split the mirror disks or to create or reuse snapshots. The application is then returned to normal operation, while the data in the replica storage version is either backed up by streaming the data to tape (ZDB tape backup) or kept in the replica storage pool (ZDB disk backup) for the instant recovery or other purposes or both (ZDB disk/tape backup).

See also **ZDB disk backup**, **ZDB tape backup**, **ZDB disk/tape backup**, and **instant recovery**.

Glossary

A

- aborting
 - Informix running session, 380
 - SAP R/3 backup session, 168
 - SAP R/3 running session, 168
 - Sybase running session, 295
- advanced options, NDMP, 452
- amount of log
 - specific Lotus Domino R5 Server backup option, 513
- application specific options
 - Informix, 367
- Archive
 - Oracle8/9 backup specification template, 38
- archive logging
 - Lotus Domino R5 Server, 497
- Archive_Delete
 - Oracle8/9 backup specification template, 38
- Archived Redo Logs only
 - Oracle8/9 backup, 3

B

- backing up clusters
 - Informix database, 404
 - SAP R/3 database, 182
 - Sybase database, 322
- backing up EMC Celerra NAS device data, 449
- backing up Informix
 - configuration files, 387
 - configuring backup, 360
 - continuous backup, 388
 - database, 378
 - database, interactively, 383
 - onbar utility, 387
 - online, 387
 - quiescent mode, 387
 - using log_full.sh script, 388
 - whole-system, 369
- backing up Lotus Domino R5 Server, 506
 - interactive backup, 515
- backing up NDMP server data, 449
- backing up NetApp NAS device data, 449
- backing up NNM, 474
 - interactive backup, 476
 - online, 465
- backing up Oracle8/9, 52
 - all archived logs and tablespaces, 63
 - allowing some corrupted blocks, 64
 - Archived Redo Logs only, 3
 - as cluster-aware, 90
 - control file, 3
 - control file, examples, 64
 - incremental, 3
 - interactive backup, 57
 - offline, 3
 - online, 3
 - particular archived logs, 63
 - Recovery Catalog, 3, 54
 - single channel, example, 62
 - three channels in parallel, 62
 - using RMAN, 60
- backing up SAP R/3
 - backup utilities, 123
 - database, 167
 - interactively, 171
 - using RMAN mode, 161
- backing up Sybase
 - database, 294
 - interactively, 299
 - using Sybase commands, 301
- backint
 - backup flow, 125
 - restore flow, 130
- backup buffer size
 - specific Lotus Domino R5 Server backup option, 514
- backup devices
 - Lotus Domino R5 Server, 508
 - NNM, 470
 - Oracle8/9, 37
- backup methods
 - Informix, 378
 - SAP R/3, 167
 - Sybase, 294
- Backup mode
 - SAP R/3 backup options, 159
- backup modes
 - SAP R/3, 167
- Backup Objects
 - SAP R/3 backup option, 157
- backup options
 - Lotus Domino R5 Server, 512
 - NNM, 473
 - Oracle8/9, 42
 - SAP R/3, 155
- backup problems
 - Informix, 408
 - Lotus Notes integration, 530
 - Oracle8/9, 100

Index

- SAP R/3, 190
- Sybase, 326
- backup procedure
 - Lotus Domino R5 Server, 489
 - NetApp NAS device, 449
 - NNM, 475
 - Oracle8/9, 54
 - SAP R/3, using backint, 125
 - Sybase, 301
- backup specification
 - configuring Lotus Domino R5 Server, 508
 - configuring NNM, 470
 - configuring Oracle8/9, 37
 - creating SAP R/3, 150
 - editing Sybase, 290
 - Informix restore option, 395
 - Lotus Domino R5 Server, 514
 - NNM, 475
 - Oracle8/9, 55
 - ownership, Informix, 350
 - ownership, Oracle8/9, 10
 - ownership, SAP R/3, 140
 - ownership, Sybase, 268
 - scheduling new, Sybase, 285
- backup templates
 - configuring new in SAP R/3, 150
- backup types
 - Informix, 378
 - Lotus Domino R5 Server, 486
 - NNM, 465
 - Oracle8/9, 3
 - Sybase, 257, 295
- balancing by load
 - SAP R/3 backup option, 157
- balancing by time
 - SAP R/3 backup option, 157
- balancing manual
 - SAP R/3 backup options, 158
- benefits
 - Informix, 340
 - Lotus Domino R5 Server, 486
 - NDMP server, 423
 - NNM, 465
 - Oracle8/9, 4
 - SAP R/3 integration, 162
 - Sybase, 258
- BOX files, 506
- BRARCHIVE
 - SAP R/3 backup option, 157
 - SAP R/3 backup utility, 123
- BRBACKUP
 - SAP R/3 backup option, 156
 - SAP R/3 backup utility, 123
 - starting SAP R/3 backup, 173
- BRRESTORE
 - command, 177
 - restore flow, 131
 - SAP R/3 backup utility, 124
- C**
 - changing
 - Informix user, 368
 - Sybase user, 284
 - checking
 - Lotus Domino R5 Server client
 - configuration, 502
 - Oracle8/9 client configuration, 33, 36
 - SAP R/3 configuration, 147, 148
 - Sybase configuration, 277
 - circular logging
 - Lotus Domino R5 Server, 497
 - cluster environment
 - Informix, 404
 - Oracle8/9, 89
 - SAP/3, 181
 - Sybase, 321
 - cluster-aware Informix
 - backing up database, 404
 - restoring database, 404
 - cluster-aware SAP R/3
 - backing up database, 182
 - restoring database, 182
 - cluster-aware Sybase
 - backing up database, 322
 - CONFIG, NDMP interface, 430
 - configuration file
 - SAP R/3, 132
 - configuration file example
 - SAP R/3, 134
 - configuration file syntax
 - SAP R/3, 133
 - configuration files
 - Lotus Domino R5 Server, 498
 - Oracle8/9, 13
 - configuration problems
 - Informix, 406
 - Lotus Notes integration, 530
 - Oracle8/9 integration, 97
 - SAP R/3, 188
 - Sybase, 323

configuring
 Data Protector client, 273
 Informix backup, 360
 Informix integration, 348
 Informix user, 350
 Lotus Domino R5 Server, 494, 501
 Lotus Domino R5 Server backup, 508
 Lotus Domino R5 Server client, 502
 Lotus Domino R5 Server user, 500
 Lotus Notes integration, 494, 497
 NDMP backup device, 441
 NDMP network appliance, 446
 NDMP server integration, 435
 NNM backup, 470
 OnLine Server, 353, 354
 Oracle8/9 backup, 37
 Oracle8/9 client, 33, 36
 Oracle8/9 integration, 20
 Oracle8/9 Server, 29, 34
 Oracle8/9 user, 27
 SAP R/3 backup, 150, 162
 SAP R/3 Database Server, 142
 SAP R/3 integration, 140
 SAP R/3 user, 140
 Sybase backup, 279
 Sybase integration, 265
 Sybase Server, 270
 Sybase user, 268
 user rights, 269
configuring Oracle8/9 Server
 on HP-UX, 34
 on SUN Solaris, 34
CONNECT, NDMP interface, 430
continuous backup
 Informix database, 388
control file
 Oracle8/9 backup, 3
 restore, 70
conventions, xv
creating
 Informix backup specification, 361
 Lotus Domino R5 Server backup
 specification, 508
 new backup template, NNM, 471
 new backup template, Oracle8/9, 37
 NNM backup specification, 471
 Oracle8/9 backup specification, 38
 SAP R/3 backup specification, 163
 Sybase backup specification, 280

D

Data Protector Command Line Interface
 configuring Oracle8/9 Server, 34
 SAP R/3, 176
Data Protector command-line interface
 configuring Lotus Domino R5 Server, 501
Data Protector Database Library, 10, 261
 linking with Oracle8/9, 21
Data Protector database, definition, 12, 125
Data Protector Restore GUI for Oracle, 66
Data Protector user rights, 269
database items restore, 66
database state
 restore, 67
Database_Archive
 Oracle8/9 backup specification template, 38
Database_Switch_Archive
 Oracle8/9 backup specification template, 38
Database_Switch_ArchiveDel
 Oracle8/9 backup specification template, 39
Direct_Database
 Oracle8/9 backup specification template, 39
disaster recovery
 data restore order, 67
 Informix database, 402
 NNM database, 477
 restoring databases after, 66
 SAP R/3 database, 178
 Sybase database, 317
Disk Agent
 installing Lotus Domino R5 Server, 492
 installing NNM, 469
 installing Oracle8/9, 18
 installing SAP R/3, 138
 installing Sybase, 263

E

editing
 Informix backup specification, 371
 Sybase backup specification, 290
EMC Celerra NAS device
 configuring backup device, 441
 creating media pool, 441
 supported configurations, 436
 supported drives, 435
 supported library devices, 435
EMC Celerra NAS device data
 backing up, 449
 restoring, 454
emergency boot file, 346

Index

F

- file history setting, NDMP, 459
- FILEHISTORY, NDMP interface, 431
- filesystem backup
 - Lotus Domino R5 Server, 494
 - Oracle8/9 Server, 29
- finding
 - information for Informix restore, 392
 - Oracle user, 141
 - Oracle8/9 user, 28
 - SAP R/3 user, 141
- full backup
 - Sybase, 257

G

- global configuration file
 - Lotus Domino R5 Server, 498
 - Oracle8/9, 13
- global configuration file example
 - Lotus Domino R5 Server, 498
 - Oracle8/9, 14
- global configuration file syntax
 - Lotus Domino R5 Server, 498
 - Oracle8/9, 14

I

- incremental backup
 - Informix, 386
 - NNM, 475
 - Oracle8/9, 56
 - SAP R/3, 167
 - SAP R/3, using RMAN, 162
- Informix backup specification
 - checking, 358
 - creating, 361
 - ownership, 350, 351
 - scheduling, 381
 - testing, 374
- Informix client
 - configuring, 353
- Informix commands, 386
- Informix configuration
 - checking, 358
 - options, 353
- Informix database
 - backing up, 378
 - disaster recovery, 402
 - restoring, 390
- Informix integration
 - as cluster-aware application, 404
 - benefits, 340
 - check before configuring, 348
 - concepts, 344
 - configuring, 348
 - installing, 346
 - limitations, 343
 - overview, 339
 - prerequisites, 342
 - restoring a database, 390
 - troubleshooting, 406
- Informix Integration software component, 346
- Informix objects
 - listing, 390
- Informix sessions
 - monitoring, 403
- Informix user
 - changing, 368
 - configuring, 350
- installation software
 - for Informix integration, 346
 - for Lotus Notes integration, 492
 - for NNM integration, 469
 - for Oracle8/9 integration, 18
 - for SAP R/3 integration, 138
 - for Sybase integration, 263
- installing
 - Informix integration, 346
 - Lotus Domino R5 Server, software components, 492
 - Lotus Notes integration, 492
 - NDMP server Integration, 434
 - NNM integration, 469
 - NNM, software components, 469
 - Oracle8/9 as cluster-aware, 89
 - Oracle8/9 integration, 18
 - Oracle8/9, software components, 18
 - SAP R/3, 138
 - Sybase integration, 263
- instance configuration file
 - Oracle8/9, 13
- instance configuration file example
 - Oracle8/9, 15
- instance configuration file syntax
 - Oracle8/9, 14
- integration concepts
 - Informix, 344
 - Lotus Domino R5 Server, 489
 - NDMP server, 427

- NNM, 467
 - Oracle8/9, 8
 - SAP R/3, 123
 - Sybase, 261
 - interactive backup
 - Informix, 383
 - Lotus Domino R5 Server, 515
 - NNM, 476
 - Oracle8/9, 57
 - SAP R/3, 171
 - Sybase, 257, 299
 - interfaces
 - NDMP, 429
 - is, 498
 - isql utility, 257, 262

 - L**
 - language environments support
 - Sybase, 320
 - limitations
 - Informix, 343
 - Lotus Domino R5 Server, 488
 - NDMP server, 425
 - NetApp NAS device, 449, 454
 - NNM, 466
 - Oracle8/9, 6
 - SAP R/3, 121
 - Sybase, 260
 - linking
 - Oracle8/9 with the Database Library, 21
 - linking library to Lotus C API, 499
 - listing Oracle8/9 configuration files
 - parameters, 17
 - listing SAP R/3 configuration file
 - parameters, 137
 - load command, 305
 - log file
 - SAP R/3 backup option, 156
 - LOGGING, NDMP interface, 431
 - Lotus C API, 491
 - linking library, 499
 - Lotus Domino R5 Server
 - configuring, 494
 - configuring by using the Data Protector command-line interface, 501
 - configuring by using the Data Protector GUI, 501
 - restoring, 517
 - Lotus Domino R5 Server backup, 506
 - configuring, 508
 - Lotus Domino R5 Server backup
 - specification, 508
 - creating, 508
 - scheduling, 514
 - Lotus Domino R5 Server client
 - checking configuration, 502
 - Lotus Domino R5 Server database
 - backing up, 506
 - Lotus Domino R5 Server sessions
 - monitoring, 487, 523
 - Lotus Domino R5 Server user
 - configuring, 500
 - Lotus Notes integration
 - backing up, 506
 - backup devices, 508
 - backup flow, 489
 - backup specification, 514
 - backup types, 486
 - concepts, 489
 - concepts, scheme, 490
 - configuration overview, 494
 - configuring, 494, 497
 - installation software, 492
 - installing, 492
 - limitations, 488
 - monitoring backup and restore, 523
 - prerequisites, 488
 - restore flow, 490
 - testing, 502
 - troubleshooting, 526
 - Lotus notes integration
 - benefits, 486
 - Lotus Notes integration agent, 485
 - Lotus Notes Integration software component, 492, 493

 - M**
 - MA. *See* Media Agent
 - manual balancing
 - creating SAP R/3 backup specification, 163
 - SAP R/3 files, 163
 - media
 - Lotus Domino R5 Server, 508
 - NNM, 470
 - Oracle8/9, 37
 - Media Agent
 - in Oracle8/9 backup procedure, 11
 - installing Informix, 346
 - installing Lotus Domino R5 Server, 492
 - installing Oracle8/9, 18
-

Index

- installing SAP R/3, 139
- installing Sybase, 263
- media management, NDMP, 458
- media pool
 - EMC Celerra NAS device, 441
 - Lotus Domino R5 Server, 508
 - NetApp NAS device, 441
 - NNM, 470
 - Oracle8/9, 37
- modifying
 - parameter file on SAP R/3 database, 160
- monitoring
 - Informix sessions, 403
 - Lotus Domino R5 Server current sessions, 523
 - Lotus Domino R5 Server previous sessions, 524
 - Lotus Domino R5 Server sessions, 523
 - NNM sessions, 479
 - Oracle8/9 sessions, 83
 - SAP R/3 sessions, 180
 - Sybase restore sessions, 309
 - Sybase sessions, 319
- MOVER, NDMP interface, 430
- N**
- NDMP**
 - advanced options,setting, 452
 - client, 423, 427
 - configuring backup device, 441
 - creating media pool, 441
 - importing server host, 439
 - interfaces, 429, 431
 - media, 428
 - Media Agent, 433
 - network appliance, configuring, 446
 - omnirc variables, 459
 - overview, 429
 - restore options,setting, 455
 - server, 427
 - Session Manager, 432
 - supported configurations, 436
 - supported drives, 435
 - supported library devices, 435
- NDMP environment configuration, 428
- NDMP environment, schematic view, 432
- NDMP Media Agent
 - installing, 434
 - integration module, 433
- NDMP omnirc variables
 - OB2NDMPFH, 459
- NDMP server
 - Add/Remove Mountpoint, 451
 - file history setting, 459
- NDMP server data
 - backing up, 449
 - restoring, 454
- NDMP server integration
 - benefits, 423
 - concepts, 427
 - configuring, 435
 - installing, 434
 - limitations, 425
 - media management, 458
 - overview, 423
 - prerequisites, 425
 - troubleshooting, 460
- NetApp NAS device
 - configuring backup device, 441
 - creating media pool, 441
 - supported configurations, 436
 - supported drives, 435
 - supported library devices, 435
- NetApp NAS device data
 - backing up, 449
 - restoring, 454
- NetApp NAS device integration
 - limitations, 449, 454
 - restore procedure, 454
 - restoring using another device, 457
 - single file restore, 455
- Network Attached Storage (NAS) devices, 423
- NNM backup
 - configuring, 470
 - of database, 474
- NNM backup specification, 470
 - creating, 471
 - creating a new template, 471
 - scheduling, 475
- NNM database
 - backing up, 474
 - disaster recovery, 477
 - restoring, 477
- NNM integration
 - backing up a database, 474
 - backup devices, 470
 - backup procedure, 475
 - backup specification, 475

- backup types, 465
- benefits, 465
- concepts, 467
- installation software, 469
- installing, 469
- limitations, 466
- monitoring backup and restore, 479
- overview, 465
- prerequisites, 466
- restore types, 465
- restoring a database, 477
- software components, 468
- testing, 473
- troubleshooting, 480

NNM Integration software component, 469

NNM sessions

- monitoring, 465, 479

Notes Data, definition, 491

Notes Storage Facility files, 506

Notes Template Facility files, 506

NOTIFY, NDMP interface, 430

NSF files. *See* Notes Storage Facility files

NTF files. *See* Notes Template Facility files

O

OB2NDMPFH

- NDMP `omnirc` variables, 459

Objects outside database

- SAP R/3 backup option, 159

offline backup

- of NNM database, 474
- of Oracle8/9 database, 3, 52

omnidb command, 303

- listing Informix objects, 390

`omnirc` variables, NDMP, 459

onbar utility

- backup, 387

on-demand backup

- Informix database, 388

online backup

- Informix, 387
- of NNM, 465
- of NNM database, 474
- of Oracle8/9 database, 3, 52
- Oracle8/9 incremental, 3

OnLine Server, 379

- configuring, 353, 354

Options Pane, 77

Oracle databases

- methods for restoring, 66

Oracle8/9 backup

- configuring, 37
- of database, 52

Oracle8/9 backup specification, 37

- creating, 38
- creating a new template, 37
- pre-defined templates, 38
- scheduling, 55

Oracle8/9 client

- checking configuration, 36

Oracle8/9 configuration files, 13

Oracle8/9 configuration files parameters

- listing, 17
- retrieving, 17
- setting, 17
- setting, retrieving and listing, 15

Oracle8/9 Database

- after recovering, 74

Oracle8/9 database

- backing up, 52
- preparing for restore, 107
- recovering, 68

Oracle8/9 global configuration file, 13

- Oracle8/9 global configuration file example, 14

Oracle8/9 global configuration file syntax, 14

Oracle8/9 instance configuration file, 13

- Oracle8/9 instance configuration file example, 15

- Oracle8/9 instance configuration file syntax, 14

Oracle8/9 integration

- as cluster-aware application, 89
- backing up a database, 52
- backup devices, 37
- backup procedure, 54
- backup specification, 55
- backup specification templates, 38
- backup types, 3
- benefits, 4
- concepts, 8
- concepts, scheme, 12
- configuration overview, 20
- configuring, 20
- installation software, 18
- installing, 18
- limitations, 6
- linking with the Database Library, 21
- monitoring backup and restore, 83
- overview, 3

Index

- prerequisites, 5
- removing, 84
- restore flow, 11
- restore types, 3
- software components, 8
- testing, 49
- troubleshooting, 92
- Oracle8/9 Integration software component, 18
- Oracle8/9 Recovery Catalog, 8
 - restoring, 80
- Oracle8/9 Recovery Manager, 8
 - backing up, 60
 - restoring, 80
- Oracle8/9 restore
 - archive log, 114
 - control file, 115
 - datafile, 113
 - example, 107
 - full database, 109
 - point in time, 110
 - tablespace, 111
- Oracle8/9 RMAN metadata, 87
- Oracle8/9 RMAN script, 44
- Oracle8/9 RMAN scripts, examples, 62, 63
- Oracle8/9 Server
 - configuring, 29
 - configuring by using Data Protector CLI, 34
 - configuring by using Data Protector GUI, 29
 - instance, 34
- Oracle8/9 sessions
 - monitoring, 4, 83
- Oracle8/9 user
 - configuring, 27
- overview
 - Informix integration, 339
 - Lotus Domino R5 Server, 485
 - NDMP, 429
 - NDMP server integration, 423
 - NNM, 465
 - Oracle8/9, 3
 - SAP R/3, 119
 - Sybase integration, 257
- ownership
 - of Informix backup specification, 350, 351
 - of Oracle8/9 backup specification, 10
 - of SAP R/3 backup specification, 140
 - of Sybase backup specification, 268

P

- parameter file
 - SAP R/3, 160
- post-exec
 - Informix object specific command, 367
 - Sybase object specific command, 283
- post-exec command
 - SAP R/3 backup options, 159
 - specific Lotus Domino R5 Server backup option, 514
 - specific Oracle8/9 backup option, 44
- pre-exec
 - Informix object specific command, 367
 - Sybase object specific command, 283
- pre-exec command
 - SAP R/3 backup options, 158
 - specific Lotus Domino R5 Server backup option, 513
 - specific Oracle8/9 backup option, 44
- preparing
 - Oracle8/9 database for restore, 107
- prerequisites
 - Informix, 342
 - Lotus Domino R5 Server, 488
 - NDMP server, 425
 - NNM, 466
 - Oracle8/9, 5
 - SAP R/3, 121
 - Sybase, 259

Q

- quiescent backups, 387

R

- recovering
 - Informix database, 402
 - NNM, 477
 - Oracle8/9 database, 68
 - SAP R/3 database, 178
 - Sybase database, 317
- Recovery Catalog
 - backup, Oracle8/9, 3, 54
- recovery catalog database
 - restore, 68
- Recovery Manager (RMAN)
 - SAP R/3 backup flow, 128
- recovery problems
 - Lotus Notes integration, 535
- removing Oracle8/9 integration, 84

- Restore
 - methods for restoring Oracle databases, 66
 - restorable Oracle items, 66
 - restore
 - control file, 70
 - database items, 66
 - database state, 67
 - database to another host, 71
 - Oracle database objects, 71
 - recovery catalog database, 68
 - tablespaces and datafiles, 75
 - restore by date
 - Informix restore option, 396
 - restore by log number
 - Informix restore option, 396
 - restore methods
 - Informix, 390
 - restore options
 - Informix, 395
 - Lotus Domino R5 Server, 520
 - NDMP, 455
 - restore problems
 - Informix integration, 417
 - Lotus Notes integration, 534
 - Oracle8/9, 103
 - SAP R/3 integration, 193
 - Sybase integration, 331
 - restore procedure
 - Lotus Domino R5 Server, 490
 - NetApp NAS device, 454
 - Oracle8/9, 11
 - SAP R/3, using backint, 130
 - SAP R/3, using RMAN, 131
 - Sybase, 303
 - restore the latest version
 - Informix restore option, 396
 - restore to client
 - Informix restore option, 396
 - restore types
 - NDMP server, 423
 - NetApp NAS device, 455, 457
 - NNM, 465
 - Oracle8/9, 3
 - Sybase, 257
 - restoring
 - archive logs, example, 204
 - cluster-aware Informix database, 404
 - cluster-aware SAP R/3 database, 182
 - EMC Celerra NAS device data, 454
 - Informix database, 390
 - Informix, methods, 390
 - Informix, to another client, 400
 - Informix, using another device, 401
 - Informix, using omnidb command, 390
 - lost files, example, 203
 - Lotus Domino R5 Server, 517
 - NDMP server data, 454
 - NetApp NAS device data, 454
 - NetApp NAS device, single file, 455
 - NetApp NAS device, using another device, 457
 - NNM database, 477
 - SAP R/3 database, 174
 - SAP R/3 database, example, 199
 - SAP R/3, using another device, 177
 - Sybase database, 303
 - Sybase from another device, 316
 - restoring Oracle8/9
 - archive log, example, 114
 - as cluster-aware, 90
 - control file, example, 115
 - datafile, example, 113
 - full database, example, 109
 - point in time, example, 110
 - Recovery Catalog, 80
 - tablespace, example, 111
 - using another device, 81
 - using RMAN, 80
 - restoring Oracle8/9 database
 - example, 107, 109
 - retrieving Oracle8/9 configuration files
 - parameters, 17
 - retrieving SAP R/3 configuration file
 - parameters, 137
 - RMAN. *See* Oracle8/9 Recovery Manager
 - See also* Recovery Manager
 - running
 - Lotus Domino R5 Server interactive
 - backup, 515
 - NNM interactive backup, 476
 - Oracle8/9 interactive backup, 57
 - SAP R/3 interactive backup, 171
 - Sybase interactive backup, 299
- ## S
- SAP R/3 backup
 - configuring, 150
 - SAP R/3 backup specification
 - scheduling, 168

Index

- SAP R/3 backup utilities, 123
- SAP R/3 commands, 177
- SAP R/3 configuration file, 132
- SAP R/3 configuration file example, 134
- SAP R/3 configuration file parameters
 - listing, 137
 - retrieving, 137
 - setting, 137
 - setting, retrieving and listing, 135
- SAP R/3 configuration file syntax, 133
- SAP R/3 data objects, 123
- SAP R/3 database
 - backing up, 167
 - configuring, 142
 - creating the parameter file, 160
 - disaster recovery, 178
 - restoring, 174
 - restoring using another device, 177
- SAP R/3 integration
 - architecture, backint mode, 126
 - architecture, RMAN mode, 130
 - as cluster-aware application, 181
 - backing up a database, 167
 - backup concept, scheme, 125
 - concepts, 123
 - configuring, 140
 - installing, 138
 - overview, 119
 - prerequisites, 121
 - restoring a database, 174
 - software components, 138
 - starting backup using commands, 173
 - testing, 165
 - troubleshooting, 183
 - upgrading, 138
- SAP R/3 restore
 - archive log restore, 204
 - example, 197, 199
 - full database restore, 199
 - partial, 202
 - preparing database, 197
 - restoring lost files, 203
- SAP R/3 sessions
 - monitoring, 180
- sapdba
 - command, 177
 - SAP R/3 administration utility, 124
 - starting SAP R/3 backup, 173
- scheduling
 - Informix backup specification, 381
 - Lotus Notes integration, 514
 - NNM integration, 475
 - Oracle8/9 integration, 55
 - SAP R/3 backup specification, 168
 - Sybase backup specification, 231, 285, 296
- SCSI, NDMP interface, 430
- server specific configuration file
 - Lotus Domino R5 Server, 498
- server specific configuration file example
 - Lotus Domino R5 Server, 499
- server specific configuration file syntax
 - Lotus Domino R5 Server, 498
- Session Manager
 - NDMP integration module, 432
- Session Manager, definition, 12, 125, 126, 490
 - setting
 - advanced options, NDMP, 452
 - file history, NDMP, 459
 - setting Oracle8/9 configuration files
 - parameters, 17
 - setting SAP R/3 configuration file
 - parameters, 137
- SM. *See* Session Manager
- SMB_Database
 - Oracle8/9 backup specification template, 39
- SMB_Database_Oracle8
 - Oracle8/9 backup specification template, 39
- Source Pane, 76
- specific backup options
 - Lotus Domino R5 Server, 512
 - NNM, 473
 - Oracle8/9, 42
- specific restore destination options
 - Lotus Domino R5 Server, 520
- sqlhosts file, 346
- supported configurations, NDMP, 436
- supported drives, NDMP, 435
- supported library devices, NDMP, 435
- syb tool command, 310
- Sybase backup
 - concept scheme, 261
 - configuring, 279
- Sybase Backup Server API, 262
- Sybase backup specification
 - creating, 280
 - editing, 290
 - ownership, 268
 - scheduling, 231, 296
- Sybase client
 - checking configuration, 277
- Sybase database

- backing up, 294
 - disaster recovery, 317
 - restoring, 303
- Sybase database objects
 - restoring using another device, 316, 401
- Sybase integration
 - as cluster-aware application, 321
 - benefits, 258
 - concepts, 261
 - configuring, 265
 - installing, 263
 - limitations, 260
 - overview, 257
 - prerequisites, 259
 - software components, 263
 - testing, 291
 - troubleshooting, 323
- Sybase Integration software component, 263
- Sybase language environments support, 320
- Sybase Server
 - backup types, 257
 - configuring, 270
- Sybase sessions
 - aborting running one, 295
 - monitoring, 319
 - monitoring restore, 309
- Sybase user
 - changing them, 284
 - configuring, 268
- system databases
 - frequent backups, 294

T

- tablespaces and datafiles
 - restore, 75
- TAPE, NDMP interface, 430
- testing
 - Informix backup configuration, 374
 - Lotus Notes integration, 502
 - NNM integration, 473
 - Oracle8/9 integration, 49
 - SAP R/3 integration, 165
 - Sybase integration, 291
- transaction backup
 - Sybase, 257
- transaction log files
 - Lotus Domino R5 Server, 506
- transaction logging
 - Lotus Domino R5 Server, 496
- transaction logs

- Lotus Domino R5 Server, 491
- troubleshooting
 - Informix integration, 406
 - Informix-related problems, 413, 420
 - Lotus Notes integration, 526
 - NDMP server integration, 460
 - NNM integration, 480
 - Oracle8/9 integration, 92
 - SAP R/3 integration, 183
 - Sybase integration, 323
- typographical conventions, xv

U

- upgrading
 - SAP R/3, 138
- use default RMAN channels
 - SAP R/3 backup options, 159
- user group
 - Informix restore option, 395
- User Interface
 - installing Informix, 346
 - installing Lotus Domino R5 Server, 492
 - installing Oracle8/9, 18
 - installing Sybase, 263
- user name
 - Informix restore option, 395
- user rights, 352
 - configuring, 269
 - in Data Protector, 269
- util_cmd command
 - Oracle8/9, 15
 - SAP R/3, 135

V

- viewing
 - current Lotus Domino R5 Server sessions, 523
 - previous Lotus Domino R5 Server sessions, 524

W

- whole database restore
 - Informix restore option, 396
- whole online delete
 - Oracle8/9 backup option, 39
- Whole_Online
 - Oracle8/9 backup specification template, 38
- Whole_Online_Delete
 - Oracle8/9 backup specification template, 38

Index

whole-system backup, Informix, 369