HP Data Protector A.06.10

Integration guide for IBM applications

Part number: B6960-96043
First edition: November 2008

Legal and notice information

Printed in the US

Contents

Figures

Tables

Publication history

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

Table 1 Edition history

Part number	Guide edition	Product
B6960-90110	October 2004	Data Protector Release A.05.50
B6960-96009	July 2006	Data Protector Release A.06.00
B6960-96043	November 2008	Data Protector Release A.06.10

About this guide

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

Intended audience

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

- Basic Data Protector functionality
- Database administration

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

Documentation set

Other documents and online Help provide related information.

Guides

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

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

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

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

- *HP Data Protector concepts guide*

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

- *HP Data Protector installation and licensing guide*

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

- *HP Data Protector troubleshooting guide*

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

- *HP Data Protector disaster recovery guide*

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

- *HP Data Protector integration guides*

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

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

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

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

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

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

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

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

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

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

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

- *HP Data Protector integration guide for HP Reporter*

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

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

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

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

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

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

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

- *HP Data Protector zero downtime backup concepts guide*

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

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

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

- *HP Data Protector zero downtime backup integration guide*

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

describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

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

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

- *HP Data Protector Media Operations user's guide*

  This guide provides tracking and management of offline storage media. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.

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

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

- *HP Data Protector product announcements, software notes, and references for integrations to HP Operations Manager, HP Reporter, HP Performance Manager, HP Performance Agent, and HP Service Information Portal*

  This guide fulfills a similar function for the listed integrations.

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

  This guide fulfills a similar function for Media Operations.

- *HP Data Protector command line interface reference*

  This guide describes the Data Protector command-line interface, command options and their usage as well as provides some basic command-line examples.

Online Help

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

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

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

Documentation map

Abbreviations

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

Abbreviation	Guide
CLI	Command line interface reference
Concepts	Concepts guide
DR	Disaster recovery guide
GS	Getting started guide
Help	Online Help
IG-IBM	Integration guide for IBM applications: Informix, DB2, and Lotus Notes/Domino
IG-MS	Integration guide for Microsoft applications: SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service
IG-O/S	Integration guide for Oracle and SAP
IG-OMU	Integration guide for HP Operations Manager for UNIX
IG-OMW	Integration guide for HP Operations Manager for Windows
IG-PM/PA	Integration guide for HP Performance Manager and HP Performance Agent
IG-Report	Integration guide for HP Reporter
IG-SIP	Integration guide for HP Service Information Portal
IG-Var	Integration guide for VMware Virtual Infrastructure, Sybase, Network Node Manager, and Network Data Management Protocol Server

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

Map

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

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

Integrations

Look in these guides for details of the following integrations:

Integration	Guide
HP Operations Manager for UNIX/for Windows	IG-OMU, IG-OMW
HP Performance Manager	IG-PM/PA
HP Performance Agent	IG-PM/PA

Integration	Guide
HP Reporter	IG-R
HP Service Information Portal	IG-SIP
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX system	MPE/iX
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG

Integration	Guide
Sybase	IG-Var
EMC Symmetrix	all ZDB
VMware	IG-Var

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: Table 2 on page 23	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
text	Emphasized monospace text

△ CAUTION:

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

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Data Protector graphical user interface

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

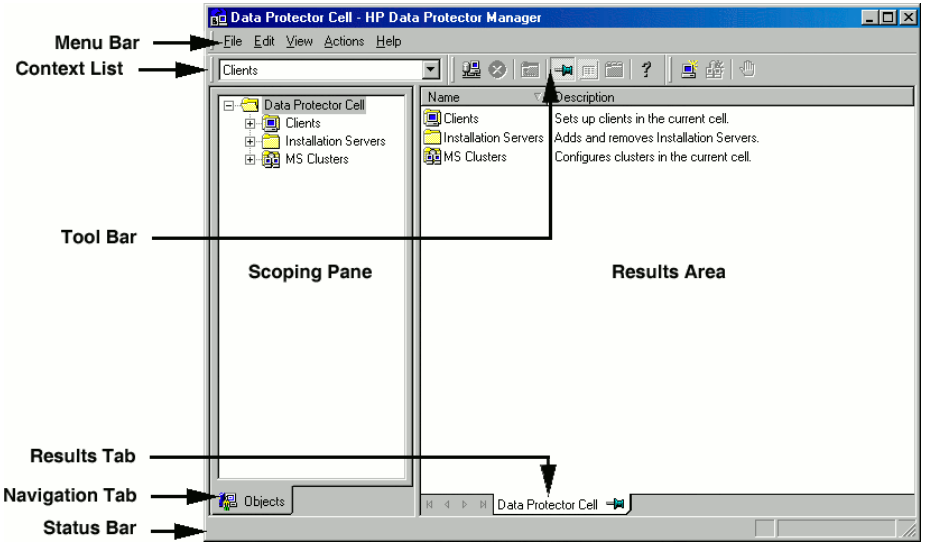

Figure 1 Data Protector graphical user interface

General information

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

HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

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

Documentation feedback

HP welcomes your feedback.

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

1 Integrating Informix Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Informix Server integration. It describes the concepts and methods you need to understand to back up and restore Informix Server database objects (**dbobjects**).

Data Protector integrates with the Informix Dynamic Server (Informix Server) to back up dbobjects online. During backup, a database server (Informix instance) is online and actively used.

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

Table 3 Informix Server backup types

Full	Full backup (level 0).
Incr1	Incremental backup (level 1). Backs up changes since the last Full backup.
Incr2	Incremental backup (level 2). Backs up changes since the last Incr1 backup.

Data Protector offers two types of restore:

Table 4 Informix Server restore types

Complete database restore	Restore from any backup. ON-Bar restores dbobjects concurrently and replays the logical logs once.

Whole-system restore	Restore from a whole-system backup. ON-Bar restores the whole system sequentially with or without restoring the logical logs. Whole-system restore is appropriate for small systems, when you do not need to restore logs, for disaster recovery, or when restoring to another client.

You can also back up and restore dbobjects using the Informix Server `onbar` command.

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

Integration concepts

Data Protector integrates with the Informix Server through the Data Protector Database Library based on a common library called Data Protector **BAR** (Backup And Restore). The Data Protector Database Library channels communication between the Data Protector Session Manager, and, via the **XBSA interface**, the Informix Server **ON-Bar utility**.Figure 2 on page 28 shows the architecture of the Data Protector Informix Server integration.

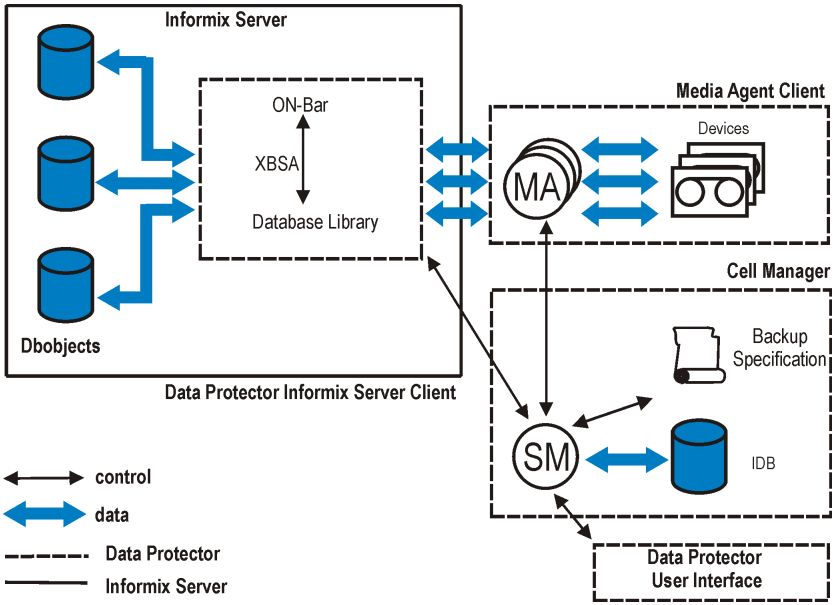

Figure 2 Data Protector Informix Server integration architecture

Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
ON-Bar	ON-Bar executes backup and restore requests from Data Protector and from the Informix Server command line.
XBSA	X/Open Backup Services Application Programmer's Interface, through which ON-Bar and Data Protector exchange control and data.
Database Library	A set of Data Protector executables that enable data transfer between an Informix instance and Data Protector.
MA	Data Protector General Media Agent.
Backup Specification	A list of objects to be backed up, backup devices, and options to be used.
IDB	The Data Protector Internal Database.

Backup is always executed on the Informix Server system via the Informix Server ON-Bar utility. ON-Bar communicates backup and restore requests to the Informix instance.

While an Informix instance is responsible for read/write operations to disk, Data Protector reads from and writes to devices and manages media.

Configuring the integration

You need to configure an Informix Server user and every Informix instance you intend to back up or restore.

Prerequisites

- Ensure that you have correctly installed and configured Informix Server.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on installing, configuring, and using Informix Server, see the Informix Server online documentation.

- Ensure that you have correctly installed Data Protector. For information on how to install Data Protector in various architectures, see the *HP Data Protector product announcements, software notes, and references*.

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

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the Informix Server system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the Informix Server system.
- ***Windows only:*** Change the Data Protector `Inet` service account. Stop the service and restart it as user `informix`. For information, see the online Help index: "changing Data Protector Inet account".

Cluster-aware clients

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

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

Windows: `set OB2BARHOSTNAME=`*`virtual_server_name`*

UNIX: `export OB2BARHOSTNAME=`*`virtual_server_name`*

Configuring Informix Server users

On UNIX, add the Informix Server administrator to the Data Protector `admin` or `operator` user group. For information, see the online Help index: "adding users".

This user is typically `informix` or `root` in the group `informix`. To determine it, check the owner of the Informix Server `onbar_d` file.

This chapter assumes that your Informix Server user is `informix` in the group `informix`.

Configuring Informix instances

You need to provide Data Protector with configuration parameters for the Informix instance:

- Name of the Informix instance.
- Pathname of the Informix Server home directory.
- ***Windows:*** Name of the system with the `sqlhosts` entry in the Windows Registry.

  UNIX: Pathname of the `sqlhosts` file.
- Name of the Informix instance ONCONFIG file.

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

To configure an Informix instance, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the Informix instance is online.

Using the Data Protector GUI

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

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

   In **Application database**, enter the Informix instance name.

   UNIX only: Enter `informix` in both **Username** and **Group name**.

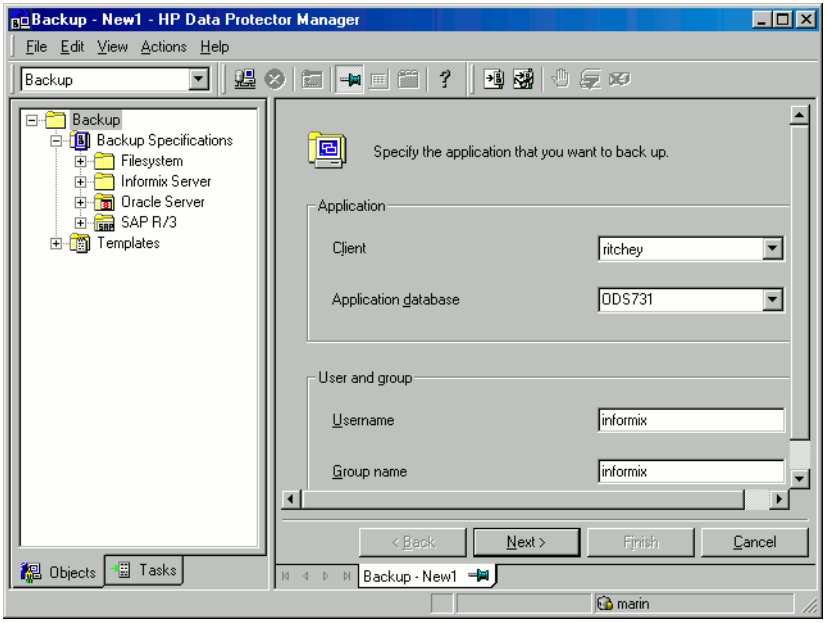

Figure 3 Specifying an Informix instance

Click **Next**.

5. In Informix Server home directory, specify the pathname of the Informix Server home directory.

   In Full pathname of sqlhosts file, enter the following:

   Windows: Name of the system with the `sqlhosts` entry in the Windows Registry. Use the UNC notation, for example: `\\computer_name`.

   UNIX: Pathname of the `sqlhosts` file.

   In **Name of ONCONFIG file**, enter the name of the Informix instance ONCONFIG file, located in the following directory:

   Windows: *`INFORMIXDIR`*`\etc`

   UNIX: *`INFORMIXDIR`*`/etc`

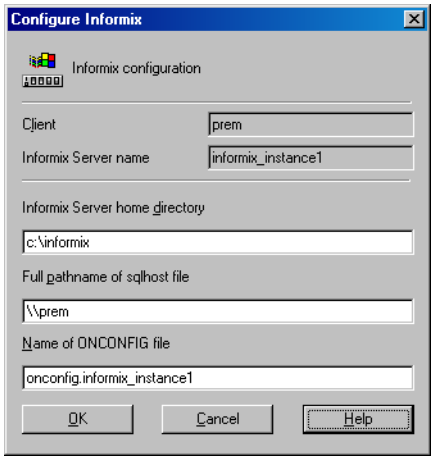

Figure 4 Configuring an Informix instance (Windows)

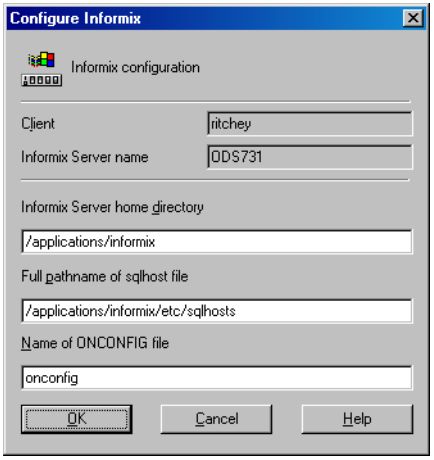

Figure 5 Configuring an Informix instance (UNIX)

Click **OK**.

6. If an error occurs, click **Details** or see "Troubleshooting" on page 66.
7. The Informix instance is configured. Exit the GUI or proceed with creating the backup specification at Step 6 on page 39.

Using the Data Protector CLI

Log in to the Informix Server system as user `informix`. From the directory:

Windows: `Data_Protector_home\bin`

HP-UX and Solaris: `/opt/omni/lbin`

Other UNIX systems: `/usr/omni/bin`

run:

Windows:

```
perl -I..\lib\perl util_informix.pl -CONFIG INFORMIXSERVER
INFORMIXDIR sqlhosts ONCONFIG
```

UNIX:

```
util_informix.pl -CONFIG INFORMIXSERVER INFORMIXDIR sqlhosts
ONCONFIG
```

Parameter description

INFORMIXSERVER	Name of the Informix instance.
INFORMIXDIR	Pathname of the Informix Server home directory.
sqlhosts	***Windows:*** Name of the system with the `sqlhosts` entry in the Windows Registry. Use the UNC notation, for example: `\\computer_name`. ***UNIX:*** Pathname of the `sqlhosts` file.
ONCONFIG	Name of the Informix instance `ONCONFIG` file.

The message `*RETVAL*0` indicates successful configuration.

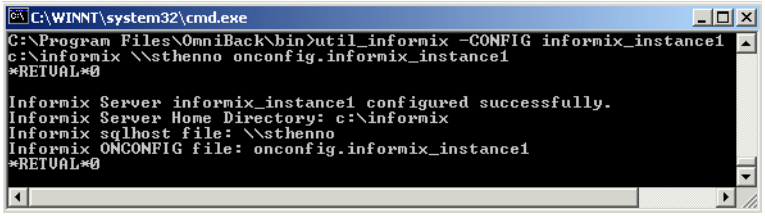

Figure 6 Configuring an Informix instance (Windows)

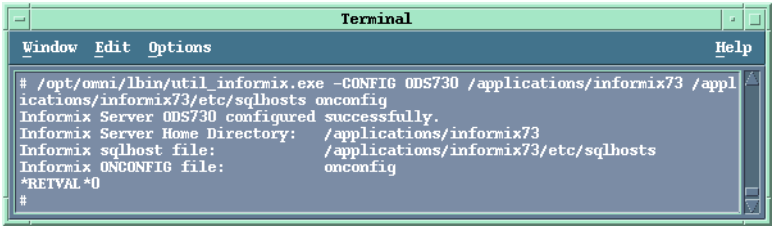

Figure 7 Configuring an Informix instance (HP-UX, Solaris)

Handling errors

If an error occurs, the error number is displayed in the form `*RETVAL*`*`error_number`*.

To get the error description:

Windows: On the Cell Manager, see the file *Data_Protector_home*\help\enu\Trouble.txt.

HP-UX and Solaris: Run:

```
/opt/omni/lbin/omnigetmsg 12 error_number
```

Other UNIX systems: Run:

```
/usr/omni/bin/omnigetmsg 12 error_number
```

Checking the configuration

You can check the configuration of an Informix instance after you have created at least one backup specification for the Informix instance. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

Using the Data Protector CLI

Log in to the Informix Server system as user `informix`. From the directory:

Windows: *Data_Protector_home*\bin

HP-UX and Solaris: /opt/omni/lbin

Other UNIX systems: /usr/omni/bin

run:

Windows:

```
perl -I..\lib\perl util_informix.pl -CHKCONF INFORMIXSERVER
```

UNIX:

```
util_informix.pl -CHKCONF INFORMIXSERVER
```

where *INFORMIXSERVER* is the name of the Informix instance.

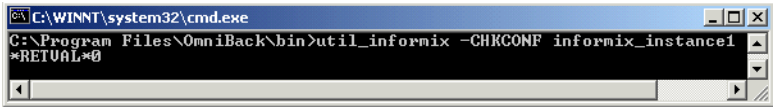

Figure 8 Checking configuration (Windows)

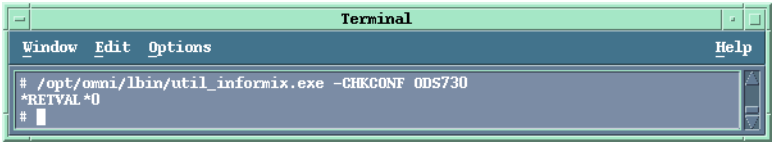

Figure 9 Checking configuration (UNIX)

A successful configuration check displays the message `*RETVAL*0`.

If an error occurs, the error number is displayed in the form `*RETVAL*`*`error_number`*. For information on how to get the error description, see "Handling errors" on page 35.

Backup

The integration provides online database backup of the following types:

Table 5 Informix Server backup types

Full	Full backup (level 0).
Incr1	Incremental backup (level 1). Backs up changes since the last Full backup.
Incr2	Incremental backup (level 2). Backs up changes since the last Incr1 backup.

For details on these types and on ON-Bar, see the *Backup and restore guide* of Informix Server.

What you must back up as filesystem

ON-Bar backs up all dbobjects *except* the following, which you *must* back up using a filesystem backup:

Table 6 What needs to be backed up as filesystem

Object	Location
The `ONCONFIG` file	***Windows:*** `INFORMIXDIR\etc` ***UNIX:*** `INFORMIXDIR/etc`
The `oncfg_SERVERNAME.SERVERNUM` file	
Emergency boot file, an Informix Server configuration file called `ixbar.server_id`, where `server_id` is the value of the `SERVERNUM` configuration parameter.	
UNIX only: The `sqlhosts` file	
Simple-large-object data in blobspaces	Disks or optical platters

IMPORTANT:

How often you need to back up these objects depends on how frequently they change. However, back up the emergency boot file at least daily and always after a critical dbspace backup.

What does not need to be backed up?

ON-Bar does not back up the following items because it automatically re-creates them during a restore:

- Dbspace pages allocated to the Informix instance but not yet allocated to a tblspace extent.
- Mirror chunks, if the corresponding primary chunks are accessible.
- Temporary dbspaces.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

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

   In **Application database**, select the Informix instance to be backed up.

   UNIX only: Type `informix` in both **Username** and **Group name**.

   Click **Next**.
5. If the Informix instance is not configured yet for use with Data Protector, the Configure Informix dialog box is displayed. Configure it as described in "Configuring Informix instances" on page 31.
6. Select the dbobjects to be backed up. For a whole-system backup, select all dbobjects.

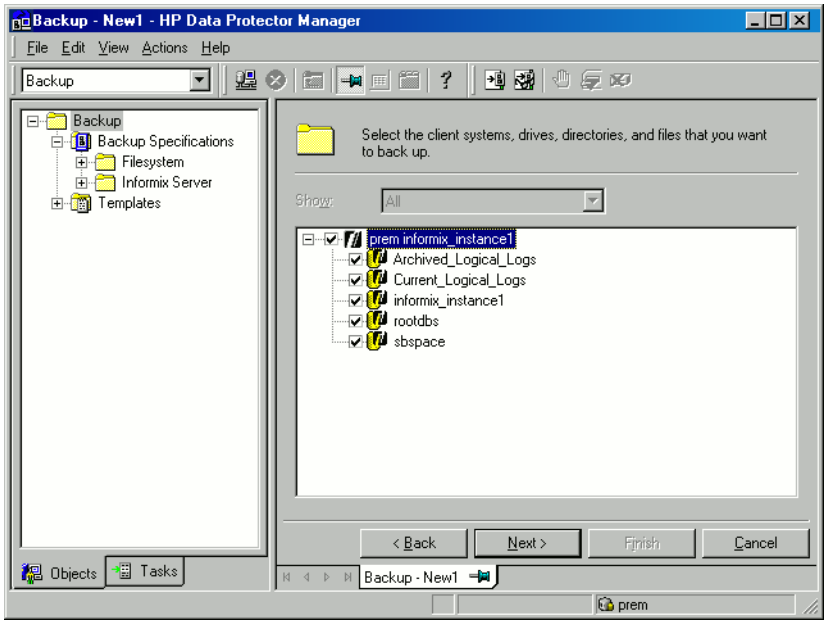

Figure 10 Selecting backup objects

Click **Next**.

7. Select devices to use for the backup.

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

NOTE:

Except for whole-system backups, ON-Bar backs up and restores dbobjects concurrently, creating a new process for each object. The number of processes is limited by the Informix Server `BAR_MAX_BACKUP` configuration parameter. Set the Informix configuration parameter `BAR_MAX_BACKUP` to the Data Protector concurrency.

To specify which resource types can be backed up to the device, click the **Informix** tab, select the desired resource types, and click **OK**. See Figure 11 on page 41.

Ensure that the selected devices cover all resource types specified for backup and are not locked when starting the backup. Ideally, back up each resource enter to a separate device.

IMPORTANT:

For a logical log backup, always use a separate device and ensure that the `LTAPEDEV` parameter in the `ONCONFIG` file is not set to `/dev/null` or `''`.

Figure 11 Specifying Informix Server resource types

Table 7 Informix Server resource types

B	Blobspace
CD	Critical dbspace (a root dbspace or a dbspace containing the physical log or a logical log file)
L	Logical log
MR	Master root dbspace
ND	Non-critical dbspace
R	Root dbspace

TIP:

Select an additional set of devices (covering all resource types specified for backup) so that they can take over if some devices in the primary group fail. Select the **Load balancing** option and set the Min and Max parameters to the number of primary devices.

Click **Next**.

8. Set backup options (Figure 12 on page 42 and Figure 13 on page 43). For information, see Table 8 on page 44.

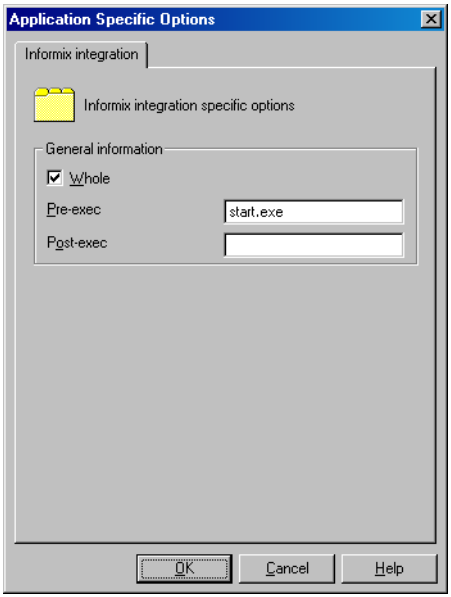

Figure 12 Informix Server specific backup options (Windows)

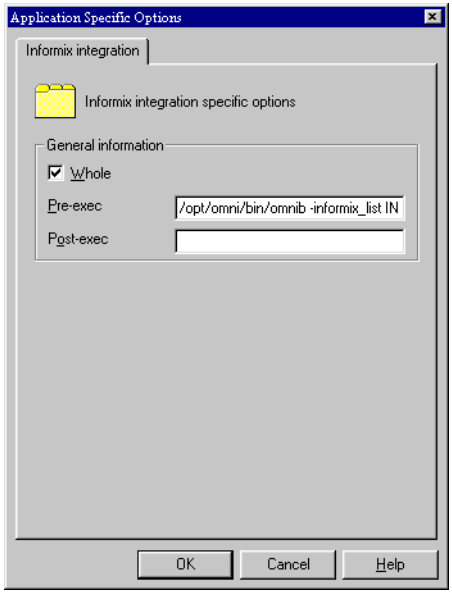

Figure 13 Informix Server specific backup options (UNIX)

Click **Next**.

9. Optionally, schedule the backup. See "Scheduling backup specifications" on page 45.

   Click **Next**.

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

TIP:

Preview backup session for your backup specification before using it. See "Previewing backup sessions" on page 46.

Table 8 Informix Server backup options

Option	Description
Whole	Select this option to perform a whole-system backup. This option is only available if you selected all dbobjects to be backed up at Step 6 on page 39. In a whole-system backup, all Informix instance's dbobjects from the `onbar` command are backed up. ON-Bar cannot back them up concurrently; they are backed up sequentially. Whole-system backup is useful for small systems, when you do not need to restore logs, for disaster recovery, or when restoring to another client.
Pre-exec **Post-exec**	Specify a command that will be started by `ob2onbar.pl` on the Informix Server system before the backup (`pre-exec`) or after it (`post-exec`). Do not use double quotes. ***Windows:*** Provide only the name of the command, which must reside in the *`Data_Protector_home`*`\bin` directory. See Figure 12 on page 42. ***UNIX:*** Provide the pathname of the command. See Figure 13 on page 43. If you selected a logical log for backup, it is sensible to add `onmode -l` as a `pre-exec` command to ensure that you always have a log file to back up. Without a log file to back up, the backup fails.

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes. See Figure 14 on page 45.

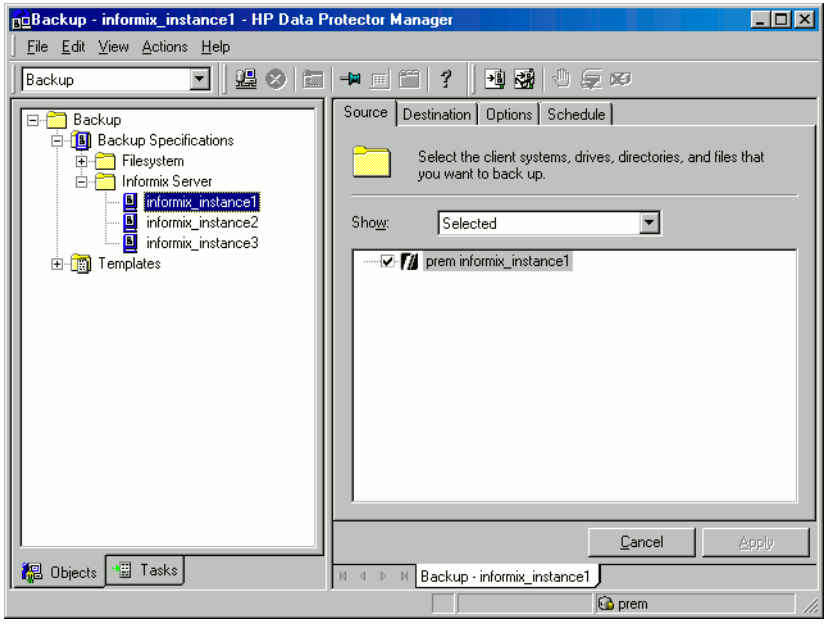

Figure 14 Modifying a backup specification

Scheduling backup specifications

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

Scheduling example

To back up logical logs at 8:00, 13:00, and 18:00 during weekdays:

1. In the Schedule property page, select the starting date in the calendar and click **Add** to open the Schedule Backup dialog box.
2. Under Recurring, select **Weekly**. Under Time options, select **8:00**. Under Recurring Options, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See Figure 15 on page 46.

   Click **OK**.
3. Repeat steps “1” on page 45 and “2” on page 45 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

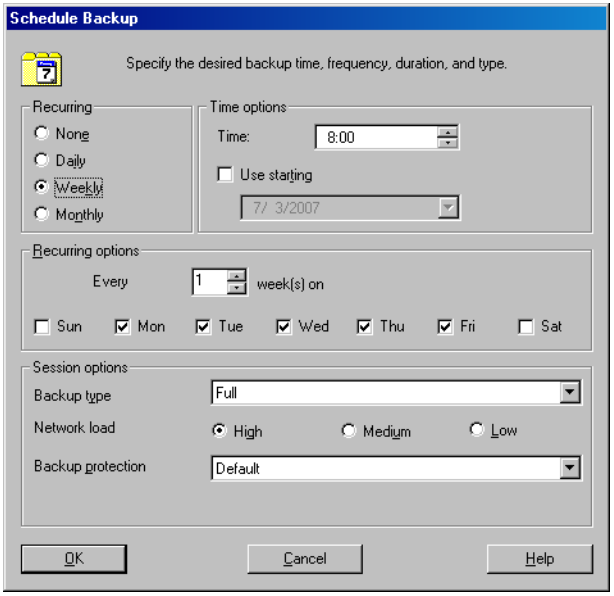

Figure 15 Scheduling a backup specification

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the directory:

Windows: *`Data_Protector_home`*`\bin`

HP-UX and Solaris: `/opt/omni/bin`

Other UNIX systems: `/usr/omni/bin`

run:

```
omnib -informix_list backup_specification_name -test_bar
```

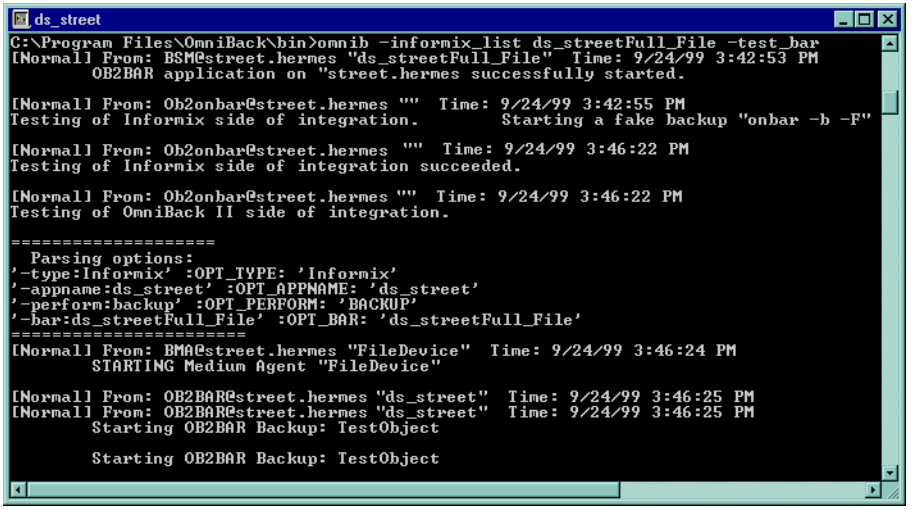

Figure 16 Previewing a backup with backup specification ds_street (Windows)

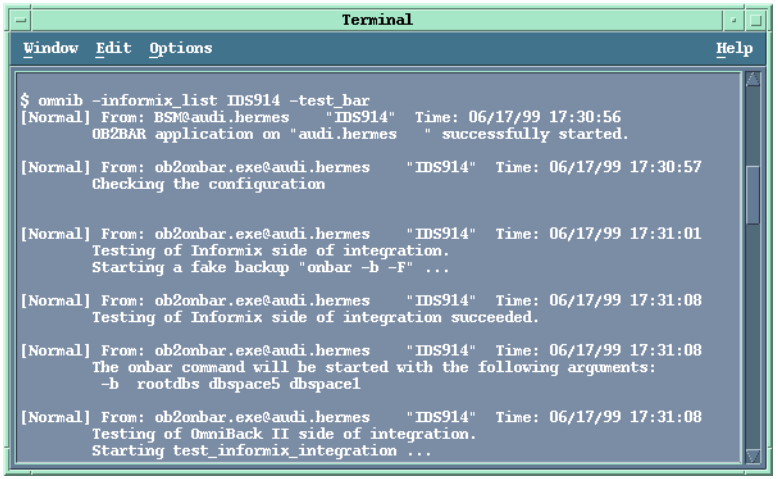

Figure 17 Previewing a backup with backup specification IDS914 (UNIX)

What happens during the preview?

1. The Informix Server `onbar` command is started with the `-F` option, which specifies a fake backup. This tests if the Informix instance is correctly configured for backup.

2. Data Protector tests the Data Protector part of the configuration. The following is tested:
   - Communication between the Informix instance and Data Protector
   - The syntax of the backup specification
   - If devices are correctly specified
   - If the necessary media are in the devices

Starting backup sessions

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

Backup methods

Start a backup of dbobjects in any of the following ways:

- Use the Data Protector GUI. See "Using the Data Protector GUI" on page 49.
- Use the Data Protector CLI. See "Using the Data Protector CLI" on page 50.
- Use the Informix Server `onbar` command. See "Using Informix Server commands" on page 51.
- ***UNIX only:*** Use the Informix Server `log_full.sh` script. See "Using Informix Server log_full.sh on UNIX" on page 52.

Before you begin

- Ensure that you have sufficient logical log space to create a backup.

  If the amount of free space in all logical log files is less than half a single log file, Informix Server does not create a backup.
- Before a Full backup, print or keep a copy of your `ONCONFIG` file, the emergency boot file, and on UNIX, also the `sqlhosts` file.
- Verify data consistency.

Using the Data Protector GUI

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

3. Select the **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

From the directory:

Windows: *`Data_Protector_home`*`\bin`

HP-UX and Solaris: `/opt/omni/bin`

Other UNIX systems: `/usr/omni/bin`

run:

```
omnib -informix_list backup_specification_name [-barmode
InformixMode] [List_options]
```

where *`InformixMode`* is one of the following:

```
full|inf_incr1|inf_incr2
```

NOTE:

Data Protector terms `full`, `inf_incr1`, and `inf_incr2` backup are equivalent to Informix Server terms level-0, level-1, and level-2 backup, respectively.

For *`List_options`*, see the `omnib` man page.

Examples

To start a full backup using the Informix Server backup specification `InformixWhole`, run:

```
omnib -informix_list InformixWhole -barmode full
```

To start an incremental backup (level 1) of the Informix Server backup specification `InformixIncr`, run:

```
omnib -informix_list InformixIncr -barmode inf_incr1
```

Using Informix Server commands

Use the Informix Server `onbar` command to start a backup of dbobjects from the Informix Server system where the relevant Informix instance is located.

Before the backup:

- Log in to the Informix Server system as user `informix`.
- Set the following variables:

Table 9 Data Protector and Informix Server variables

`ONCONFIG`	Name of the Informix instance `ONCONFIG` file.
`INFORMIXSQLHOSTS`	***Windows:*** System on which the `sqlhosts` entry in the Windows Registry exists. ***UNIX:*** Pathname of the `sqlhosts` file, for example `/applications/informix/etc/sqlhosts`.
`INFORMIXDIR`	Pathname of the Informix Server home directory.
`INFORMIXSERVER`	Name of the Informix instance.
`OB2APPNAME`	Name of the Informix instance.
`OB2BARLIST`	For *backup*, name of the backup specification to be used for the backup. For *restore*, name of the backup specification to be used for salvaging logical logs.

- Ensure that the Informix instance is in online or quiescent mode. Once you start a backup, do not change the mode until the backup finishes; changing the mode terminates the backup. Only online dbspaces and blobspaces are backed up. To list online dbobjects, run:

  Windows: *`INFORMIXDIR`*`\bin\onstat -d`

  UNIX: *`INFORMIXDIR`*`/bin/onstat -d`

Table 10 Backup modes

Online	Use online mode if your Informix instance must be accessible during the backup. An online backup may impact performance.

Quiescent	Use quiescent mode to eliminate partial transactions in a backup. Quiescent backup may not be practical if you need continuous access to Informix instances.

- Keep a copy of your `ONCONFIG` file, the emergency boot file, and on UNIX, also the `sqlhosts` file, after you create a full backup. You need this information to restore dbobjects.

To back up a list of dbspaces, run:

```
onbar -b dbspace_list
```

For example, to back up dbspaces `dbspace1` and `dbspace3`, run:

```
onbar -b dbspace1, dbspace3
```

To back up the current logical log file and switch to the next logical log file, run:

Informix Server 7.3: `onbar -l -c`

Informix Server 9.4: `onbar -b -l -c`

For more information, see the *Backup and restore guide* of Informix Server.

Using Informix Server log_full.sh on UNIX

On UNIX, `log_full.sh` is used to start a backup of logical log files when the Informix Server issues a log-full event alarm on the Informix Server. For information on logical log file backups, see "Manual and continuous logical log backups" on page 53.

To enable Informix Server backups from the `log_full.sh` script:

1. Add the following line to the Informix instance `ONCONFIG` file:

   ```
   ALARMPROGRAM INFORMIXDIR/etc/log_full.sh.
   ```

2. If the Data Protector User Interface is not installed on the Informix Server system, create an Informix Server backup specification to back up only logical logs, and edit *INFORMIXDIR*`/etc/log_full.sh`.

   Add the following at the beginning of the file:

   ```
   export OB2BARLIST=backup_specification_name
   ```

   ```
   export OB2APPNAME=INFORMIXSERVER
   ```

3. If the Data Protector User Interface is installed on the Informix Server system, create an Informix Server backup specification to back up logical logs only.

Manual and continuous logical log backups

To back up logical log files that are full and ready to be backed up, start:

- a manual logical log backup to back up all full logical log files and stop at the current logical log file.
- a continuous logical log backup to back up each logical log file automatically as it becomes full. Use this backup 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.

To perform a manual logical log backup, set the `OB2APPNAME` and `OB2BARLIST` environment variables as described in Table 9 on page 51 and run:

```
onbar -l
```

To back up the current logical log file and switch to the next logical log file, run:

Informix Server 7.3: `onbar -l -c`

Informix Server 9.4: `onbar -b -l -c`

For more information, see the *Backup and restore guide* of Informix Server.

Restore

The Data Protector Informix Server integration provides two types of restore:

Table 11 Informix Server restore types

Complete database restore	Restore from any backup. ON-Bar restores dbobjects concurrently and replays the logical logs once.

Whole-system restore	Restore from a whole-system backup. ON-Bar restores the whole system sequentially with or without restoring the logical logs. Whole-system restore is appropriate for small systems, when you do not need to restore logs, for disaster recovery, or when restoring to another client.

Restore methods

Restore dbobjects in any of the following ways:

- Use the Data Protector GUI. See "Restoring using the Data Protector GUI" on page 58.
- Use the Data Protector CLI. See "Restoring using the Data Protector CLI" on page 63.
- Use the Informix Server `onbar` command. See "Restoring using Informix Server commands" on page 64.

Before you begin

- Before restoring the root dbspace or performing a whole-system restore, shut down the Informix instance (cold restore). Log in to the Informix Server system as user `informix` and run:

  Windows: *INFORMIXDIR*`\bin\onmode -ky`

  UNIX: *INFORMIXDIR*`/bin/onmode -ky`

NOTE:

Once the Informix instance is offline, you cannot restore only non-critical (user) dbspaces. The root dbspace must also be selected for restore.

- To restore only non-critical dbspaces, ensure that the Informix instance is online or in a quiescent mode (warm restore), and that the non-critical dbspaces to be restored are offline.

  To check whether dbspaces are offline, run:

  Windows: *INFORMIXDIR*`\bin\onstat -d`

  UNIX: *INFORMIXDIR*`/bin/onstat -d`

Finding information for restore

To restore dbobjects, first find the needed media and the session ID of the last full backup session. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

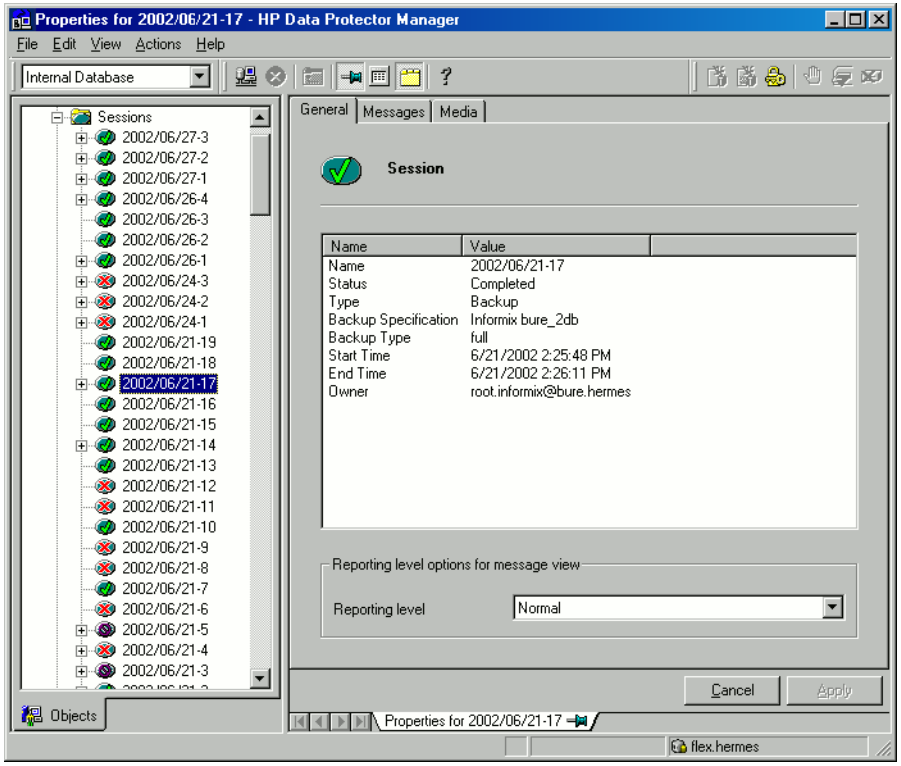

Figure 18 Example of session properties

Using the Data Protector CLI

Localized database names: If the names of backed up objects contain characters from different Unicode language groups (for example, if you are using Japanese and

latin characters), you must redirect the output of Data Protector utilities to use UTF-8 encoding:

- Set the environment variable `OB2_CLI_UTF8` to `1`.
- Set the encoding used on the terminal to UTF-8.

If you are using localized databases, and the system locale uses the same Unicode language group, no changes are required.

1. Go to the directory:

   Windows: *`Data_Protector_home`*`\bin`

   HP-UX and Solaris: `/opt/omni/bin`

   Other UNIX systems: `/usr/omni/bin`

2. Get a list of Informix Server backed up objects:

```
omnidb -informix
```

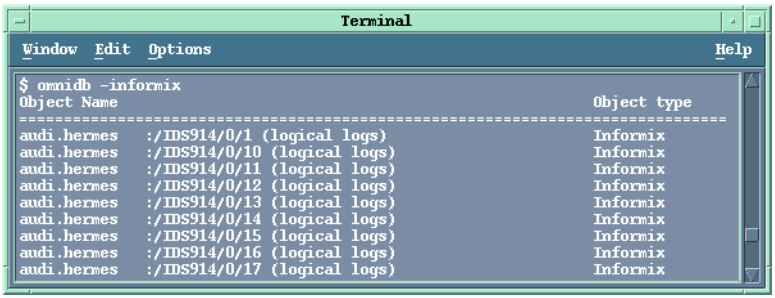

Figure 19 Example of a list of Informix Server backed up objects

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

```
omnidb -informix object_name
```

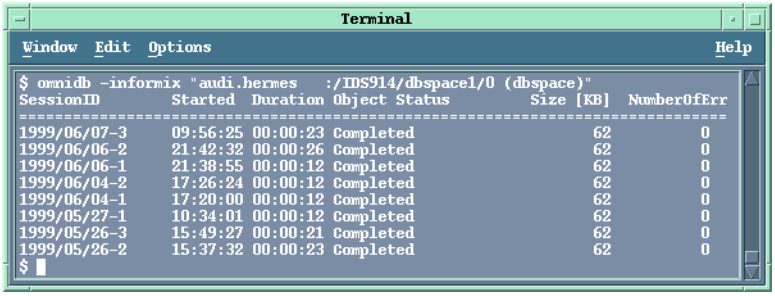

Figure 20 Example of a list of backup sessions for a specific object

IMPORTANT:

For object copies, use the object backup ID (which equals the object backup session ID). Do not use the object copy session ID.

To get information on the object backup ID, run:

```
omnidb -session session_id -detail
```

4. Get a list of media needed for restore:

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

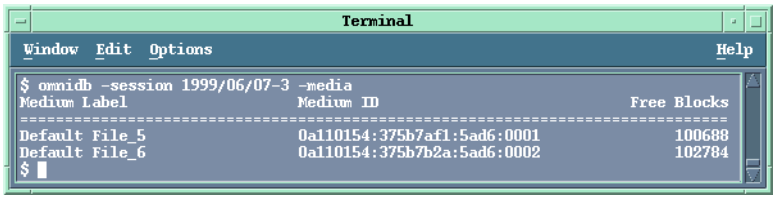

Figure 21 Example of finding media needed for restore

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

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Informix Server**, expand the client from which the data to be restored was backed up, and then click the Informix instance you want to restore.

3. In the **Source** page, select objects for restore. To restore the complete database or for a whole-system restore, select **Restore complete database**.

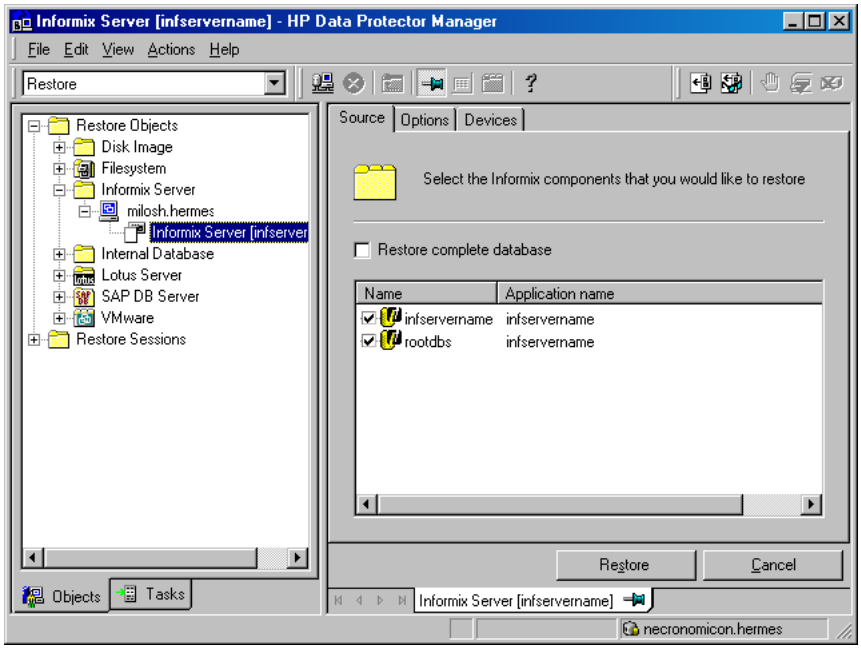

Figure 22 Selecting objects for restore

4. In the **Options** page, set the Informix Server specific restore options. For information, see Table 12 on page 62 or press **F1**.

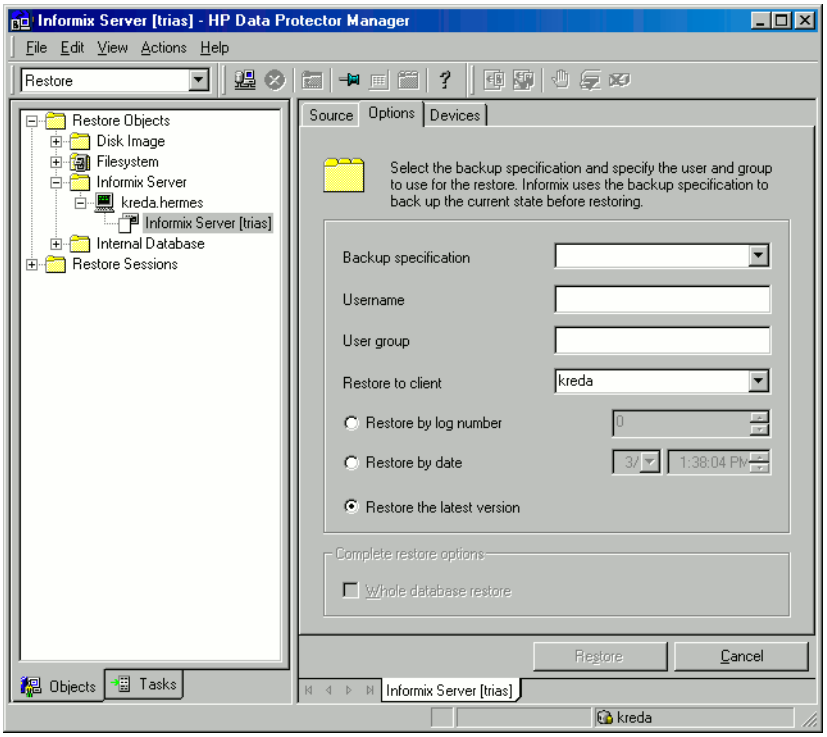

Figure 23 Informix Server restore options

5. In the **Devices** page, select devices to use for the restore.

   The **Automatic device selection** option is selected by default, but it is recommended to select the **Original devices selection** option.

IMPORTANT:

If you decide to use the **Automatic device selection** option, ensure that the number of available devices is equal or greater than the number of devices that were used for backup.

6. If you perform a whole-system restore and the Informix instance is in online mode, take the Informix instance offline by running:

   ```
   onmode -ky
   ```

   Click **Restore**.

7. In the Start Restore Session dialog box, click **Next** .

8. Specify the **Report level** and **Network load**.

   Click **Finish** to start the restore.

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

9. If you performed a whole-system restore, bring the Informix instance online by running:

```
onmode -m
```

Table 12 Informix Server restore options

Option	Description
Backup Specification	The backup specification to be used for salvaging logical log files still on the disk before restoring. Preferably, specify the backup specification used for the backup of logical logs.
Username	***UNIX only:*** User name of the Informix Server backup owner. `onbar` is started under the account of the specified user.
User group	***UNIX only:*** User group of the Informix Server backup owner.
Restore to client	The client to restore to. By default, you restore to the original backup client. This option is only valid for a whole-system restore.
Restore by log number	This option is only available if you selected **Restore complete database** in the Source page. Use this option to restore data up to a specific log number. If further logs exist, ON-Bar does not restore them. This option invokes `onbar -r -n last_log_number`. For details, see the *Backup and restore guide* of Informix Server.
Restore by date	This option is only available if you selected **Restore complete database** in the Source page. Use this option to restore data to a specific point in time. This option invokes `onbar -r -t time`. For details, see the *Backup and restore guide* of Informix Server.
Restore the latest version	Select this option to restore the latest backup version.
Whole database restore	This option is only available if you selected **Restore complete database** in the Source page. Select this option to perform a whole-system restore. Only use this option when restoring from a whole-system backup. Data Protector does not automatically detect if a whole-system backup exists. Data Protector searches for the last whole-system backup and restores from that. This option invokes `onbar -r -w`. For details, see the *Backup and restore guide* of Informix Server.

Restoring using the Data Protector CLI

From the directory:

Windows: *Data_Protector_home*\bin

HP-UX and Solaris: /opt/omni/bin

Other UNIX systems: /usr/omni/bin

run:

```
omnir -informix -barhost ClientName -barcmnd ob2onbar.pl -user
User:Group -appname INFORMIXSERVER -bararg
OnBarRestoreArguments [INFORMIX_OPTIONS]
```

ClientName	Name of the Informix Server system. In a cluster environment, name of the virtual server.
INFORMIXSERVER	Name of the Informix instance.
User, *Group*	***UNIX only:*** The user name and its group name.
OnBarRestoreArguments	ON-Bar restore arguments. Put each argument in double quotes.
INFORMIX_OPTIONS	A subset of general restore options. For information, see the omnir man page.

Example

To restore the Informix instance informix_instance1 on the UNIX system computer with the bar argument -r rootdbs, run:

```
omnir -informix -barhost computer -barcmnd ob2onbar.pl -user
informix:informix -appname informix_instance1 -bararg "-r
rootdbs"
```

Restoring using Informix Server commands

Before restoring:

- Log in to the Informix Server system as user `informix`.
- Set Data Protector and Informix Server variables as described in Table 9, page 51.
- If a disk failure occurs, salvage logical log files that are still on the disk by running:

  ```
  onbar -l -s
  ```

The following are examples of the `onbar` command syntax for restore. For further options, see the *Backup and restore guide* of Informix Server.

Restoring dbspaces, blobspaces, and logical logs

1. If the Informix instance to be restored is in online mode, take it offline:

   ```
   onmode -ky
   ```

2. Restore `dbspaces`, `blobspaces`, and appropriate logical logs:

   Complete database restore: `onbar -r`

   Whole-system restore: `onbar -r -w`

3. After the restore, bring the Informix instance online:

   ```
   onmode -m
   ```

Restoring dbspaces and blobspaces only

To restore `dbspaces` and `blobspaces` without the logical log, run:

```
onbar -r -p
```

Restoring a particular dbspace or blobspace

To restore a specific dbspace, for example `dbspace_1` , run:

```
onbar -r dbspace_1
```

Restoring to another Informix Server

To restore data to an Informix Server system other than that from which the backup was made:

1. Install the Data Protector Informix Integration software component on the client to which you want to restore (target client).
2. Create user `informix` on the target client.
3. Create an Informix instance with the same name and number as the original Informix instance by using the Informix Server `ON-monitor` utility on the target client.
4. Ensure that the Informix instance is online.
5. Configure the Informix instance as described in "Configuring Informix instances" on page 31.
6. Take the Informix instance offline.
7. Copy the original Informix Server configuration files (`ONCONFIG`, the emergency boot file, `oncfg_`*`SERVERNAME`*`.`*`SERVERNUM`*, and on UNIX, also the `sqlhosts` file) to the target client. Change the client name in the files to the target client name.
8. Start a whole-system restore of dbobjects as described in "Restoring using the Data Protector GUI" on page 58.

Restoring using another device

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

Using the Data Protector GUI

For information on how to specify another device for restore using the Data Protector GUI, see the online Help index: "restore, selecting devices for".

Using the Data Protector CLI or Informix Server commands

If you are restoring using the Data Protector CLI or Informix Server commands, specify the new device in the file:

Windows: *`Data_Protector_home`*`\Config\Server\cell\restoredev`

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

Use the format:

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

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

IMPORTANT:

Delete this file after use.

On Windows, use the Unicode format for the file.

Monitoring sessions

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

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

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

When ON-Bar encounters an error or a condition that warrants a warning, it writes a message to the Informix Server 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, see the Backup and Restore Guide of Informix Server.

To abort a backup or restore session successfully, set the ON-Bar `BAR_RETRY` configuration parameter to `0`. This parameter specifies how many times ON-Bar retries a backup or restore if the first attempt fails.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Informix Server integration. Start at “Problems” on page 71 and if you cannot find a solution there, perform general checks and verifications.

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

Before you begin

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

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

Checks and verifications

If your configuration, backup, or restore failed:

- On the Informix Server system, examine system errors reported in the `debug.log` and `informix.log` files, located in the directory:

  Windows: *`Data_Protector_home`*`\log`

  HP-UX and Solaris: `/var/opt/omni/log`

  Other UNIX systems: `/usr/omni/log`

- Make a test backup and restore of any filesystem on the problematic client. For information, see online Help.
- ***Windows only:*** Ensure that the Data Protector `Inet` service is running under the account `informix`.
- ***UNIX only:*** Verify that the `onbar_d` command has the switch ownership(s) bit set and that it is owned by the Informix Server user, for example, `informix:informix` or `root:informix`.

  Verify that this user is also the owner of the backup specification, or in the case of a restore failure, verify that this user is specified for the restore session, and that it is in the Data Protector `operator` or `admin` group.

  If this user is in the Data Protector `operator` group, ensure that the **See private objects** user right of this group is selected. For information, see the online help index: "user rights, changing".

  Now test if this user, for example user *`informix`*, has appropriate rights in Data Protector. Log in to the Informix Server system as user `informix`. From the directory:

  HP-UX and Solaris: `/opt/omni/bin/utilns`

  Other UNIX systems: `/usr/omni/bin/utilns`

  run:

  `testbar -type:informix -perform:checkuser -bar:` *`backup_specification_name`*

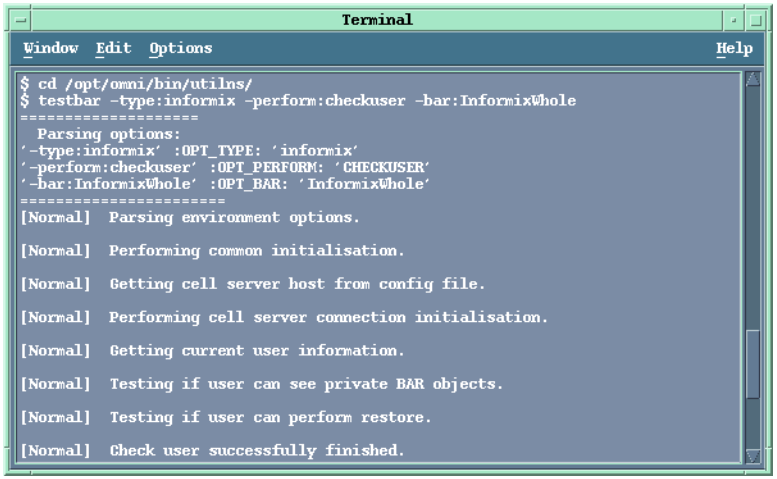

Figure 24 Example of checking the Informix Server user

In this example, the user has all the necessary rights for the backup specification named `InformixWhole`.

If the user `informix` on the Informix Server system `computer.hp.com` does not have the necessary rights, an error similar to the following will be displayed:

```
[Critical] From: OB2BAR@computer.hp.com "" Time: 08/06/2005
17:51:41[131:53]
User "informix.users@computer.hp.com" is not allowed
to perform a restore.
```

- In a cluster environment, ensure that the environment variable `OB2BARHOSTNAME` is set to the virtual server name before performing procedures from the Data Protector CLI. When the Data Protector GUI is used, this is not required.

Additionally, if your configuration or backup failed:

- Ensure that the Informix instance is online.

Additionally, if your backup failed:

- Check the configuration of the Informix instance as described in Checking the configuration, page 36.
- Test the backup specification as described in "Previewing backup sessions" on page 46.
  - If this fails, check if the Informix Server part of the test failed:

    Run the `onbar -b -F` command. If the test fails, see the Informix Server documentation for further instructions.

- If the Data Protector part of the test failed, create an Informix Server backup specification to back up to a null or file device.

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

- If the test succeeds, start the backup directly from the Informix Server system using Informix Server commands. For information, see "Using Informix Server commands" on page 51.

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

Additionally, if your backup or restore failed:

- Test the Data Protector data transfer using the `testbar` utility. Log in to the Informix Server system as user `informix`. From the directory:

  Windows: *`Data_Protector_home`*`\bin`

  HP-UX and Solaris: `/opt/omni/bin/utilns`

  Other UNIX systems: `/usr/omni/bin/utilns`

  - if your backup failed, run:

    `testbar -type:Informix -appname:`*`INFORMIXSERVER`* `-bar:`*`backup_specification_name`* `-perform:backup`

    where *`INFORMIXSERVER`* is the name of the Informix instance.

  - if your restore failed, run:

    `testbar -type:Informix -appname:`*`INFORMIXSERVER`*
    `-bar:`*`backup_specification_name`* `-perform:restore`
    `-object:`*`OBJECT_NAME`* `-version:`*`OBJECT_VERSION`*

    where *`INFORMIXSERVER`* is the name of the Informix instance, *`OBJECT_NAME`* is the name of the object to be restored, *`OBJECT_VERSION`* is the object version.

  If the test fails:

  1. Troubleshoot errors reported by the `testbar` utility using the Data Protector troubleshooting file, located on the Cell Manager in:

     Windows: *`Data_Protector_home`*`\help\enu\Trouble.txt`

     UNIX: `/opt/omni/gui/help/C/Trouble.txt`

2. On the Informix Server system, examine system errors reported in the file:

   Windows: *Data_Protector_home*\log\debug.log

   HP-UX and Solaris: /var/opt/omni/log/debug.log

   Other UNIX systems: /usr/omni/log/debug.log

Additionally, if your restore failed:

- Ensure that the backup specification used for salvaging logical logs is properly configured.

Checking the Informix Server side

The following checks may help you solve some Informix Server related problems.

If your backup or restore failed:

- Check the following Informix Server files for error descriptions:

  bar_act.log

  bar_dbg.log

  online.log

  Locations of these files are specified in the Informix Server ONCONFIG file.

Additionally, if your backup failed:

- Start a backup, not using Data Protector:
  1. Set the BAR_BSALIB_PATH shell variable to:

     Windows: *ISMDIR*\bin\libbsa.dll

     where *ISMDIR* is the path to the ISM.

     UNIX: *INFORMIXDIR*/lib/ibsad001.sl

     where *INFORMIXDIR* is the home directory of Informix Server.
  2. Use the onbar command to start the backup.

Additionally, if your restore failed:

- For a cold restore, verify if the dbspaces you want to restore are offline:
  1. Log in to the Informix Server system as user informix.

2. Run:

   Windows: *INFORMIXDIR*\bin\onstat -d

   UNIX: *INFORMIXDIR*/bin/onstat -d

   where *INFORMIXDIR* is the home directory of Informix Server.

- Ensure that the Informix Server configuration files (ONCONFIG, the emergency boot file, oncfg_*INFORMIXSERVER*.*SERVERNUM*, and on UNIX, also the sqlhosts file) are not corrupt. If they are corrupt, restore them manually.

Problems

Problem

Restore to another client fails

If you backed up data to one client, 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 Server emergency boot file (ixbar.*server_id*, where *server_id* is the value of the SERVERNUM 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 following format:

```
ODS730  rootdbs  R  1  7  0  9  1999008018  2005-08-18  18:10:25  1
```

Entries 7 and 9 make up make up the Data Protector session ID. Entry 9 is the date and entry 7 the unique session number.

Here, the session ID is 2005/08/18-9. Note that the delimiter in the date field is "-" in the emergency boot file and "/" in the Data Protector session ID.

The value of the SERVERNUM configuration parameter is given in entry 4.

Problem

Restore fails because the emergency boot file is too large

Action

Use the ON-Bar `onsmsync` utility to remove expired backups from the Informix Server `sysutils` database and emergency boot file. For information on the `onsmsync` utility, see the *Backup and restore guide* of Informix Server.

2 Integrating IBM DB2 UDB and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector IBM DB2 UDB (**DB2**) integration. It describes concepts and methods you need to understand to back up and restore DB2 databases.

Data Protector integrates with IBM DB2 Universal Database Server (**DB2 Server**) to back up DB2 database objects online and offline.

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

Table 13 Backup types

Full	Backs up complete DB2 objects.
Incremental	Backs up changes since the last Full backup.
Delta	Backs up changes since the last backup of any type.

The basic backup unit is a table space. Only table spaces or databases (DB2 objects) can be selected for backup.

When restoring a database or table space, you can specify restore options to perform:

- Rollforward recovery
- Version recovery
- Restore to a new database (database only)
- Restore to another instance (database only)
- Restore to another system (database only)
- Automatic restore from incremental or delta backups

Databases are restored offline, table spaces online.

Table or datafile backup and restore are not supported. Neither are backup or restore using Data Protector media with the DB2 Command Line Processor or the DB2 Control Center.

This chapter provides information specific to the Data Protector DB2 Server integration. For general Data Protector procedures and options, see online Help.

Integration concept

Data Protector integrates with the DB2 Server through a set of modules responsible for data backup and restore. Figure 25 on page 74 shows the architecture of the Data Protector DB2 integration.

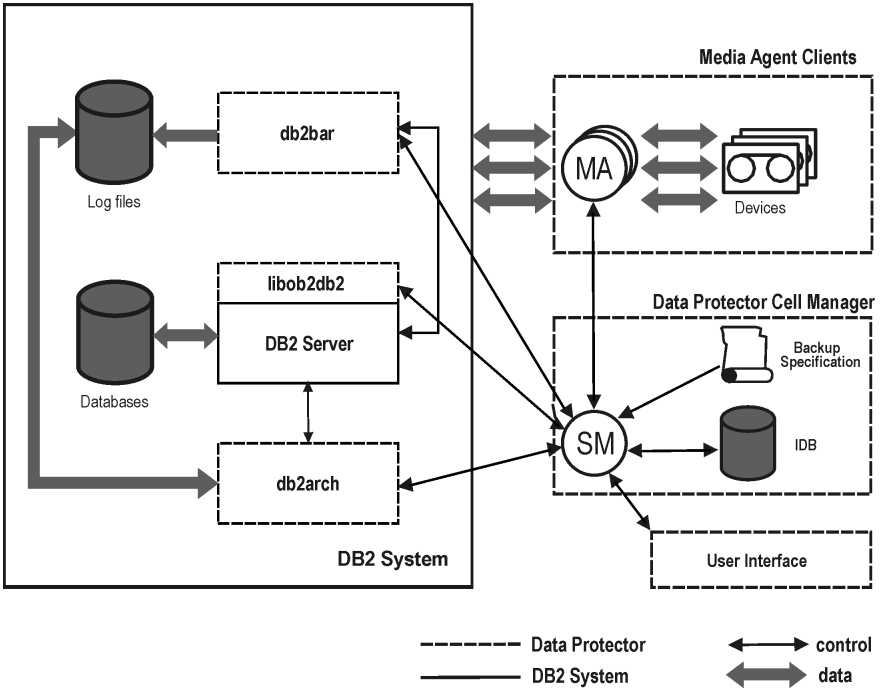

Figure 25 DB2 integration architecture

Table 14 Legend

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

db2bar	Data Protector module, used for controlling activities between the DB2 Server and Data Protector backup and restore.
db2arch	Program that backs up and restores DB2 log files.
libob2db2	Data transferring module, called by DB2 Server.
MA	Data Protector General Media Agent.
Backup Specification	A list of objects to be backed up, backup devices, and options to be used.
IDB	The Data Protector Internal Database.

While the DB2 Server is responsible for read/write operations to disk, Data Protector reads from and writes to devices and manages media.

Configuring the integration

You need to configure DB2 users and every DB2 instance you intend to back up or restore to.

Prerequisites

- Ensure you have correctly installed and configured DB2 Server.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on DB2 Server, see the *DB2 administration guide* and *DB2 server books online*.
- Ensure you have correctly installed Data Protector. For information on how to install the Data Protector IBM DB2 UDB integration in various architectures, see the *HP Data Protector installation and licensing guide*.

  Every DB2 Server system you intend to back up from or restore to must have the Data Protector `DB2 Integration` and `Disk Agent` components installed.

  In a partitioned environment, ensure that the `DB2 Integration` and `Disk Agent` components are installed on all the physical nodes on which the DB2 database resides.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the DB2 Server system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on the DB2 Server system.

Partitioned environment

In a physically partitioned environment, configure the integration on every physical node separately.

Ensure that the configuration parameter `MaxBSession` in the Data Protector global options file is set to at least twice the number of nodes of the partitioned database.

Configuring DB2 users

Ensure the DB2 user has appropriate authorities to perform DB2 backups and restores (either `SYSADM`, `SYSCTRL`, or `SYSMAINT`).

Add user `root` (UNIX only) and the DB2 user to both the Data Protector and DB2 `admin` user groups. For more information, see the online Help index: “user groups” and “adding users”.

Provide this user in configuration and restore procedures. This user is needed by Data Protector to start the Data Protector Inet service (Windows) or process (UNIX).

Configuring DB2 instances

Provide Data Protector with the DB2 instance configuration parameters:

- DB2 user
- DB2 user password
- DB2 instance home directory (only in a partitioned environment)

Data Protector then creates a DB2 instance configuration file on the Cell Manager and verifies the connection to the instance.

These parameters are used for connecting to the DB2 Server system to perform backups, restores, and other operations, such as listing objects for backup.

To configure a DB2 instance, use the Data Protector GUI or CLI.

Before you begin

- Ensure the DB2 instance is online.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **DB2 Integration**, and click **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. In Client, select the DB2 Server system.

   In a cluster environment, select the virtual server

   In **Application database**, type the DB2 instance name.

   UNIX only: Type a username and its group name.

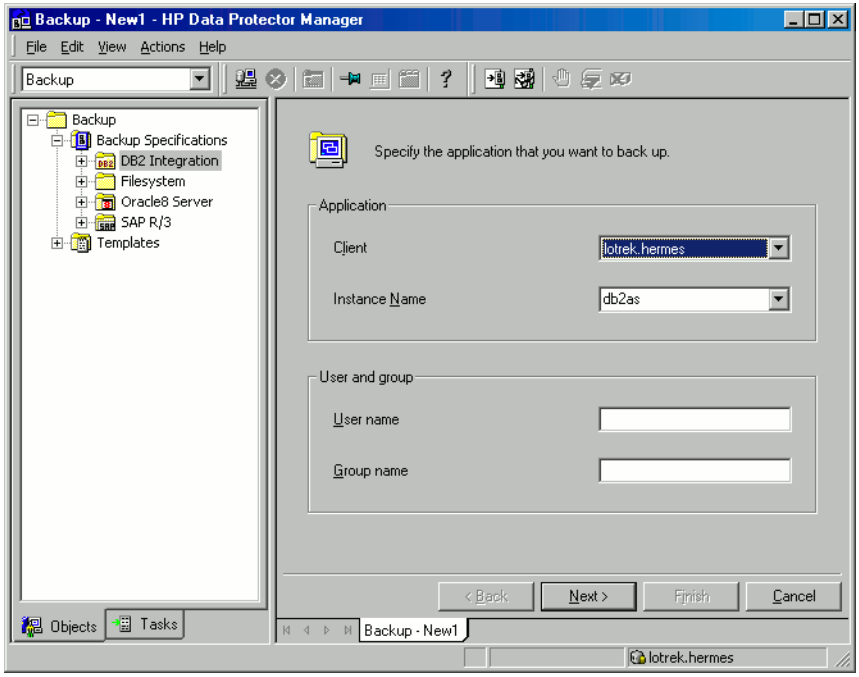

Figure 26 Specifying a DB2 instance

Click **OK**

5. Click **Next**. The Configure DB2 dialog box is displayed.
6. Type the name of the DB2 user and its password. This user must be configured as described in "Configuring DB2 users" on page 76.

   In a partitioned environment, select **DB2 EEE** and specify the pathname of the DB2 instance home directory.
7. The DB2 instance is configured. Exit the GUI or proceed with creating a backup specification at Step 6 on page 82.

Using the Data Protector CLI

From the directory:

Windows: `Data Protector\bin`

HP-UX: `/opt/omni/lbin`

Othe UNIX systems: `/usr/omni/bin`

run:

`util_db2 -CONFIG` *`DB2_instance username password`* `[`*`DB2_instance_home`*`]`

Parameter description

`DB2_instance`	Name of the DB2 instance.
`username`	DB2 user.
`password`	DB2 user password.
`DB2_instance_home`	Home directory (pathname) of the DB2 instance (only in a partitioned environment).

The message `*RETVAL*0` indicates successful configuration.

Checking the configuration

You can check the configuration of a DB2 instance after you have created at least one backup specification for the DB2 instance. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **DB2 Integration**. Click a backup specification for the DB2 instance.
3. In the Results Area, right-click the DB2 instance and click **Check configuration**.

Using the Data Protector CLI

From the directory:

Windows: `Data Protector\bin`

HP-UX: `/opt/omni/lbin`

Othe UNIX systems: `/usr/omni/bin`

run:

`util_db2.exe -CHKCONF` *`DB2_instance`*

Backup

The Data Protector DB2 integration provides three backup types and two backup modes.

Table 15 Backup types

Full	Backs up complete DB2 objects.
Incremental	Backs up changes since the last full backup.
Delta	Backs up changes since the last backup of any type.

Table 16 Backup modes

Online	Database is online.
Offline	Database is unavailable for use.

To configure a DB2 backup:

1. Create a backup specification for DB2 objects, using the `DB2 Database Backup` template.
2. To back up archived logs, create a backup specification for the archived logs, using the `Archived_Logs_Backup` template. Specify a different device than the one for backing up DB2 objects. Otherwise, archived logs cannot be backed up because the device is locked by the online backup session of DB2 objects.

IMPORTANT:

Archived logs are automatically backed up whenever a new offline redo log appears, for example, after the online backup of DB2 objects completes. Therefore, do not start an online backup of DB2 objects before creating an archived logs backup specification.

Delete any old archived logs backup specification before creating a new one.

Physically partitioned environment

In a physically partitioned environment, create one backup specification for DB2 database objects and one for archived logs for each physical node (system) on which the DB2 objects reside.

Ensure that the same DB2 database objects are selected for backup on all the physical nodes.

Since two devices are required to back up DB2 objects and archived logs from a single system, the total number of devices (drives) required is twice the number of physical nodes.

For information on how to run these backup specifications, see "Starting backups of physically partitioned DB2 objects" on page 88.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **DB2 Integration**, and click **Add Backup**.

3. Select a template and click **OK**.

Table 17 Backup templates

`DB2 Database Backup`	Used for backing up only DB2 database objects.
`Archived_Logs_Backup`	Used for backing up only archived logs. This type of backup specification can be saved, but not started or scheduled. It is used every time the User Exit program starts the backup of archived logs.

4. In **Client**, select the DB2 Server system; in a cluster environment, select the virtual server.

   In **Application database**, select the DB2 instance to be backed up and click **Next**.

   UNIX only: Type a username and its group name. This user will be the backup owner.

5. If the DB2 Instance is not configured for use with Data Protector, the Configure DB2 dialog box is displayed. Configure it as described in "Configuring DB2 instances" on page 76.

6. Select the DB2 objects you want to back up and click **Next**. The basic backup unit is a table space. Only table spaces and databases can be selected for backup. See Figure 27 on page 82.

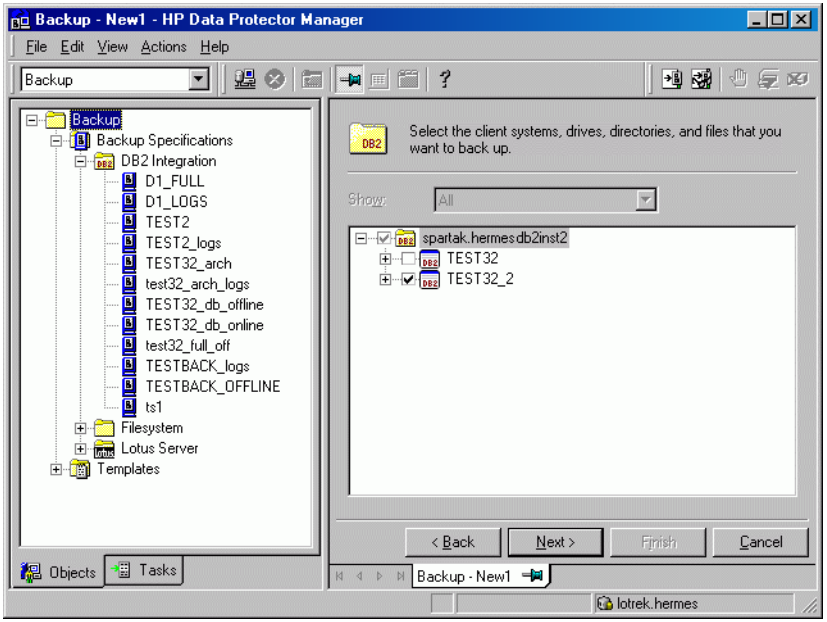

Figure 27 Selecting DB2 objects

If you select only DB2 temporary table spaces, the backup fails. To back up DB2 temporary table spaces, select the whole database.

IMPORTANT:

In a physically partitioned environment, select only one database or table spaces of the same database.

Click **Next**.

7. Select devices to use for the backup.

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

8. Set backup options and click **Next**.

   For information on application specific options, see Table 18 on page 84.

9. Optionally, schedule the backup and click **Next**. For more information, see "Scheduling backup specifications" on page 45.

10. To perform an offline backup of a particular DB2 object, right-click the object and click **Properties**. In the Object Properties dialog box, select **Offline Backup** and click **OK**. See Figure 28 on page 83.

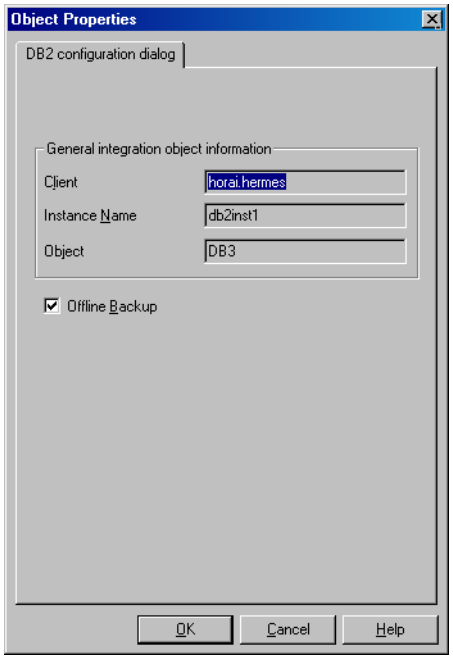

Figure 28 Selecting offline backup

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

TIP:

Use consistent names for the backup specifications of a physically partitioned DB2 object. For example, `MyObject1`, `MyObject2` and so on.

TIP:

Preview backup session for your backup specification before using it. See "Previewing backup sessions" on page 46.

Table 18 DB2 backup options

Pre-exec **Post-exec**	Specify a command to be started by `db2bar` on the DB2 Server system before the backup of every selected DB2 object (`pre-exec`) or after it (`post-exec`). Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in: ***Windows:*** *`Data Protector`*`\bin` ***HP-UX:*** `/opt/omni/lbin` ***Othe UNIX systems:*** `/usr/omni/bin`
Parallelism	Specify the number of data streams for backing up a database from a node. In a partitioned environment, `Parallelism` must equal the device concurrency. Default: 1.

Modifying backup specifications

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

Scheduling backup specifications

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

Example

To back up table spaces at 8:00, 13:00, and 18:00 during weekdays:

1. In the Schedule property page, select the starting date in the calendar and click **Add** to open the Schedule Backup dialog box.
2. Under Recurring, select **Weekly**. Under Time options, select **8:00**. Under Recurring Options, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See Figure 29 on page 85. Click **OK**.
3. Repeat Step 1 on page 84 and Step 2 on page 84 to schedule another backup at 13:00, and another one at 18:00.
4. Click **Apply** to save the changes.

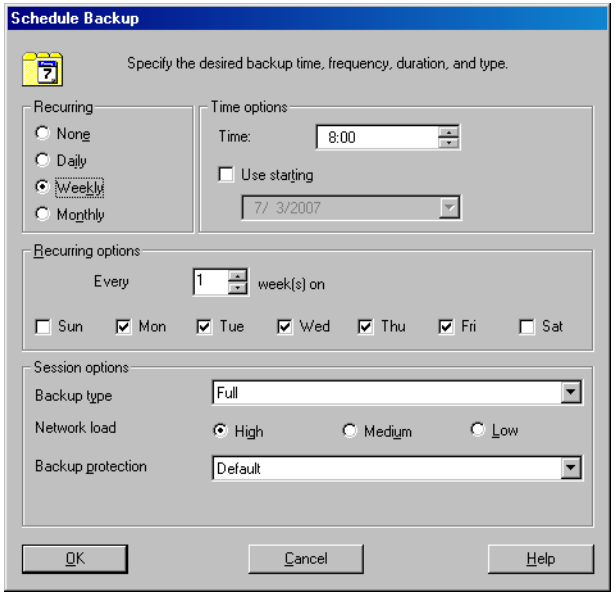

Figure 29 Scheduling a backup specification

Previewing backup sessions

Preview the backup session to test it. Use the Data Protector GUI or CLI.

The preview creates a file *backup_specification_name_TEST_FILE* in the *Data Protector*\tmp directory on the DB2 Server system. Delete it after the test.

Using the Data Protector GUI

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

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

Using the Data Protector CLI

From the directory:

Windows: `Data Protector\bin`

HP-UX and Solaris: `/opt/omni/bin`

Othe UNIX systems: `/usr/omni/bin`

run:

```
omnib -db2_list backup_specification_name -test_bar
```

What happens during the preview?

The `db2bar` command is started, which starts the Data Protector `testbar2` command to test:

- Communication within the Data Protector cell
- The syntax of the backup specification
- If devices are correctly specified
- If necessary media are in the devices

Then, the DB2 instance is checked for the presence of selected DB2 objects and whether they are in an appropriate state for backup.

Starting backup sessions

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

You can start a backup of DB2 objects using the Data Protector GUI or CLI.

Before you begin

- To enable online backups of DB2 objects, set the DB2 *`logretain`* and *`userexit`* parameters to ON (in a partitioned environment, on every node on which the object resides). Then restart the database for the new parameters to take effect and perform a full offline database backup.
- To enable incremental or delta backups of DB2 objects, set the DB2 *`trackmod`* parameter to ON:
  1. Run:

     ```
     db2 update db cfg for db_name USING TRACKMOD ON
     ```

     In a partitioned environment, run the command on every node on which the DB2 object resides.

2. Restart the database.
3. Perform a full offline database backup to non-Data Protector media by running:

   `backup db db_name`

- To enable offline backups of one or several DB2 table spaces (not the whole database), set the DB2 `logretain` parameter to ON.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **DB2 Integration**. Right-click the backup specification you want to start and click **Start Backup.**
3. Select the **Backup type** and **Network load**. Click **OK**.

   Successful backup displays the message `Session completed successfully`, providing the backup size, which is the size of full and incremental/delta backups together.

Using the Data Protector CLI

From the directory:

Windows: `Data Protector\bin`

HP-UX and Solaris: `/opt/omni/bin`

Othe UNIX systems: `/usr/omni/bin`

run:

`omnib -db2_list backup_specification_name [-barmode db2_mode] [options] [-preview]`

Parameter description

db2_mode	Backup type: {`-full` | `-incr` | `-delta`}
options	For information, see the `omnib` man page.

Example

To perform a full DB2 backup, using the backup specification `MyObjects`, and to set data protection to 10 weeks, run:

```
omnib -db2_list MyObjects -barmode -full -protect weeks 10
```

Starting backups of physically partitioned DB2 objects

1. Run the backup specification for the part of DB2 objects residing on the system with the catalog node. Use the Data Protector GUI or CLI.
2. Run the backup specifications for the other parts of the DB2 objects in any order.

   The order in which you run the backup specifications is only important if the object resides on the catalog node.

TIP:

To the first backup specification, add a post-exec script that will automatically run the other backup specifications. For more information, see the online Help index: “pre- and post-exec commands for backup specifications”.

Restore

Restore DB2 objects using the Data Protector GUI or CLI.

IMPORTANT:

Databases are restored offline.

Table spaces are restored online. Only table spaces that are not being restored are available for use.

A dropped table space can only be restored from a full database backup.

For information on how to restore a DB2 database to a new database, see “Restoring to a new database or another DB2 instance” on page 96.

For information on how to restore partitioned DB2 objects, see “Restore in a partitioned environment” on page 100.

Restoring using the Data Protector GUI

1. In the Context List, select **Restore**.

2. In the Scoping Pane, expand **DB2 Integration**, expand the client from which the data to be restored was backed up, and then click the DB2 instance you want to restore.

3. In the Source page, specify whether you want to restore database/tablespaces or archived logs and then browse for and select desired DB2 objects. See Figure 30 on page 90.

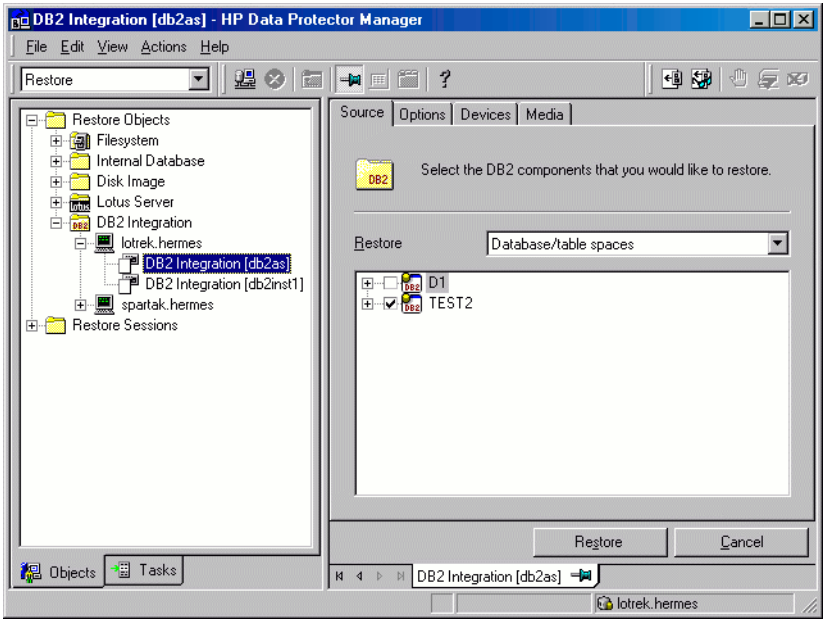

Figure 30 Selecting objects for restore

IMPORTANT:

In a physically partitioned environment, select only one database or several table spaces of the same database.

By default, the latest backup version is restored. To restore a DB2 object from a specific backup version, right-click the object, click **Properties** , and specify the backup version in the Properties for *DB2_object* dialog box. See Figure 31 on page 91.

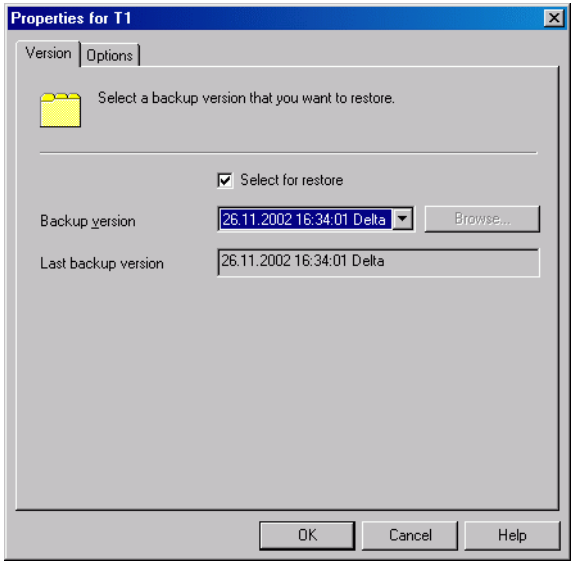

Figure 31 Selecting a version

To restore a database to a new database, right-click the database, click **Properties**, and then click the **Options** tab. Select **Restore to a new database** and specify a name for the new database. See Figure 32 on page 92.

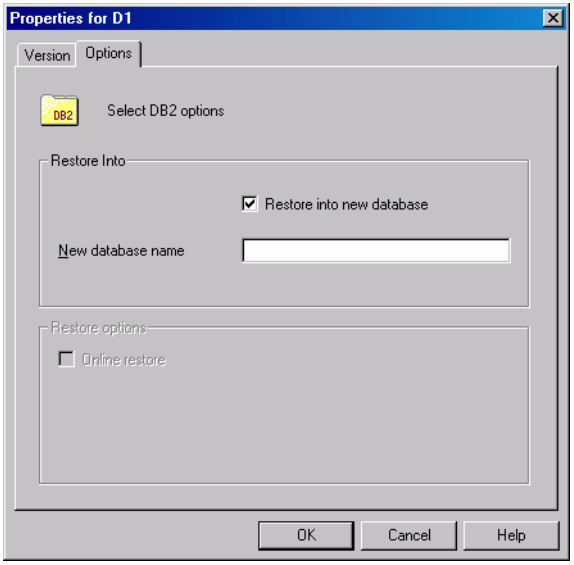

Figure 32 Restoring to a new database

4. In the Options page, set the DB2 restore options. For information, see Table 19 on page 94 or press **F1**.

NOTE:

For rollforward recovery, the latest backup version of log files is used. To perform a rollforward recovery using an older version of log files, first restore the desired log files and then restore the databases/tablespaces with the Rollforward option cleared. In a partitioned environment, connect to the catalog node. Finally, perform a rollforward recovery using DB2 tools.

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

   The **Automatic device selection** option is selected by default, but it is recommended to select the **Original device selection** option.

IMPORTANT:

If you decide to select the **Automatic device selection** option, ensure that the number of available devices is equal to or greater than the number of devices that were used for backup.

Click **Restore**.

6. In the Start Restore Session dialog box, click **Next**.
7. Specify the **Report level** and **Network load**.

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

Table 19 DB2 restore options

Restore to client	The client to restore to. By default, DB2 objects are restored to the source client. This option is only valid when restoring the whole database.
Username **User group** **Password**	DB2 user of the target DB2 instance, its group, and password.
Restore to instance	The DB2 instance to restore to. By default, DB2 objects are restored to the source DB2 instance. The instance must be configured for use with Data Protector as described in "Configuring Informix instances" on page 31. For details, see "Restoring to a new database or another DB2 instance" on page 96.
Rollforward	Select this option to perform a rollforward recovery. The database/tablespace is restored to its state at a specific time. During a rollforward recovery, both databases/tablespaces and archived logs are restored, and then the changes recorded in the archived logs are applied to the database/tablespace. The latest backup version of log files is used for this purpose. Specify the rollforward recovery by selecting `Rollforward to the end of the logs` or `Rollforward to date`. When specifying `Rollforward to date`, use local time (as set on the DB2 target server), not coordinated universal time (UTC). Rollforward recovery of the system catalog can only be performed to the end of the logs. You cannot restore other table spaces of the same database from the same session simultaneously. To perform a rollforward recovery in a physically partitioned environment, restore all the parts with `Rollforward` cleared (see "Restore in a partitioned environment" on page 100), connect to the catalog node, and then start a rollforward recovery using the DB2 Command Line Processor. To perform a version recovery, clear this option. The database/tablespace is restored to its state at the time of the backup. For a version recovery, you need a full offline database backup. When restoring from an online backup with `Rollforward` cleared, the database enters the rollforward pending state and becomes unavailable for use. To make it

	available, start a rollforward recovery using the DB2 Command Line Processor or Command Center (in a partitioned environment, the rollforward recovery must be started from the catalog node).

Restoring using the Data Protector CLI

From the directory:

Windows: *Data Protector*\bin

HP-UX and Solaris: /opt/omni/bin

Othe UNIX systems: /usr/omni/bin

run:

```
omnir -db2 -barhost source_client [-destination target_client]
-instance target_instance -dbname source_db [-session
session_ID] [-newdbname new_db] [-frominstance
source_instance] -tsname table_space [-session session_ID]
-logfile log_file [-session session_ID] [-rollforward [-time
YYYY-MM-DD.hh.mm.ss]]
```

Parameter description

source_client	The DB2 Server system from which DB2 objects were backed up. In a cluster environment, the name of the virtual server.
target_client	The target DB2 Server system (only if you are not restoring to the source client).
source_instance	The DB2 instance whose DB2 objects were backed up.
target_instance	The target DB2 instance.
source_db	The database you want to restore.
new_db	The target database (specify only if not the source database).

table_space	The table space you want to restore.
log_file	The log file you want to restore.
session_ID	Backup version ID. For object copies, use the object backup ID (which equals the object backup session ID). Do not use the object copy session ID.

For more information, see the omnir man page.

Example

To restore the DB2 database TEMP from the instance DB2Inst on the DB2 Server system degas and to roll it forward until the 10th of January 2006, 9:15 a.m., run:

```
omnir -db2 -barhost degas -instance DB2Inst -dbname TEMP
-rollforward time: 2006-01-10.09.15.00
```

Restoring to a new database or another DB2 instance

To restore a database to a new database in the source DB2 instance or another instance:

1. Find the containers of the source database:
   - To list table spaces of a particular database that reside on a particular node, connect to that node, then connect to the database, and run:

     ```
     db2 list tablespaces
     ```
   - To list the containers for a particular table space, run:

     ```
     db2 list tablespace containers for table_space_number
     ```

NOTE:

When restoring a corrupt database, you cannot use these commands because the data is lost.

2. Define new table space containers for the non-system table spaces by adding options for redirection to the DB2 configuration file. From the directory:

   Windows: `Data Protector\bin`

   HP-UX: `/opt/omni/lbin`

   Other UNIX systems: `/usr/omni/bin`

   run:

   ```
   util_cmd -putopt DB2 target_instance "old_container"
   "new_container" -sublist Redirection/source_db
   ```

   for every pair of table space containers.

   ## Parameter description

target_instance	The target instance.
source_db	The backed up database. The DB2 user of the target instance must have read and write permissions for the new containers.

3. In a physically partitioned environment, repeat Step 1 on page 96 and Step 2 on page 97 on every system.
4. Restore the source database to the new database without specifying rollforward recovery. Use the Data Protector GUI or CLI.

   In a physically partitioned environment, first restore the part of the database that resided on the system with the catalog node and then restore the other parts in any order.

   After the restore, the new database enters the rollforward pending state.

5. If you have restored from an offline backup, perform a rollforward recovery using DB2 tools:

   - In a non-partitioned environment, run:

     ```
     db2 rollforward db db_name stop
     ```

   - In a partitioned environment, run:

     ```
     db2 terminate
     export DB2NODE=catalog_node_number
     db2 rollforward db db_name stop
     ```

   If you have restored from an online backup, restore the archived logs, using the Data Protector GUI, and then perform a rollforward recovery, using DB2 tools:

   a. Log in to the source instance.

   b. Ensure that you have permissions to write to the archived logs directory and restore the archived logs, using the Data Protector GUI.

   The archived logs are restored to the same directory from which they were backed up.

   c. Copy the archived and redo logs of the source database to the corresponding log path directories of the new database (in a partitioned environment, to every node of the target instance).

   If the `SQLLPATH.TAG` file exists in the target log file directory, delete it to avoid possible database inconsistencies.

   d. If you are restoring to another instance, grant the ownership of the copied logs to the DB2 user of the target instance and log in to the target instance.

   e. Perform a rollforward recovery using DB2 tools:

   - In a non-partitioned environment, run:

     ```
     db2 rollforward db db_name [to time | to end of logs]
     [and complete]
     ```

   - In a partitioned environment, run:

     ```
     db2 terminate
     export DB2NODE=catalog_node_number
     db2 rollforward db db_name [to time | to end of logs]
     [and complete]
     ```

The following examples are from a non-partitioned environment.

Example 1

To restore the `db2db_old` database to the `db2db_new` database from an online backup (both databases reside in the `db2inst` instance, the log files of `db2db_old` are located in the `/db2_db/db2inst/NODE0000/SQL00003/SQLLOGDIR` directory and `"/tmp/db2cont1"` is the container for one of the table spaces:

1. Define a new container, `"/tmp/db2cont2"`, for the table space, using the Data Protector CLI:

   ```
   util_cmd -putopt DB2 db2inst "/tmp/db2cont1" \
   "tmp/db2cont2" -sublist Redirection/db2db_old
   ```

2. Restore the `db2db_old` database to the `db2db_new` database, using the Data Protector CLI:

   ```
   omnir -db2 -barhost source_client -instance db2inst -dbname
   db2db_old -newdbname db2db_new
   ```

3. Restore all archived logs needed for rollforward recovery using the Data Protector GUI.
4. Copy the archived and redo logs of the source database to the corresponding log path directories of the new database.
5. Perform a rollforward recovery to the end of logs, using the DB2 CLI:

   ```
   db2 rollforward db db2db_new to end of logs
   ```

Example 2

To restore the `db2db` database, from the instance `inst1`, to the `db2db` database in the `inst2` instance:

1. Define a new container `/tmp/db2cont2` for the table space, using the Data Protector CLI:

   ```
   util_cmd -putopt DB2 inst2 "/tmp/db2cont1" "/tmp/db2cont2"
   -sublist Redirection/db2db
   ```

2. Restore the db2db database to the inst2 instance, using the Data Protector CLI:

   ```
   omnir -db2 -barhost source_client [-destination
   target_client] -instance inst2 -dbname db2db -frominstance
   inst1
   ```

NOTE:

When restoring to another instance on another system, use the `db2 list tables for all` command to list tables.

Restore in a partitioned environment

You can restore a partitioned DB2 object to the original database or to a new database (on another DB2 instance).

Limitations

- You can restore an object from a non-partitioned environment to a partitioned environment (or the reverse), only if the partitioned environment has only one node (single partition).
- In a physically partitioned environment, automatic recovery is not possible.

Restoring to the original database

Corrupt database

To restore a corrupt database:

1. Connect to the node that was the catalog node of the corrupt database.
2. Create a new database with the same name.
3. Continue with the restore as described in "Restoring to a new database or another DB2 instance" on page 96.

Physically partitioned environment

To restore a physically partitioned DB2 object (residing on more than one system):

1. Restore the part of the DB2 object that resided on the system with the catalog node, without specifying rollforward recovery. Use the Data Protector GUI or CLI.
2. Restore all other parts of the DB2 object to the corresponding systems in any order, without specifying rollforward recovery.

3. Connect to the catalog node and perform a rollforward recovery, using DB2 tools:

```
db2 terminate

export DB2NODE=catalog_node_number

db2 rollforward db db_name [[stop]|[to time|to end of logs]
[and complete]]
```

NOTE:

The order in which you restore the parts of a DB2 object is only important if the object resides on the catalog node.

Logically partitioned environment

To restore a logically partitioned DB2 object (residing on only one system):

- For a version recovery:
  1. Restore the object, without specifying rollforward recovery. Use the Data Protector GUI or CLI.
  2. Connect to the catalog node and perform a rollforward:

     ```
     db2 terminate

     export DB2NODE=catalog_node_number

     db2 rollforward db db_name stop
     ```

- For a rollforward recovery, restore the object, specifying rollforward. Use the Data Protector GUI or CLI.

Restoring to a new database or another instance

To restore a database to a new database in the original DB2 instance, see "Restoring to a new database or another DB2 instance" on page 96.

To restore a database to a new database in another DB2 instance:

1. Log in to the target instance.
2. Ensure that the instance has the same node structure (number of nodes, node groups) as the source instance.

3. Connect to the node with the same node number as the catalog node of the source database:

   ```
   EXPORT DB2NODE=catalog_node_of_the_source_database
   ```

4. Create a database with the same name as the source database:

   ```
   db2 create db source_db
   ```

5. Continue with the restore as described in "Restoring to a new database or another DB2 instance" on page 96.

Monitoring sessions

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

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

For information on how to monitor a session, see the online Help index: "viewing currently running sessions".

NOTE:

All DB2 timestamps in messages displayed during rollforward recovery are by DB2 design in Universal Coordinated Time (UCT) format.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector DB2 integration. Start at "Problems" on page 71 and if you cannot find a solution there, perform general checks and verifications.

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

Before you begin

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

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

Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors reported in the `debug.log` and `db2.log` files, located in the directory:

  Windows: `Data Protector\log`

  HP-UX and Solaris: `/var/opt/omni/log`

  Other UNIX systems: `/usr/omni/log`

Additionally, if your backup failed:

- Test the backup specification as described in "Previewing backup sessions" on page 46.
  - If the DB2 part of the preview fails, see the DB2 documentation.
  - If the Data Protector part of the preview fails, create a DB2 backup specification to back up to a null or file device. Successful backup implies that the problem is related to devices. For information on troubleshooting devices, see online Help.

Additionally, if your backup or restore failed:

- Try performing:
  - A Data Protector filesystem backup and restore. For information, see online Help. After troubleshooting the filesystem backup, restart the DB2 Server and start a backup of DB2 objects again.
  - A backup and restore of DB2 objects using DB2 tools.

Additionally, if your restore failed:

- Ensure the target DB2 instance is online and configured for use with Data Protector.

Problems

Problem

Online backup is not allowed

DB2 reports:

```
Online backup is not allowed because either logretain or
userexit for roll-forward is not activated, or a backup
pending condition is in effect for the database.
```

Action

After configuring the DB2 database for rollforward recovery (`userexit` and `logretain ON`), first back up the database offline. If online backup is started first, the above error is reported.

Problem

Offline backup of one or several tablespaces is not allowed

When backing up DB2 tablespaces (not the whole database) offline, DB2 reports that offline backup is not allowed because the DB2 `logretain` option is not activated or that a backup pending condition is in effect for the database.

Action

Set the DB2 `logretain` option to `ON`.

Problem

Archived logs are not backed up

If you have created several archived logs backup specifications and deleted the one created last, the remaining backup specifications are not used and archived logs are not backed up.

Action

Create a new archived logs backup specification.

Problem

Incremental backup is not enabled for the database

If you start an incremental backup before a full backup has been performed, Data Protector reports:

```
Incremental backup is not enabled for this database.
```

Action

1. Activate modification tracking by running:

   ```
   db2 update db cfg for database_name USING TRACKMOD ON
   ```

2. Restart the database.
3. Perform a full database backup.

Problem

Error occurred while accessing an object

DB2 reports:

```
SQL2048N An error occurred while accessing object object.
Reason code: code_number
```

The following can be a reason (code number):

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 not in the appropriate state for the operation or some containers of the table space are not available. (`LIST TABLESPACES` lists the current table space state.)
7. A delete object operation failed.
8. Tried 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. Consider using the `QUIESCE` command to bring the database to a quiesced state to ensure access.

Problem

Cannot list table spaces

Data Protector reports:

```
Cannot list table spaces.
```

Action

- Ensure that the database is not in a backup/restore/rollforward pending state.
- Ensure that user `root` (UNIX only) and the DB2 user are in both the DB2 and Data Protector `admin` groups.

Problem

Restore from an object copy hangs

Action

Before restarting the restore:

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

Problem

Restore finishes successfully, but rollforward fails

When performing a rollforward recovery from an online backup, restore finishes successfully, but rollforward fails.

Action

Ensure that the archived logs are available. If they are not, restore them from the last backup.

3 Integrating Lotus Notes/Domino Server and Data Protector

Introduction

This chapter explains how to configure and use the Data Protector Lotus Notes/Domino Server integration. It describes the concepts and methods you need to understand to back up and restore Lotus Notes/Domino Server.

Data Protector integrates with Lotus Notes/Domino Server to back up databases and transaction logs online. During backup, the database can be actively used.

Data Protector backs up all types of databases: storage databases, templates, and mailboxes (NSF, NTF, and BOX files). You can back up and restore individual databases or the whole server (all databases under Lotus Notes/Domino Server).

You can also back up:

- Archived transaction logs when `archived` logging is in effect.
- The current transaction log if Lotus Notes/Domino Server 5.0.4 or later is installed.

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

Table 20 Lotus Notes/Domino Server backup types

Full	Backs up all the selected Lotus Notes/Domino Server databases. If archived logs are selected, it also backs up the archived logs that have not been backed up yet, including the log currently in use.

Incremental	Backs up the selected Lotus Notes/Domino Server databases that meet the following condition: the size of the changes made to a database since it was last backed up exceeds the size set in the **Amount of log** option. Databases that do not meet this condition are not backed up. If archived logs are selected, it also backs up the archived logs that have not been backed up yet.

Data Protector offers the following restore options:

- Restore without recovery.
- Restore of a specific backup version of a Lotus Notes/Domino Server database and the possibility of applying changes made since the backup from the transaction log.
- Recovery of Lotus Notes/Domino Server databases to a specific point in time or to the latest possible consistent state.
- Restore of databases to a Lotus Notes/Domino Server location other than originally backed up from.
- Automatic restore of archived transaction logs in the case of recovery.

A database restore is possible even while Lotus Notes/Domino Server is running, with no impact on other databases currently in use. To enable a recovery using the logs from an online backup, Lotus Notes/Domino Server must be set to use `archived` transaction logging.

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

Integration concepts

The Data Protector Lotus Notes/Domino Server integration provides online backup, restore, and recovery of Lotus Notes/Domino Server, using the Lotus C API. Figure 33 on page 109 shows the architecture of the integration.

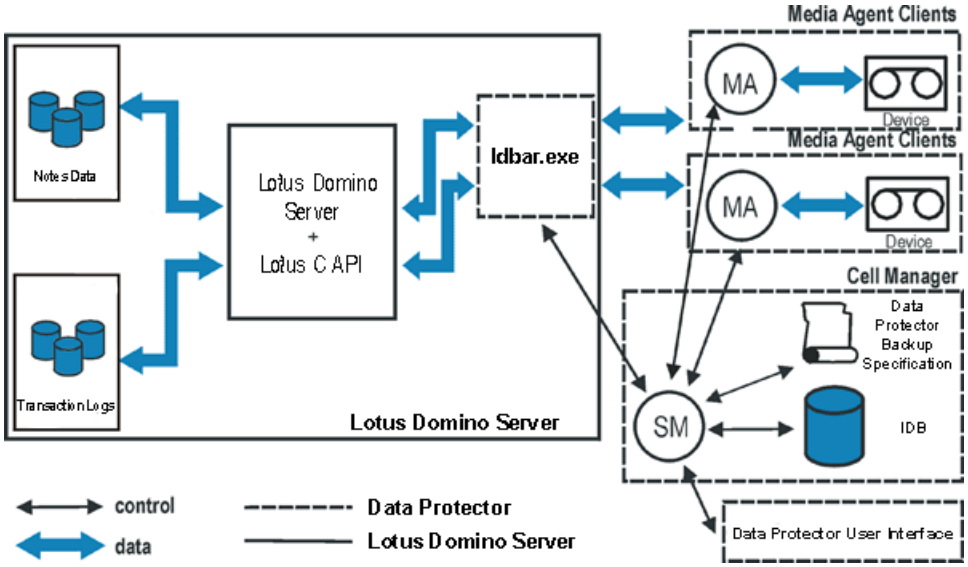

Figure 33 Data Protector Lotus Notes/Domino Server integration architecture

Table 21 Legend

SM	Data Protector Session Manager: Backup Session Manager during backup and Restore Session Manager during restore.
ldbar.exe	The central component of the integration, installed on the Lotus Notes/Domino Server system, which controls activities between Lotus Notes/Domino Server and Data Protector backup and restore processes.
Lotus C API	The Lotus-defined interface that enables data transfer between Data Protector and the Lotus Notes/Domino Server.
Notes Data	A set of Lotus Notes/Domino Server databases, where users create, update, store, and track documents in various formats.
MA	Data Protector General Media Agent.
Backup Specification	A list of objects to be backed up, backup devices, and options to be used.
IDB	The Data Protector Internal Database.

Lotus Notes/Domino Server databases are backed up in parallel streams, each stream transferring multiple databases. The number of streams equals the sum of concurrencies of all the devices used. The concurrency is defined in the backup specification.

Lotus Domino Cluster

Data Protector supports Lotus Domino Cluster. Unlike operating system clusters (MSCS, MC/ServiceGuard cluster, and Veritas cluster), the Lotus Domino cluster is an end-application cluster. This means it does not provide the cluster resources failover to a secondary cluster node if the primary cluster node becomes unavailable; it just ensures a Lotus client can access a replica database on another Domino server if the Domino database on the initial Domino server becomes unavailable for connection. All servers in a Domino cluster continually communicate with each other to keep updated on the status of each server and to keep database replicas synchronized.

The Domino cluster also lets you set limits for workload balancing, track the availability of servers and databases, and add servers and databases to the cluster. To take advantage of failover and workload balancing, databases and replicas are distributed throughout the cluster. It is not necessary to maintain replicas of every database on every server; the number of replicas created for a database depends on how busy the database is and how important it is for users to have constant access to that database.

NOTE:

- The Lotus Domino cluster must be part of the Lotus Domino 6 Enterprise Server or the Lotus Domino 6 Utility Server.
- Lotus Domino 6 covers both Lotus Domino 6.0.x and 6.5.x versions. The cluster functionality is the same in both versions.

Replicas

Replicas make a database available to users in different locations, on different networks, or in different time zones. If a replica is available on one or more local servers, users do not need to connect to the single central server.

All replicas share a *replica ID*, assigned when the database is first created. Although replicas can have different file names, can contain different documents, and have different database designs, as long as they have the same replica ID, replication can occur between them. A replica is not the same as a copy of a database. A copy may

look the same as the original, but because it does not share a replica ID with the original database, it cannot replicate with it.

Replication in a cluster

Cluster replication is *event-driven*, rather than schedule-driven:

- When the Cluster Replicator (a Lotus Domino cluster component) is aware of a change in a database, it immediately pushes that change to other replicas in the cluster.
- If there is a backlog of replication events, the Cluster Replicator stores these in memory until it can push them to the other cluster servers.
- If a change to the same database occurs before a previous change has been sent, the Cluster Replicator pools these changes and sends them together to save processing time.

Because Domino stores replication events only in memory, both the source and destination servers must be available for the replication to complete successfully. If a destination server is not available, the Cluster Replicator continues to store the events in memory, and attempts periodically to push them to the destination server until it becomes available. The interval between these attempts starts at one hour and increases over time to a maximum of one day.

If the source server shuts down before replication completes, the replication events in memory are lost. For this reason, you should use standard replication (the REPLICA task) to perform immediate replication with all members of the cluster whenever you restart a cluster server. It is also a good idea to schedule regular replication between cluster servers, such as several times per day, to ensure the databases remain synchronized. The Cluster Replicator always attempts to make all replicas identical so that users who fail over do not notice that they failed over.

Failover in a cluster

A cluster's ability to redirect requests from one server to another is called *failover*. If you try to access a database on a server that is unavailable or under heavy load, Domino directs you to a replica of the database on another server in the cluster, so that failover is essentially transparent to you.

Example

This example describes the process that Domino uses when it fails over. This cluster contains three servers. Server 1 is currently unavailable. The Cluster Managers on Server 2 and Server 3 are aware that Server 1 is unavailable.

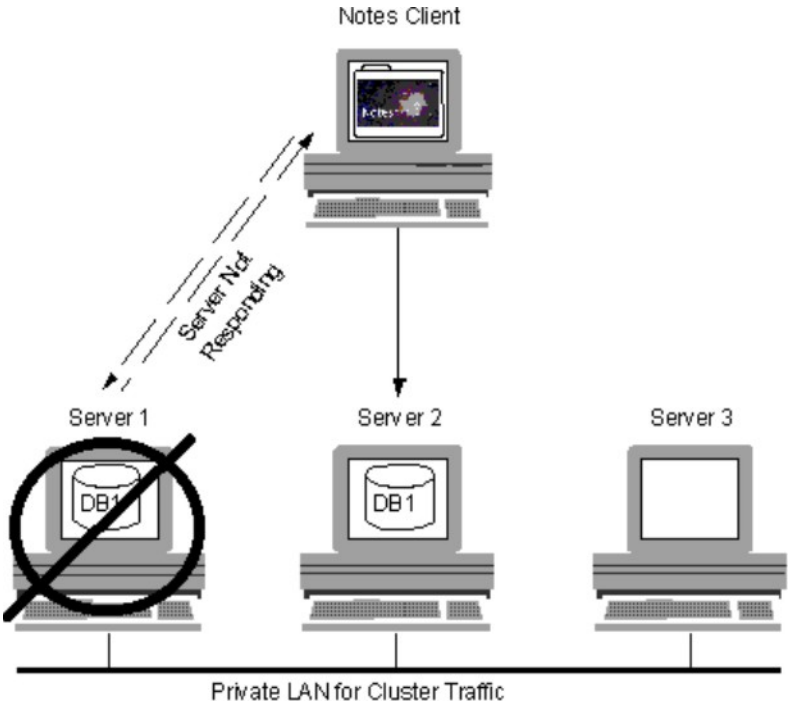

1. A Notes user attempts to open a database on Server 1.
2. Notes realizes that Server 1 is not responding.
3. Instead of displaying a message that says the server is not responding; Notes looks in its cluster cache to see if this server is a member of a cluster and to find the names of the other servers in the cluster.
4. Notes sends a query to the Cluster Manager, which looks in the Cluster Database Directory to find which servers in the cluster contain a replica of the desired database, and finds the availability of the servers.
5. The Cluster Manager sends a list of the servers it has found to Notes, sorted in order of availability.
6. Notes opens the replica on the first server in the list. If that server is no longer available, Notes opens the replica on the next server in the list. In this example, Server 2 was the most available server.

When the Notes client shuts down, it stores the contents of the cluster cache in the file `CLUSTER.NCF`. Each time the client starts, it populates the cluster cache from the information in `CLUSTER.NCF`.

Configuring the integration

You need to configure a Lotus Notes/Domino Server user and every Lotus Notes/Domino Server you intend to back up or restore.

Prerequisites

- Ensure that you have correctly installed and configured Lotus Notes/Domino Server.
  - For supported versions, platforms, devices, and other information, see the *HP Data Protector product announcements, software notes, and references* or http://www.hp.com/support/manuals.
  - For information on installing, configuring, and using Lotus Notes/Domino Server, see the Lotus Notes/Domino Server documentation.

  Lotus Domino Cluster: When configuring the Lotus Domino cluster, decide if you need a private LAN for the cluster. The main benefit is to separate the network traffic created by the cluster when it uses cluster replication and server probes, thus leaving more bandwidth available on primary LAN. If you anticipate a lot of cluster replication activity, create a private LAN. To do this, install an additional network interface card in each cluster server and connect these cards through a private hub or switch.
- Ensure that you have correctly installed Data Protector. For information on how to install Data Protector in various architectures, see the *HP Data Protector installation and licensing guide*.

  Every Lotus Notes/Domino Server system you intend to back up from or restore to must have the Data Protector Lotus Integration component installed.

Before you begin

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

Transaction logging of Lotus Notes/Domino Server

To enable recovery from an online backup, Lotus Notes/Domino Server must be set to use transaction logging. This way, transactions are stored to the transaction log

directory and can be used to apply or undo database transactions during database recovery.

You can perform daily full backups of transaction logs instead of full database backups.

After enabling transaction logging, all databases are automatically logged. With transaction logging enabled, multiple `S0000000.TXN` files may appear in the log directory.

Table 22 Transaction logging styles

Linear (circular) logging	The default mode. Lotus Notes/Domino 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 can recover only transactions stored in the transaction log. Archiving of transaction logs is not possible.
Archived logging	Lotus Notes/Domino Server does not reuse log extents until they are backed up. The system uses transaction logs to apply or undo database transactions not flushed to disk for databases that were open when system failure occurred.

IMPORTANT:

To back up log files in an incremental backup, transaction logging must be set to `archived` logging.

Enabling transaction logging

Use Lotus Domino Administrator on the Lotus Notes/Domino Server system. Alternatively, use Web Administrator or edit the `notes.ini` file.

In a cluster environment, enable transaction logging on all cluster nodes.

To enable transaction logging and set `archived` logging:

1. Start Lotus Domino Administrator.
2. Log on to Lotus Notes/Domino Server and select the **Configuration** tab.

3. Expand **Server**, select **All Server Documents**, and select the desired Lotus Notes/Domino Server. See Figure 34 on page 115.

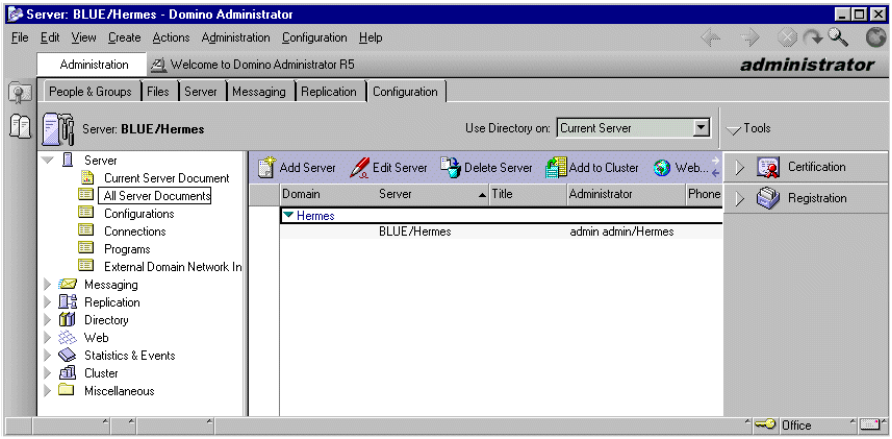

Figure 34 Browsing Lotus Notes/Domino Server

4. Select the **Transactional Logging** tab and set appropriate values. See Figure 35 on page 115.

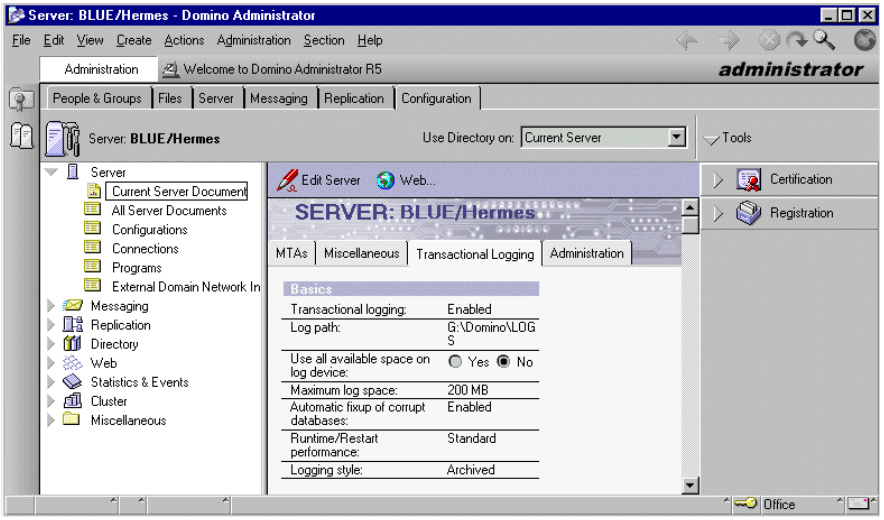

Figure 35 Enabling archived transactional logging

5. Save the settings and restart Lotus Notes/Domino Server for the changes to take effect.

Configuring Lotus Notes/Domino Server users

On UNIX, add the Lotus Notes/Domino Server administrator to the Data Protector `admin` or `operator` user group. You need to specify this user in backup specifications. By default, this user is `notes` in the group `notes`.

Additionally, add the operating system user `root` on the Lotus Notes/Domino Server system to the Data Protector `admin` or `operator` user group.

For information, see the online Help index: “adding users”.

Configuring Lotus Notes/Domino Server systems

Using the Data Protector GUI

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

4. In Client, select the**Lotus Notes/Domino Server system**. In a cluster environment, select the virtual server.

   In **Application database**, select the name of the Lotus Notes/Domino Server to be backed up.

   UNIX only: Enter the username and user group of the Lotus Notes/Domino Server administrator. This user will be the owner of the backup.

   See Figure 36 on page 117.

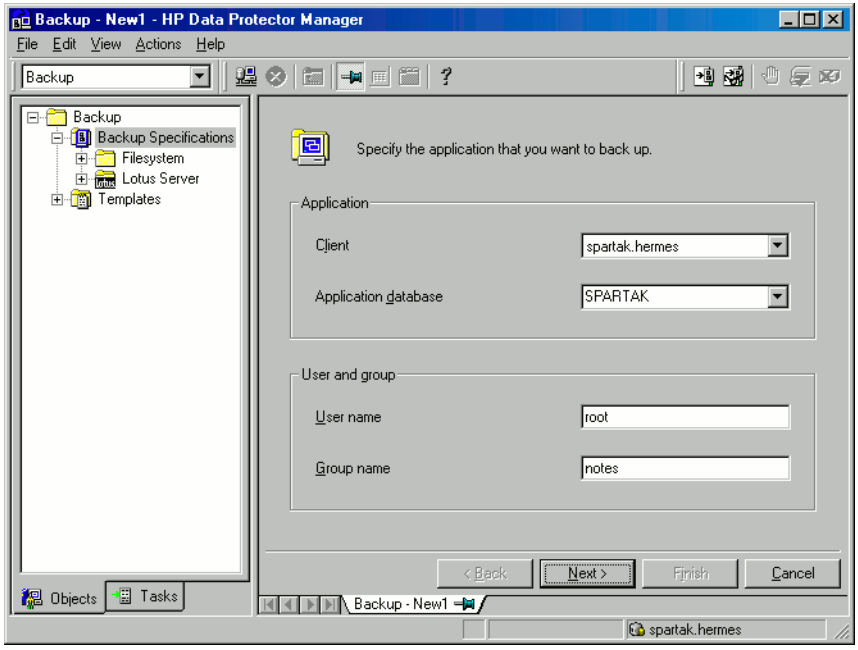

Figure 36 Specifying the Lotus Notes/Domino Server system

Click **Next**.

5. In the Configure Lotus dialog box, specify the pathname of the `notes.ini` file on the Lotus Notes/Domino Server system.

   Review and, if necessary, update other automatically determined options.

   See Figure 37 on page 118.

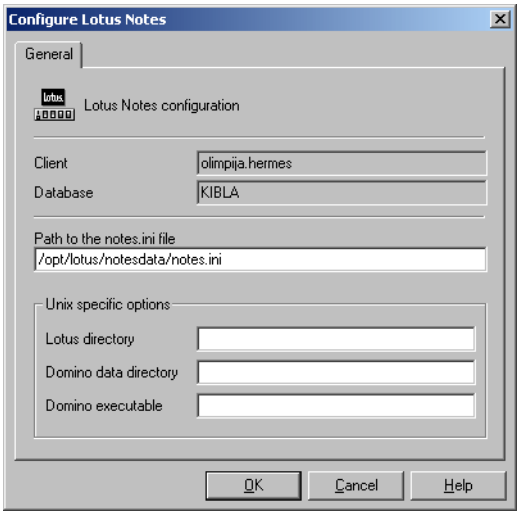

Figure 37 Specifying Lotus Notes/Domino Server data

   Click **OK**.

   If an error occurs, click **Details** or see "Troubleshooting" on page 139.

6. The integration is configured. Exit the GUI or proceed with creating the backup specification at Step 6 on page 123.

Using the Data Protector CLI

On the Lotus Notes/Domino Server system, run:

Windows:

```
Data_Protector_home\bin\util_notes.exe -CONFIG
-SERVER:SRV_NAME -INI:notes.ini_file
```

Solaris:

```
/opt/omni/lbin/util_notes.exe -CONFIG -SERVER:SRV_NAME
-INI:notes.ini_file [-HOMEDIR:Lotus_home_directory]
[-DATADIR:Domino_data_directory]
[-EXECDIR:Domino_executables_directory]
```

AIX:

```
/usr/omni/bin/util_notes.exe -CONFIG -SERVER:SRV_NAME
-INI:notes.ini_file [-HOMEDIR:Lotus_home_directory]
[-DATADIR:Domino_data_directory]
[-EXECDIR:Domino_executables_directory]
```

Parameter description

SRV_NAME	Lotus Notes/Domino Server name.
notes.ini_file	Pathname of the Lotus Notes/Domino Server `notes.ini` file.
Lotus_home_directory	Pathname of the Lotus Notes/Domino Server home directory.
Domino_data_directory	Pathname of the Lotus Notes/Domino Server data directory.
Domino_executables_directory	Pathname of the Lotus Notes/Domino Server executables directory.

NOTE:

UNIX: If the `-HOMEDIR`, `-DATADIR`, and `-EXECDIR` options are not specified, the values are automatically read from the `notes.ini` file.

The message `*RETVAL*0` indicates successful configuration.

Examples

Windows:

```
Data_Protector_home\bin\util_notes.exe -CONFIG -SERVER:BLUE
-INI:d:\Lotus\Domino\BLUE\notes.ini
```

Solaris:

```
/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

You can check the configuration of the Lotus Notes/Domino Server using the Data Protector GUI after you have created at least one backup specification for the Lotus Notes/Domino Server. If you use the Data Protector CLI, a backup specification is not needed.

Using the Data Protector GUI

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

Using the Data Protector CLI

On the Lotus Notes/Domino Server system, from the directory:

Windows: `Data_Protector_home\bin`

Solaris: `/opt/omni/lbin`

AIX: `/usr/omni/bin`

run:

```
util_notes.exe -CHKCONF -SERVER:SRV_NAME
```

Data Protector checks the path to the specified directories and files.

The message `*RETVAL*0` indicates successful configuration.

Handling errors

If an error occurs, the error number is displayed in the form `*RETVAL*error_number`.

To view the error description:

Windows: *On the Cell Manager, see the file* `Data_Protector_home\help\enu\Trouble.txt`

Solaris: Run:

```
/opt/omni/lbin/omnigetmsg 12 error_number
```

AIX: Run:

```
/usr/omni/bin/omnigetmsg 12 error_number
```

Backup

The integration provides backup of the following types:

Table 23 Lotus Notes/Domino Server backup types

Full	Backs up all the selected Lotus Notes/Domino Server databases. If archived logs are selected, it also backs up the archived logs that have not been backed up yet, including the log currently in use.
Incremental	Backs up the selected Lotus Notes/Domino Server databases that meet the following condition: the size of the changes made to a database since it was last backed up exceeds the size set in the **Amount of log** option. Databases that do not meet this condition are not backed up. If archived logs are selected, it also backs up the archived logs that have not been backed up yet.

What is backed up?

Lotus Notes/Domino Server databases consist of the following files:

- Notes Storage Facility files (**NSF** files)
- Notes Template Facility files (**NTF** files) - templates for creating new NSF databases
- Mailbox files (**BOX** files) - files used by the mail router
- Transaction log files, named `SXXXXXXX.TXN`, where `XXXXXXX` is a 7–digit number that is automatically incremented for every new transaction file

  Lotus Notes/Domino Server automatically recycles archived transaction logs after backup.

IMPORTANT:

Back up archive logs frequently. Once they are backed up, Lotus Notes/Domino Server overwrites them with new log entries when needed. Otherwise, new log files are created, which consume additional disk space. Since the archive logging style does not have any size limit as far as the amount of log files is concerned, you may run out of disk space.

To delete all backed up archive logs, restart the Lotus Notes/Domino Server instance. Manual deletion of archive logs is not recommended.

TIP:

To speed up a Lotus Notes/Domino Server backup, exclude NTF files from the backup specification. Create a separate backup specification to back up NTF files. These files do not need to be backed up frequently because they do not change.

What is not backed up?

You *must* back up the following non-database files using a filesystem backup:

- `notes.ini`
- `desktop.dsk`
- all `*.id` files

Considerations

- **Lotus Domino Cluster:** Back up the replica database from a Domino server in the same way as a normal Domino database.

  Unlike operating system clusters, there are no virtual servers or virtual IP addresses involved with a Domino cluster, so when creating a Data Protector backup specification, select common physical hostnames for the backed up source databases.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context list, click **Backup**.

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

   In **Application database**, select the Lotus Notes/Domino Server to be backed up.

   UNIX only: Type the username and user group of the Lotus Notes/Domino Server administrator, who will be the owner of the backup.

   Click **Next**.
5. If Lotus Notes/Domino Server is not configured yet for use with Data Protector, the Configure Lotus dialog box is displayed. Configure the integration as described in "Configuring Lotus Notes/Domino Server systems" on page 116.
6. Select the Lotus Notes/Domino Server objects to be backed up. See Figure 38 on page 123.

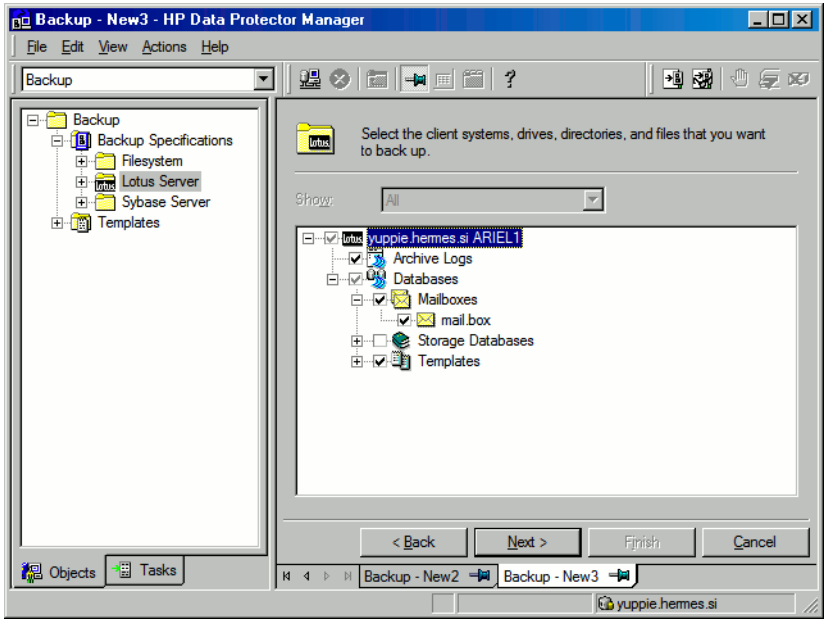

Figure 38 Selecting backup objects

Click **Next**.

7. Select devices to use for the backup.

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

   Click **Next**.

8. Set backup options. For information on the application specific options (Figure 39 on page 124), see Table 24 on page 125 or press **F1**.

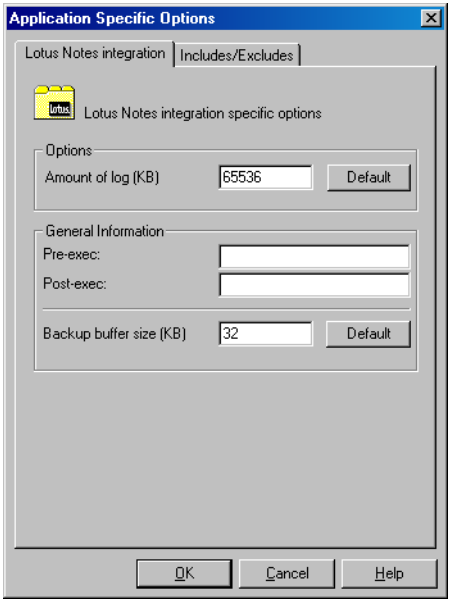

Figure 39 Application specific options

Click **Next**.

9. Optionally, schedule the backup and click **Next**. See "Scheduling backup specifications" on page 125.

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

TIP:

Preview backup session for your backup specification before using it. See "Previewing backup sessions" on page 126.

Table 24 Lotus Notes/Domino Server backup options

Amount of log	Applies to incremental backups. If the database to be backed up has a smaller log amount than specified, the backup skips the database. If the database exceeds the specified log amount, a full backup of the database is performed.
Pre-exec, Post-exec	Specify a command that will be started by `ldbar.exe` on the Lotus Notes/Domino Server system before the backup (`pre-exec`) or after it (`post-exec`). The command must reside in the directory: ***Windows:*** *`Data_Protector_home`*`\bin` ***Solaris:*** `/opt/omni/lbin` ***AIX:*** `/usr/omni/bin` In the backup specification, provide only the filename.
Backup buffer size	The size of the buffer used for reading and writing data during the backup.

Modifying backup specifications

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

Scheduling backup specifications

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

Scheduling example

To back up Lotus Notes/Domino Server at 9:00, 13:00, and 18:00 during weekdays:

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

2. Under Recurring, select **Weekly**. Under Time options, select **9:00**. Under Recurring Options, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**. See Figure 40 on page 126.

   Click **OK**.

3. Repeat Step 1 on page 125 and Step 2 on page 126 to schedule backups at 13:00 and 18:00.

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

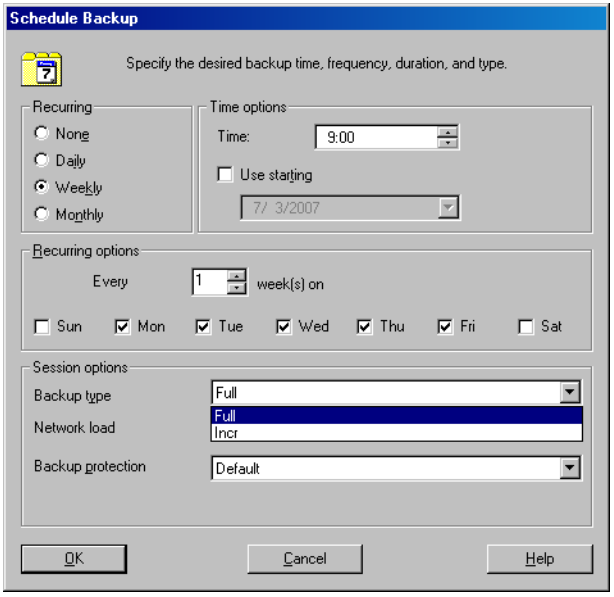

Figure 40 Scheduling backups

Previewing backup sessions

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

Using the Data Protector GUI

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

3. Specify the **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

A test can be performed on the Lotus Notes/Domino Server system or on any Data Protector client system within the same Data Protector cell with the Data Protector User Interface installed.

From the directory:

Windows: *`Data_Protector_home`*`\bin`

Solaris: `/opt/omni/bin`

AIX: `/usr/omni/bin`

Run:

```
omnib -lotus_list backup_specification_name -test_bar
```

What happens during the preview?

The command tests the Data Protector part of the configuration. The following is tested:

- Communication between Lotus Notes/Domino Server and Data Protector.
- The syntax of the backup specification.
- If devices are correctly specified.
- If the necessary media are in the devices.

Starting backup sessions

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

You can start the backup using:

- The Data Protector GUI.
- The Data Protector CLI. See the `omnib` man page.

Using the Data Protector GUI

1. In the Context List, click **Backup**.

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

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

Restore

You can restore databases directly to the Lotus Notes/Domino Server system. When you restore a database, the database is brought offline, restored, and brought online. Transaction logs are also restored if needed. If recovery is selected, the restore of archived logs is performed automatically during the recovery process.

You can restore a database restore while the server is online, if the database is not being accessed. A newly-restored Lotus Notes/Domino Server database is not active. If you access it, it will automatically be brought online, but a recovery using the backed up logs will not be performed. To get the last possible consistent state of the databases or to perform a recovery to a specific point in time, use the Recover option.

You can restore a database to:

- Its original location at backup time.

  Select this to replace a corrupted or deleted database.
- A different location.

  Select this to keep the original database intact.

Recovery to a different client system is not possible.

To restore Lotus Notes/Domino Server databases, use the Data Protector GUI or CLI.

Finding information for restore

You can find details on backup sessions and the media used in the Data Protector IDB. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

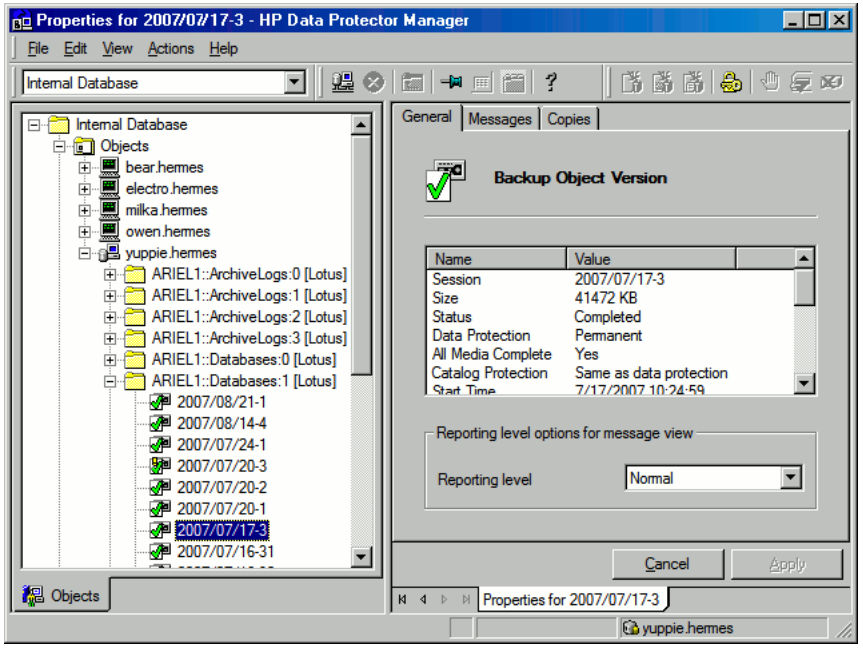

Figure 41 Example of session properties

TIP:

To see which files are contained in the backup object, click the **Messages** tab. Backup objects with the same name (for example, `ARIEL:Databases:1 [Lotus]`), created in different sessions, may contain different files.

Using the Data Protector CLI

1. Go to the directory:

   Windows: *`Data_Protector_home`*`\bin`

   Solaris: `/opt/omni/bin/`

   Other UNIX systems: `/usr/omni/bin/`

2. Get a list of Lotus Notes/Domino Server objects created in a particular session:

   ```
   omnidb -session session_id
   ```

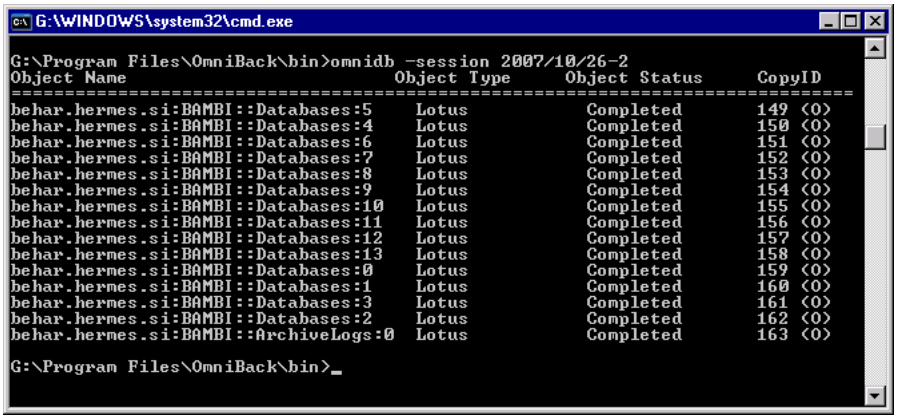

Figure 42 Lotus Notes/Domino Server objects from a particular session

3. See which Lotus Notes/Domino Server databases are contained in a particular Lotus Notes/Domino Server object from a particular session:

   ```
   omnidb -lotus client:Lotus_instance::stream_id -session
   session_id -catalog
   ```

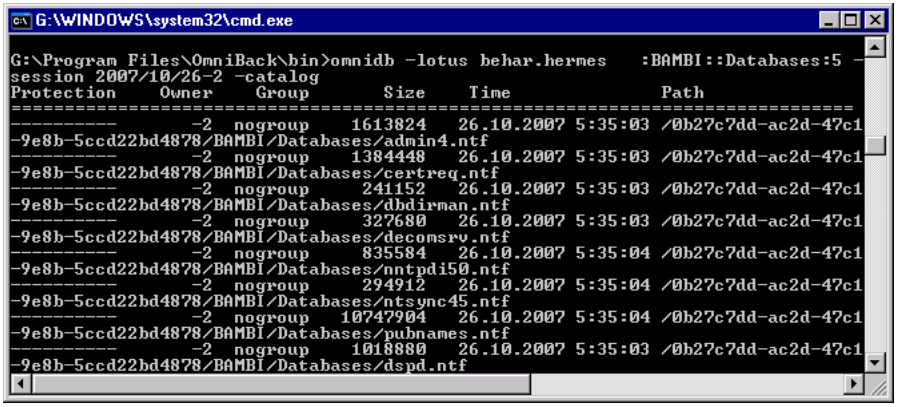

Figure 43 Lotus Notes/Domino Server databases of a particular object

For details, see the `omnidb` man page or the *HP Data Protector command line interface reference*.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Lotus Server**, expand the client from which the data was backed up, and select the instance you want to restore.

3. In the Source page, select objects for restore. See Figure 44 on page 132.

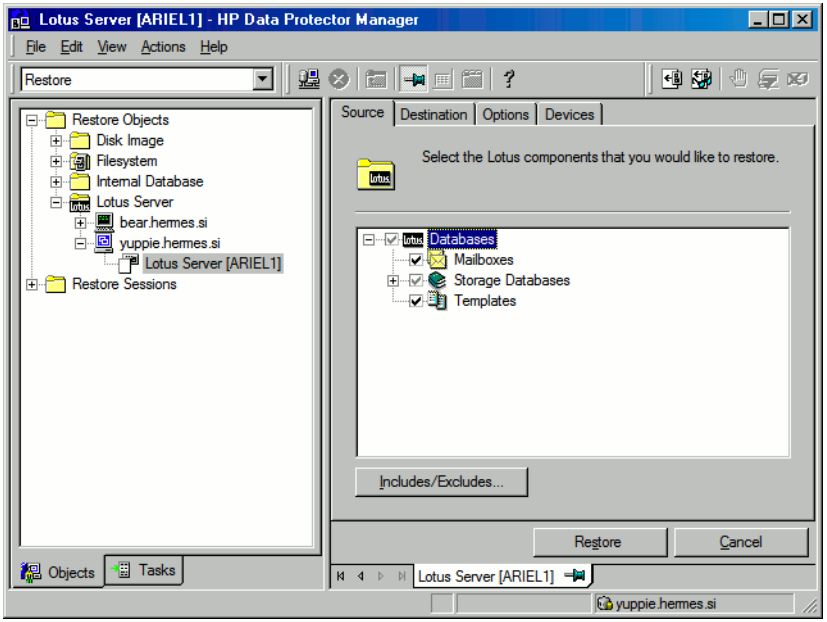

Figure 44 Selecting objects for restore

NOTE:

In the Source page all backed up databases are listed. When restoring multiple databases from a specific backup session, ensure that the databases were backed up in the selected backup session. If not, the warning `Object not found in the database` appears at 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, the Lotus Integration Agent finds the latest backup version of each database for restore.

You can select the backup version in the Options page (Figure 45 on page 134). Click **Browse** to select a different version of backup.

4. In the Destination page, set the destination options. For information, see Table 25 on page 135 or press **F1**.

IMPORTANT:

If you restore to a location where a database with the same file name resides as the one being restored, then this database is taken offline and deleted.

5. In the Options page, set the restore options (Figure 45 on page 134). For information on the application specific options, see Table 26 on page 136 or press **F1**.

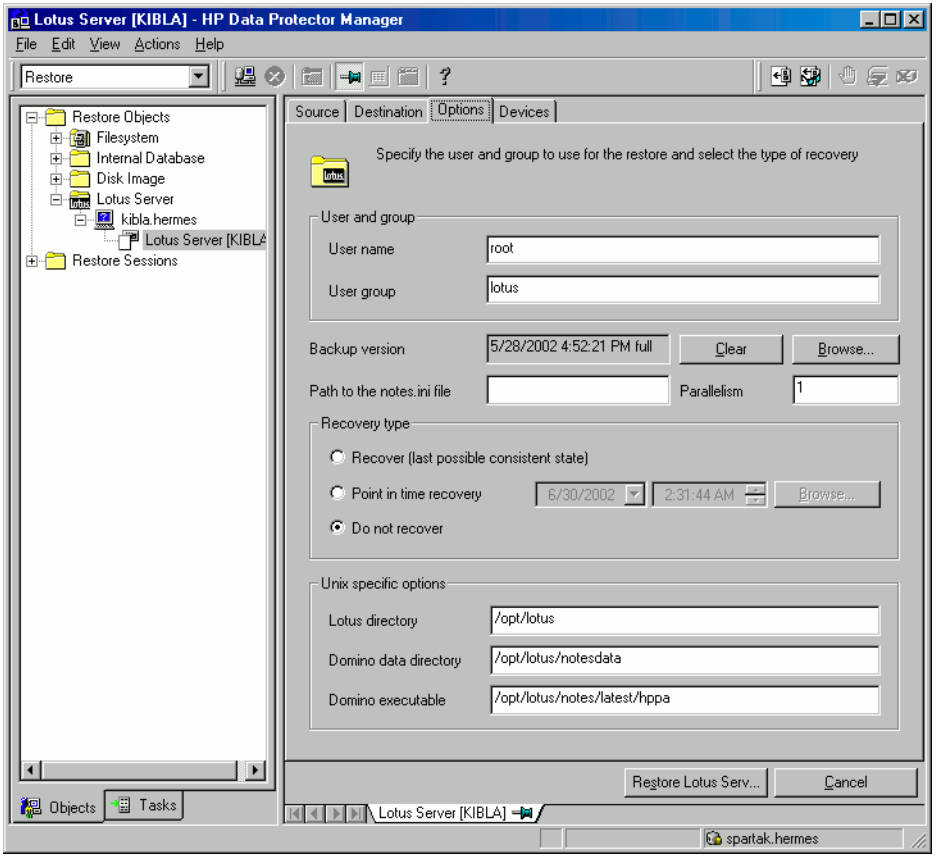

Figure 45 Lotus Notes/Domino Server restore options

6. In the Devices page, the devices and media for restore are automatically selected.You can also restore using a device other than that used for backup. For information on how to specify another device for restore, see the online Help index: “restore, selecting devices for”.

   Click **Restore**.

7. In the Start Restore Session dialog box, click **Next**.

8. Specify the **Report level** and **Network load**.

   Click **Finish** to start the restore.

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

Restoring using the Data Protector CLI

For details, see the `omnir` man page.

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

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

If not set, names of backup objects returned by the Data Protector CLI commands (for example `omnidb`) may not be usable when providing the parameters to other Data Protector commands (for example `omnir`).

Restore options

Specify destination and restore options specific to the Data Protector Lotus Notes/Domino Server integration. If the target system is a UNIX system, specify UNIX specific options as well.

Table 25 Destination options

Restore to client	By default, Lotus Notes/Domino Server databases are restored to the same client from which they were backed up. To restore to another client, select the new client from the drop-down list or type its name in the text box. The client must be part of the Data Protector cell and have the Lotus Notes/Domino Server integration installed.
Restore to instance	By default, Lotus Notes/Domino Server databases are restored to the same Lotus Notes/Domino Server instance from which they were backed up. To restore to another instance, select the new instance from the drop-down list or type its name in the text box. The instance must be configured for use with this integration.
Restore to the original location	By default, databases are restored to the same directory from which they were backed up (either on the original system or on some other system you selected).

Restore to a new location	This option enables you to restore your data to another directory. Specify the relative path to the Lotus Notes/Domino Server data directory where you want to restore your data.

Example

Lotus Notes/Domino Server data directory is located in:

Windows: `C:\Lotus\Domino\BLUE\`

UNIX: `/opt/lotus/notesdata/BLUE/`

To restore a database to the directory:

Windows: `C:\Lotus\Domino\BLUE\restore_dir\`

UNIX: `/opt/lotus/notesdata/BLUE/restore_dir/`

select **Restore to new location** and enter `type restore_dir`. The restored database filenames are the same as they were at backup time.

Table 26 Restore options

Restore options	**Username**	***UNIX only:*** Username of the Lotus Notes/Domino Server backup owner, for example, `notes`.
	User group	***UNIX only:*** User group of the Lotus Notes/Domino Server backup owner, for example, `notes`.
	Backup version	By default, a restore is performed from the last full backup of the database. Click `Browse` to select a backup version other than the last one.
	Parallelism	Specify how many parallel streams should be used to restore your data. Default: 1.
Recovery type options	**Recover (last possible consistent state)**	Select this to recover the database to the last possible consistent state. This also includes the restore of archived transaction logs if needed during recovery.

	Point in time recovery	The point in time to which the database state should be recovered. Click **Browse** to specify the desired date and time. Only transactions written before the specified date and time are applied to the database.
	Do not recover	The default option. Select this to restore databases without recovering them from the backed up logs. Transactions made after the backup are not reflected in the restored databases.

Restore in Lotus Domino Cluster environment

The following are typical cases to consider when restoring a Domino database.

Restoring a replica database without recovery

In this case, the replica database is restored to the state it was in at the time of backup. The contents of archive logs are ignored, so recovery to the latest possible state is not performed.

The Lotus Domino Cluster Server containing the replicated database preserves its latest state even if you use the "Push" replication style when replicating the restored replica database from the restored target Domino Cluster Server to the Domino Cluster Server containing the replicated database.

If you use the "Pull" or "Push/Pull" replication style to replicate the restored replica database, the restored replica database is recovered to the latest state just like the replicated database. The state gathered after the restore will be lost.

If the restored replica database is not to be replicated and recovered to the last consistent state, then it should never be replicated with the "Pull" or "Push/Pull" replication style from the restored Domino Cluster Server.

Restoring with recovery to the latest possible state

In this case, the database is restored and recovered to the latest possible state by applying the archive logs from the target system. If another Lotus Domino Cluster Server contains a replica of the restored database, this replica will already be in the latest state.

If the archive logs from the restore target Lotus Domino Cluster Server do not allow recovery to the latest state, use the "Pull" or "Push/Pull" replication style from the restored target Domino Cluster Server to the other Domino Cluster Servers containing the replicas in order to replicate the restored database and bring it up to the latest state

Point-in-time recovery

In this case, the database is restored to the point-in-time state as it was at the selected backup time, no matter what the latest archive logs contains.

If another Lotus Domino Cluster Server contains a replica of the restored database that is in a more recent state than the restored one, the replica will preserve its latest state even if you use the "Push" replication style when replicating the database from the restored target Domino Cluster Server to the other Domino Cluster Server containing the replica database.

If you use the "Pull" or "Push/Pull" replication style to replicate the point-in-time recovered database from the restored target Domino Cluster Server to the other Domino Cluster Server containing the replica database, the point-in-time recovered database will be recovered to the latest state just like the replica database. The state gathered after the point-in-time recovery will be lost.

If the point-in-time recovered database is not to be replicated and recovered to the last consistent state, then it should never be replicated with the "Pull" or "Push/Pull" replication style from the restored Domino Cluster Server. You can also achieve this as follows:

1. Delete all replica databases of the replicated database from other Domino Cluster Servers before the restore.
2. Restore the replicated database as described above.
3. Create new replicas of the replicated database on the Domino Cluster Servers from which you deleted replicas in step 1.

In this way, the replicas will contain the restored point-in-time state, not the latest state.

Restoring to a new location

In this case, the database is restored to a new location with the same ID as the original replicated database and its replicas. The new database is treated as a replica database. The restored state depends on the type of restore/recovery you select in **Options -> Recovery type**.

To decide in which state you want the databases to be restored or recovered, see "Restoring a replica database without recovery" on page 137, "Restoring with recovery to the latest possible state" on page 137, and "Point-in-time recovery" on page 138.

Performance tuning

The time needed for *backup* can be significantly reduced by fine-tuning the following backup device parameters:

- Concurrency
- Block size

Concurrency has a much greater impact on backup performance than block size. Tests have shown that better results are achieved when using lower concurrency values and a medium block size (256 KB). The optimum values still depend on your environment.

For information on the concurrency and block size parameters, see the online Help index: "concurrency", "block size", and "backup devices, advanced options".

The *restore* performance can be additionally improved by setting the **Parallelism** option as high as possible. As a result, Data Protector automatically creates the optimum number of streams.

Monitoring sessions

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

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

For information on how to monitor a session, see the online Help index: "viewing currently running sessions".

Troubleshooting

This section lists Lotus Notes/Domino Server checks, general checks and verifications, plus problems you might encounter when using the Data Protector Lotus Notes/Domino Server integration. Start at "Problems" on page 143 and if you cannot find a solution there, go through the checks and verifications.

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

Before you begin

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

Checking the Lotus Notes/Domino Server side

If you encounter errors when performing the following checks, contact Lotus Notes/Domino Server support. For more information on these procedures, see the Lotus Notes/Domino Server documentation.

Windows:

- Check if the `nNotes.dll` library is linked. Run:

  `Data_Protector_home``\bin\util_notes.exe -chkconf`

Checks and verifications

If your configuration, backup, restore, or recovery failed:

- Examine system errors reported in the `debug.log` file on the Lotus Notes/Domino Server system, located in the directory:

  Windows: *`Data_Protector_home`*`\log`

  Solaris: `/var/opt/omni/log/`

  AIX: `/usr/omni/log/`

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

  For details, see "Verifying Data Protector Client Installation" in the *HP Data Protector installation and licensing guide*.

- Check whether the Data Protector Lotus Integration Agent `ldbar.exe` is installed on the system.
- ***Windows only:*** Verify the inet startup parameters on the Lotus Notes/Domino Server system.

Make sure the Data Protector Inet service is running under a user that is a member of the Data Protector `admin` user group. For information, see the online Help index: "Inet, changing account".

- Check the `omnirc` environment settings.

  For information on how to use the `omnirc` file, see the *HP Data Protector troubleshooting guide*.

- Check errors during the backup or restore session.

  Error related to Lotus Notes/Domino Server take the following form:

  ```
  Lotus ERROR [error #]: Error description
  ```

  Examine the error description and take appropriate actions.

Additionally, if your backup failed:

- Check your Lotus Notes/Domino Server configuration as described in "Checking the configuration" on page 120.
- Perform a filesystem backup of the Lotus Notes/Domino Server system.

  Observe session messages and examine system errors reported in the `debug.log` file on the

  - Data Protector Lotus Notes/Domino Server client if the Lotus Notes/Domino Server part of the filesystem backup fails.
  - Data Protector Cell Manager system if the Data Protector part of the filesystem backup fails.

- Verify Data Protector internal data transfer using the `testbar` utility.

  1. From the directory:

     Windows: *`Data_Protector_home`*`\bin`

     Solaris: `/opt/omni/bin/utilns`

     AIX: `/usr/omni/bin/utilns`

     run:

     ```
     testbar -type:Lotus -appname:SRV_NAME
     -bar:backup_specification_name -perform:backup
     ```

  2. Create a Lotus Notes/Domino Server backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices.

- Start a backup session using `ldbar.exe`.

  You can start a backup of a single database using the Data Protector CLI, specifying backup options as `ldbar.exe` command line options.

On the Data Protector Lotus Notes/Domino Server client, from the directory:

Windows: *Data_Protector_home*\bin

Solaris: /opt/omni/bin

AIX: /usr/omni/bin

run:

Windows:

```
ldbar.exe -perform:backup -db: DB_NAME -server: SRV_NAME
[-ini:Path_to_notes.ini_file] -bar:
backup_specification_name
```

UNIX:

```
ldbar.exe -perform:backup -db:DB_NAME -server:SRV_NAME
[-ini:Path_to_Notes.ini_file] -bar:backup_specification_name
[-homedir:PathToLotusHome] [-datadir:path to Domino data]
[-execdir:PathToDominoExecutables]
```

The `-bar` option is mandatory because `ldbar.exe` reads the device options from the backup specification as opposed to other options in the backup specification, which are ignored. Command line options are used instead.

For other `ldbar.exe` parameters, run `ldbar.exe -help`.

- ***Windows only:*** When Lotus Notes/Domino Server and Windows Terminal Services coexist on the same system and Lotus Notes/Domino Server is started from the terminal client program, Lotus Notes/Domino Server backup cannot be performed.

  Windows Terminal Services should not be used to manage Lotus Notes/Domino Server. However, Lotus Notes/Domino Server backup can be performed when using the terminal service client program to start the Data Protector GUI on the system where Lotus Notes/Domino Server is running. Lotus Notes/Domino Server can be managed locally or with a VNC program.

Additionally, if your restore failed:

- Perform a test restore of any filesystem on the problematic client.
- Test a restore session using the `ldbar.exe` command on the Data Protector Lotus Notes/Domino Server system. From the directory:

  Windows: *Data_Protector_home*\bin

  Solaris: /opt/omni/bin

  AIX: /usr/omni/bin

  run:

  ```
  ldbar.exe -perform:restore -db:DB_NAME -server:SRV_NAME
  -ini:Path_to_notes.ini_file
  ```

For other `ldbar.exe` parameters, run `ldbar.exe -help`.

Additionally, if your recovery failed:

- Check if the recovery time parameter is set in a 24 hour format:
  `yyyy/mm/dd.hh:mm:ss`

Example

```
2006/01/25.18:15:00
```

Problems

Problem

Script failed error

While configuring or starting a backup using the Data Protector GUI, the following error is displayed:

```
Script failed. Cannot get information from remote host.
```

Action

For information on how to solve this problem, see "Checking the Lotus Notes/Domino Server side" on page 140.

Problem

Lotus Notes/Domino Server freezes during backup

Lotus Notes/Domino 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:

- The Lotus Notes C API initialization failed.
- ***UNIX only:*** If Lotus Notes/Domino Server is not online and the Lotus Notes/Domino Server daemon `logasio` is not running, then while the Lotus Integration Agent is initializing the Lotus C API, the `logasio` daemon automatically starts. Since the environment for user `notes` is not set because the `.profile` is not executed, the `logasio` server could fail to start.

Action

Kill the `ldbar.exe` or `logasio` processes:

1. ***UNIX only:*** Log in to the Lotus Notes/Domino Server system as user `root`.
2. ***Windows only:*** Kill all the `ldbar.exe` processes using `Task Manager`.
3. ***UNIX only:*** Kill all the `ldbar.exe` and `logasio` processes.
4. If Lotus Notes/Domino Server is running, restart it. Before restarting, ensure that no Lotus Notes/Domino Server processes are still running.
5. Log in as user `notes` and check if Lotus Notes/Domino Server recovered. From the directory:

   Windows: *`Data_Protector_home`*`\bin`

   Solaris: `/opt/omni/lbin`

   AIX: `/usr/omni/bin`

   run: `util_notes.exe -box -ini:`*`path_to_notes.ini`*

   If everything is working properly, the `*RETVAL*0` message is displayed.

NOTE:

On UNIX, you need to clean up shared memory and semaphores before restarting Lotus Notes/Domino Server.

Problem

Restore to another client fails

Action

Ensure that Lotus Notes/Domino Server is installed on the target system and that it has the same non-database files as the Lotus Notes/Domino Server system whose backup is to be restored. These files must be restored first from a filesystem backup.

Problem

Restore of a database fails

During a restore session, some of the selected Lotus/Notes Domino Server databases are not restored, for which Data Protector reports an error similar to the following:

```
[Major] From: OB2BAR@ice.hermes "BLUE"  Time: 8/22/2008 4:07:09 PM
Lotus Notes C API 'NSFTakeDatabaseOffline' returned error 5098:
The database is in use and cannot be taken off-line..
```

Action

1. Disconnect all users that are accessing the databases you want to restore.
2. Restart the restore.

Problem

Recovery of restored Lotus Notes/Domino Server NSF database fails

During recovery, the following error message is displayed:

```
[Critical] From: OB2BAR@ice.hermes "BLUE" Time: 19.10.01
17:24:23

Lotus Notes C API 'NSFGetTransLogStyle' returned error
5114:Recovery Manager: Recovery only supported for Backup
Files.
```

This indicates that at least one database from the restore list was accessed before the recovery ended, either by Lotus Notes/Domino Server, a user, or a process.

Action

1. Restart the Lotus Notes/Domino Server system and perform the restore again.
2. Restore the failed database to a location other than the one it was backed up from.

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.

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

AML *(EMASS/GRAU specific term)* Automated Mixed-Media library.

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

application system *(ZDB specific term)* A system the application or database runs on. The application or database data is located on source volumes.
See also backup system and source volume.

archived redo log *(Oracle specific term)* Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of

an archived redo log is determined by the mode the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also online redo log.

archive logging *(Lotus Domino Server specific term)* Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager (in `Data_Protector_home\Config\Server\dr\asr` on a Windows Cell Manager or in `/etc/opt/omni/server/dr/asr/` on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

Audit Logs Data files to which auditing information is stored.

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

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

autochanger *See* library.

autoloader *See* library.

Automatic Storage Management *(Oracle specific term)* Automatic Storage Management is an Oracle 10g/11g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity

associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.

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

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

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

backup chain *See* restore chain.

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

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

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

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

- Client name: Hostname of the Data Protector client where the backup object resides.
- Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is

located (drive on Windows and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.

- Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration objects — displays the integration type (for example, SAP or Lotus).
- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".

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

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

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

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

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

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

See also application system, target volume, and replica.

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

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

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

BC *(EMC Symmetrix specific term)* Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also BCV.

BC *(HP StorageWorks Disk Array XP specific term)* The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets should be connected to the backup system.
See also HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.

BC EVA *(HP StorageWorks EVA specific term)* Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware.

See also replica, source volume, snapshot, and CA+BC EVA.

BC Process *(EMC Symmetrix specific term)* A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also BCV.

BC VA *(HP StorageWorks Virtual Array specific term)* Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.
See also HP StorageWorks Virtual Array LUN, application system, and backup system.

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also BC and BC Process.

Boolean operators The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

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

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

See also BRBACKUP, and BRRESTORE.

BRBACKUP

(SAP R/3 specific term) An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also BRARCHIVE, and BRRESTORE.

BRRESTORE

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

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

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

BSM

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

CA

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

CA+BC EVA

(HP StorageWorks EVA specific term) The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.
See also BC EVA, replica, and source volume.

CAP
(StorageTek specific term) Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

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

CDB
The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions,, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also MMDB.

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

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

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

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

Centralized Media Management
See CMMDB.

Database (CMMDB)

Change Journal *(Windows specific term)* A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider *(Windows specific term)* A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

channel *(Oracle specific term)* An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type 'disk'
- type 'sbt_tape'

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

chunking *(IAP specific term)* The process of dividing data into blocks (chunks), where each chunk gets a unique content address. This address is then used to determine whether a particular chunk is already backed up to the IAP appliance. If the duplicate data is identified (two addresses are identical, that is the address is the same as for another data chunk already stored into IAP), it is not backed up. This way, the data redundancy is eliminated and the optimal data storage is achieved.
See also backup to IAP.

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

client backup A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or client system

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

cluster-aware application

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

cluster continuous replication

(Microsoft Exchange Server specific term) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also Exchange Replication Service and local continuous replication.

CMD Script for Informix Server

(Informix Server specific term) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB

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

between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Registration Database

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

command-line interface (CLI)

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

Command View (CV) EVA

(HP StorageWorks EVA specific term) The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.
See also HP StorageWorks EVA SMI-S Agent and HP StorageWorks SMI-S EVA provider.

Command View VLS

(VLS specific term) A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.
See also Virtual Library System (VLS).

concurrency

See Disk Agent concurrency.

control file

(Oracle and SAP R/3 specific term) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set

(HP StorageWorks EVA specific term) A pair that consists of the source volumes on a local EVA and their replica on a remote EVA.
See also source volume, replica, and CA+BC EVA

CRS

The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems,

the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

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

data file *(Oracle and SAP R/3 specific term)* A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

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

data stream Sequence of data transferred over the communication channel.

Data_Protector_home On Windows Vista and Windows Server 2008, the directory containing Data Protector program files. On other Windows operating systems, the directory containing Data Protector program files and data files. Its default path is `%ProgramFiles%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_program_data.

Data_Protector_program_data On Windows Vista and Windows Server 2008, the directory containing Data Protector data files. Its default path is `%ProgramData%\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_home.

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

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

Data Replication (DR) group *(HP StorageWorks EVA specific term)* A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided

they have common characteristics and share a common CA EVA log.
See also copy set.

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

Dbobject *(Informix Server specific term)* An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located on the Cell Manager in the directory *`Data_Protector_program_data`*`\db40` (Windows Server 2008), *`Data_Protector_home`*`\db40` (other Windows systems), or `/var/opt/omni/server/db40` (UNIX systems). You can create more DC directories and use a custom location. Up to 50 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

DCBF The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the file system settings.

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

device A physical unit which contains either just a drive or a more complex unit such as a library.

device chain A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)* A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than

a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

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

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

differential backup *(Microsoft SQL Server specific term)* A database backup that records only the data changes made to the database after the last full database backup.
See also backup types.

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

direct backup A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
See also XCopy engine.

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

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

Disk Agent A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

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

disk discovery The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group *(Veritas Volume Manager specific term)* The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing

the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. *See also* virtual full backup.

Distributed File System (DFS) A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller A server in a network that is responsible for user security and verifying passwords within a group of other servers.

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

DR OS A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It

can also read the data from the medium and send it to the computer system.

drive-based encryption

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

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See client backup with disk discovery.

EMC Symmetrix Agent (SYMA) ***(EMC Symmetrix specific term)***

See Symmetrix Agent (SYMA).

emergency boot file

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

enhanced incremental backup

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

Enterprise Backup Environment

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

Event Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the `Admin` group

and to Data Protector users who are granted the `Reporting and notifications` user rights. You can view or delete all events in the Event Log.

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

Exchange Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that represents storage groups that were replicated using either Local Continuous Replication (LCR) or Cluster Continuous Replication (CCR) technology.
See also cluster continuous replication and local continuous replication.

exchanger Also referred to as SCSI Exchanger.
See also library.

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

Extensible Storage Engine (ESE) *(Microsoft Exchange Server specific term)* A database technology used as a storage system for information exchange in Microsoft Exchange Server.

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

failover *(HP StorageWorks EVA specific term)* An operation that reverses the roles of source and destination in CA+BC EVA configurations.
See also CA+BC EVA.

FC bridge *See* Fibre Channel bridge.

Fibre Channel An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed

bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

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

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

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

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

file tree walk *(Windows specific term)* The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

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

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

first-level mirror *(HP StorageWorks Disk Array XP specific term)* HP StorageWorks Disk Array XP allows up to three mirror copies of a primary volume and each of these copies can have additional two copies. The three mirror copies are called first-level mirrors.

See also primary volume and MU number.

flash recovery area

(Oracle specific term) Flash recovery area is an Oracle 10g/11g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and recovery (recovery files).
See also recovery files.

fnames.dat

The `fnames.dat` files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting

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

free pool

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

full backup

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

full database backup

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

full mailbox backup

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

full ZDB

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

global options file

A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located on the Cell Manager in

	the directory *Data_Protector_program_data*\Config\Server\Options (Windows Server 2008), *Data_Protector_home*\Config\Server\Options (other Windows systems), or /etc/opt/omni/server/options (HP-UX or Solaris 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 graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI that runs on Windows, Data Protector also provides a Java-based graphical user interface with the same look and feel, which runs on numerous platforms.
hard recovery	*(Microsoft Exchange Server specific term)* A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.
heartbeat	A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.
Hierarchical Storage Management (HSM)	A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.
Holidays file	A file that contains information about holidays. You can set different holidays by editing the Holidays file on the Cell Manager in the directory *Data_Protector_program_data*\Config\Server\holidays (Windows Server 2008), *Data_Protector_home*\Config\Server\holidays (other Windows systems), or /etc/opt/omni/server/Holidays (UNIX systems).
host backup	*See* client backup with disk discovery.

hosting system

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

HP Operations Manager

HP Operations Manager provides powerful capabilities for operations management of a large number of systems and applications in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for HP Operations Manager management servers on Windows, HP-UX, Solaris, and Linux. Earlier versions of HP Operations Manager were called IT/Operation, Operations Center, Vantage Point Operations, and OpenView Operations.

HP Operations Manager SMART Plug-In (SPI)

A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager, extending the managed domain. Through the Data Protector integration, which is implemented as an HP Operations Manager SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager.

HP StorageWorks Disk Array XP LDEV

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also BC, CA *(HP StorageWorks Disk Array XP specific term)*, and replica.

HP StorageWorks EVA SMI-S Agent

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
See also Command View (CV) EVA and HP StorageWorks SMI-S EVA provider.

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for

information or method invocation, and returns standardized responses.
See also HP StorageWorks EVA SMI-S Agent and Command View (CV) EVA.

HP StorageWorks Virtual Array LUN
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also BC VA and replica.

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

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

IDB recovery file
An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

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

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

incremental backup
(Microsoft Exchange Server specific term) A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.

See also backup types.

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

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

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

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

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

Inet A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

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

Informix Server *(Informix Server specific term)* Refers to Informix Dynamic Server.

initializing *See* formatting.

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

instant recovery *(ZDB specific term)* A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

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

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

IP address An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)* A Sybase utility used to perform system administration tasks on Sybase SQL Server.

Java GUI Client The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities and requires connection to the Java GUI Server to function.

Java GUI Server The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.

jukebox *See* library.

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

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

Key Management Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that provides encryption functionality for enhanced security.
See also Information Store and Site Replication Service.

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

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

LBO *(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library

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

lights-out operation or unattended operation

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

LISTENER.ORA

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

load balancing

By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

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

local continuous replication

(Microsoft Exchange Server specific term) Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.

An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.
See also cluster continuous replication and Exchange Replication Service.

lock name

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

log_full shell script

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

logging level

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

logical-log files

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

login ID

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

login information to the Oracle Target Database

(Oracle and SAP R/3 specific term) The format of the login information is *user_name*/*password*@*service*, where:

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

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *user_name*/*password*@*service*, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API

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

LVM

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

Magic Packet

See Wake ONLAN.

mailbox

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

mailbox store *(Microsoft Exchange Server specific term)* A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)* An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. *See also* BC (HP StorageWorks Disk Array XP specific term), CA *(HP StorageWorks Disk Array XP specific term)*, and HP StorageWorks Disk Array XP LDEV.

Manager-of-Managers (MoM) *See* MoM.

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

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

MAPI *(Microsoft Exchange Server specific term)* The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU *See* Main Control Unit (MCU).

Media Agent A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium.

During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

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

media condition The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

medium ID A unique identifier assigned to a medium by Data Protector.

media label A user-defined identifier used to describe a medium.

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

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

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

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

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

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

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

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

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

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

Microsoft Volume Shadow Copy Service (VSS) A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also shadow copy, shadow copy provider, replica, and writer.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* *See* target volume.

mirror rotation *(HP StorageWorks Disk Array XP specific term)* *See* replica set rotation.

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

MMDB The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also CMMDB, CDB.

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

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

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

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

MU number *(HP StorageWorks Disk Array XP specific term)* Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror.
See also first-level mirror.

multi-drive server A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat *See* IDB recovery file.

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

object *See* backup object.

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

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

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

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

object copying The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

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

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

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

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

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup

period (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.

- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

See also zero downtime backup (ZDB) and online backup.

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

offline redo log *See* archived redo log.

ON-Bar *(Informix Server specific term)* A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- the `onbar` command
- Data Protector as the backup solution
- the XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

ONCONFIG *(Informix Server specific term)* An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file in the directory *`INFORMIXDIR`*`\etc` (on Windows) or *`INFORMIXDIR`*`/etc/` (on UNIX).

online backup A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished.

- For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
See also zero downtime backup (ZDB), and offline backup.

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

OpenSSH A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard *(Oracle specific term)* Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance *(Oracle specific term)* Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)* A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *ORACLE_SID*. The *ORACLE_SID* is included in the CONNECT DATA parts of the connect descriptor in a `TNSNAMES.ORA` file and in the definition of the TNS listener in the `LISTENER.ORA` file.

original system The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also merging.

ownership

Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.
If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.
If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.
If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

P1S file

P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into `Data_Protector_home\Config\Se ver\dr\p1s` directory on a Windows Cell Manager or in `/etc/opt/omni/server/dr/p1s` directory on a UNIX Cell Manager with the filename `recovery.p1s`.

package

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

pair status

(HP StorageWorks Disk Array XP specific term) A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

parallel restore Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

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

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

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

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

prealloc list A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec

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

primary volume (P-VOL)

(HP StorageWorks Disk Array XP specific term) Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and also catalog protection.

public folder store

(Microsoft Exchange Server specific term) The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID

Redundant Array of Inexpensive Disks.

RAID Manager Library

(HP StorageWorks Disk Array XP specific term) The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP

(HP StorageWorks Disk Array XP specific term) The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This

instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup *See* disk image backup.

RCU *See* Remote Control Unit (RCU).

RDBMS Relational Database Management System.

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

RDS The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)* A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts

Recovery Catalog Database *(Oracle specific term)* An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files *(Oracle specific term)* Recovery files are Oracle 10g/11g specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. *See also* flash recovery area.

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

Recovery Manager (RMAN)

(Oracle specific term) An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle

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

redo log

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

Remote Control Unit (RCU)

(HP StorageWorks Disk Array XP specific term) The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database

(Windows specific term) A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point

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

replica

(ZDB specific term) An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects

is replicated. However, if a volume manager is used on UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.
See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also replica and replica set rotation.

replica set rotation *(ZDB specific term)* The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also replica and replica set.

restore chain All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.

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

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

RMAN *(Oracle specific term)* *See* Recovery Manager.

RSM The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM

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

scan

A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning

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

Scheduler

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

secondary volume (S-VOL)

(HP StorageWorks Disk Array XP specific term) secondary volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. *See also* primary volume (P-VOL) and Main Control Unit (MCU)

session

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

session ID

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

session key

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

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

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

shadow copy set *(Microsoft VSS specific term)* A collection of shadow copies created at the same point in time.
See also shadow copy and replica set.

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

SIBF The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

single instancing *(IAP specific term)* The process of recognizing redundancy of data, at both a whole object and a chunk level. It computes a strong hash for each data chunk and uses it as a unique content address needed to determine whether attempts to store duplicates are being made.
See also backup to IAP.

Site Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also Information Store and Key Management Service.

slot A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a

number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB *See* split mirror backup.

smart copy *(VLS specific term)* A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management.
See also Virtual Library System (VLS).

smart copy pool *(VLS specific term)* A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library.
See also Virtual Library System (VLS) and smart copy.

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

snapshot *(HP StorageWorks VA and HP StorageWorks EVA specific term)* A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also replica and snapshot creation.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)* *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation *(HP StorageWorks VA and HP StorageWorks EVA specific term)* A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for

use. However background copying processes normally continue after creation.
See also snapshot.

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

source volume *(ZDB specific term)* A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents of the source volumes.
See also replica and split mirror creation.

split mirror backup *(EMC Symmetrix specific term)* *See* ZDB to tape.

split mirror backup *(HP StorageWorks Disk Array XP specific term)* *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also split mirror.

split mirror restore *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.
See also ZDB to tape, ZDB to disk+tape, and replica.

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

SRD file The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

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

SSE Agent *(HP StorageWorks Disk Array XP specific term)* A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

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

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

for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

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

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

Storage Group *(Microsoft Exchange Server specific term)* A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

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

storage volume *(ZDB specific term)* A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover *See* failover.

Sybase Backup Server API *(Sybase specific term)* An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

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

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)* The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

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

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

System Backup to Tape *(Oracle specific term)* An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)* The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

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

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

SysVol *(Windows specific term)* A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

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

tapeless backup *(ZDB specific term)* *See* ZDB to disk.

target database *(Oracle specific term)* In RMAN, the target database is the database that you are backing up or restoring.

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

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

target volume *(ZDB specific term)* A storage volume to which data is replicated.

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

thread *(Microsoft SQL Server specific term)* An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)* A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

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

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

transaction backup *(Sybase and SQL specific term)* A backup of the transaction log providing a record of changes made since the last full or transaction backup.

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

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

transaction logs *(Data Protector specific term)* Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table *(Sybase specific term)* A system table in which all changes to the database are automatically recorded.

transportable snapshot *(Microsoft VSS specific term)* A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed.
See also Microsoft Volume Shadow Copy Service (VSS).

TSANDS.CFG file *(Novell NetWare specific term)* A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the `SYS:SYSTEM\TSA` directory on the server where `TSANDS.NLM` is loaded.

UIProxy The Java GUI Server (`UIProxy` service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.

unattended operation *See* lights-out operation.

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

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

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

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

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

user rights User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical

user rights. Users have the access rights of the user group to which they belong.

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

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

Virtual Controller Software (VCS) *(HP StorageWorks EVA specific term)* The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
See also Command View (CV) EVA.

Virtual Device Interface *(Microsoft SQL Server specific term)* This is a SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)* A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also source volume and target volume.

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

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

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

virtual tape

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

Virtual Tape Library (VTL)

(VLS specific term) An emulated tape library that provides the functionality of traditional tape-based storage.
See also Virtual Library System (VLS).

VMware management client

(VMware integration specific term) The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).

volser

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

volume group

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

volume mount point

(Windows specific term) An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy Service

See Microsoft Volume Shadow Copy Service.

VSS

See Microsoft Volume Shadow Copy Service.

VSS compliant mode

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

See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS Veritas Journal Filesystem.

VxVM (Veritas Volume Manager) A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

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

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

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

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

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

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

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

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

XCopy engine *(direct backup specific term)* A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
See also direct backup.

ZDB *See* zero downtime backup (ZDB).

ZDB database *(ZDB specific term)* A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.
See also zero downtime backup (ZDB).

ZDB to disk *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape

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

zero downtime backup (ZDB)

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

Index

C

D

E

F

H

I

L

M

N

O

P

Q

R

S

T

W