

HP Data Protector

Software Version: 8.10

Integration guide for IBM Applications

Document Release Date: November 2016

Software Release Date: November 2016



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

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

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

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

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

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

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

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

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

Support

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

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

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

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

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

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

To find more information about access levels, go to:

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

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

Contents

Contents	3
Chapter 1: Data Protector Informix Server integration	9
Introduction	9
Integration concepts	9
Configuring the integration	11
Prerequisites	11
Before you begin	11
Cluster-aware clients	12
Configuring Informix Server users	12
Configuring Informix instances	12
Before you begin	12
Using the Data Protector GUI	12
Using the Data Protector CLI	16
Handling errors	17
Checking the configuration	17
Using the Data Protector GUI	17
Using the Data Protector CLI	17
Backup	18
What you must back up as filesystem	18
What does not need to be backed up?	19
Creating backup specifications	19
Modifying backup specifications	24
Scheduling backup sessions	25
Scheduling example	25
Previewing backup sessions	26
Using the Data Protector GUI	26
Using the Data Protector CLI	26
What happens during the preview?	27
Starting backup sessions	28
Backup methods	28

Before you begin	28
Using the Data Protector GUI	28
Using the Data Protector CLI	29
Using Informix Server commands	29
Using Informix Server log_full.sh on UNIX	30
Manual and continuous logical log backups	31
Restore	31
Restore methods	32
Before you begin	32
Finding information for restore	32
Using the Data Protector GUI	32
Using the Data Protector CLI	33
Restoring using the Data Protector GUI	35
Restoring using the Data Protector CLI	38
Restoring using Informix Server commands	39
Restoring dbspaces, blobspaces, and logical logs	39
Restoring dbspaces and blobspaces only	40
Restoring a particular dbspace or blobspace	40
Restoring to another Informix Server	40
Restoring using another device	41
Using the Data Protector GUI	41
Using the Data Protector CLI or Informix Server commands	41
Monitoring sessions	42
Troubleshooting	42
Before you begin	42
Checks and verifications	42
Checking the Informix Server side	46
Problems	47
Chapter 2: Data Protector DB2 UDB integration	48
Introduction	48
Integration concept	49
Configuring the integration	50

Prerequisites	50
Before you begin	50
Partitioned environment	50
Configuring DB2 users	50
Configuring DB2 instances	51
Before you begin	51
Using the Data Protector GUI	51
Using the Data Protector CLI	52
Checking the configuration	53
Using the Data Protector GUI	53
Using the Data Protector CLI	53
Backup	53
Physically partitioned environment	54
Creating backup specifications	54
Modifying backup specifications	58
Scheduling backup sessions	58
Previewing backup sessions	59
Using the Data Protector GUI	59
Using the Data Protector CLI	59
What happens during the preview?	60
Starting backup sessions	60
Before you begin	60
Using the Data Protector GUI	61
Using the Data Protector CLI	61
Starting backups of physically partitioned DB2 objects	61
Restore	61
Restoring using the Data Protector GUI	62
Restoring using the Data Protector CLI	66
Restoring to a new database or another DB2 instance	67
Restore in a partitioned environment	70
Restoring to the original database	71
Corrupt database	71

Physically partitioned environment	71
Logically partitioned environment	71
Restoring to a new database or another instance	72
Monitoring sessions	72
Troubleshooting	73
Before you begin	73
Checks and verifications	73
Problems	74
Chapter 3: Data Protector Lotus Notes/Domino Server integration	78
Introduction	78
Integration concepts	79
Lotus Domino Cluster	80
Replicas	80
Replication in a cluster	81
Failover in a cluster	81
Example	81
Configuring the integration	82
Prerequisites	83
Before you begin	83
Transaction logging of Lotus Notes/Domino Server	84
Enabling transaction logging	84
Configuring Lotus Notes/Domino Server users	85
Configuring Lotus Notes/Domino Server systems	86
Using the Data Protector GUI	86
Using the Data Protector CLI	87
Checking the configuration	88
Using the Data Protector GUI	88
Using the Data Protector CLI	89
Handling errors	89
Backup	89
What is backed up?	90
What is not backed up?	90

Considerations	91
Creating backup specifications	91
Modifying backup specifications	94
Scheduling backup sessions	94
Scheduling example	94
Previewing backup sessions	95
Using the Data Protector GUI	95
Using the Data Protector CLI	96
What happens during the preview?	96
Starting backup sessions	96
Using the Data Protector GUI	96
Restore	96
Finding information for restore	97
Using the Data Protector GUI	97
Using the Data Protector CLI	98
Restoring using the Data Protector GUI	99
Restoring using the Data Protector CLI	101
Restore options	102
Restore in Lotus Domino Cluster environment	103
Restoring a replica database without recovery	104
Restoring with recovery to the latest possible state	104
Point-in-time recovery	104
Restoring to a new location	105
Performance tuning	105
Monitoring sessions	105
Troubleshooting	106
Before you begin	106
Checking the Lotus Notes/Domino Server side	106
Checks and verifications	106
Problems	109
Glossary	112
Index	153

We appreciate your feedback! 159

Chapter 1: Data Protector Informix Server integration

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 1: 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 2: 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.

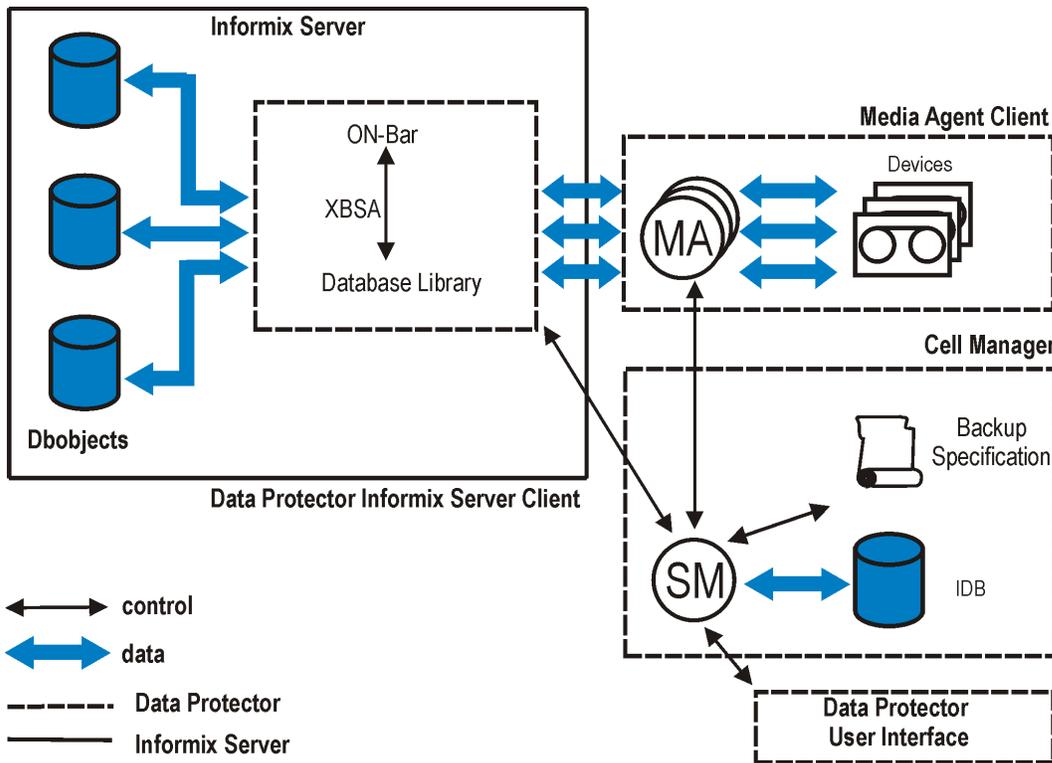
You can also back up and restore dbobjects using the Informix Server onbar command.

This chapter provides information specific to the Data Protector Informix Server integration. For general Data Protector procedures and options, see the *HP Data Protector 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**. [Data Protector Informix Server integration architecture on the next page](#) shows the architecture of the Data Protector Informix Server integration.

Figure 1: 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://support.openview.hp.com/selfsolve/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 systems:**
 - On Windows Server 2003 system, you need to restart the Data Protector Inet service under a Windows domain user account that has the appropriate Informix Server permissions for running backups and restores. Stop the service and restart it as user `informix`.

For information on changing the user account under which the Data Protector Inet service is running, see the *HP Data Protector Help* index: "Inet, changing account".

- On other Windows operating systems, configure the Data Protector Inet service user impersonation for the user that has the appropriate Informix Server permissions for running backups and restores.

For details, see the *HP Data Protector Help* index: "Inet user impersonation".

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 systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `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 *HP Data Protector 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 systems:** Name of the system with the `sqlhosts` entry in the Windows Registry.
- **UNIX systems:** 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.

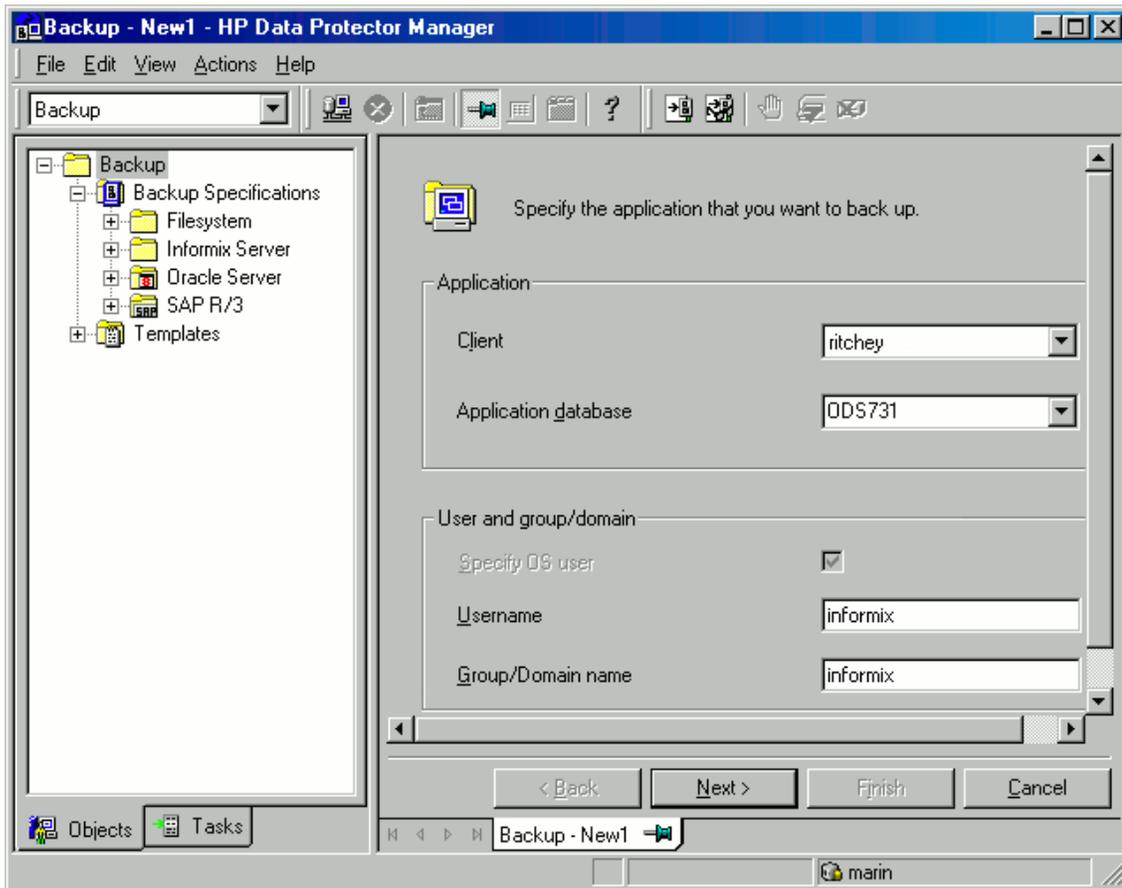
In **Application database**, enter the Informix instance name.

In the **User and group/domain** options, specify the account under which you want the backup session to run. These options are available on UNIX and Windows Server 2008 clients. On Windows Server 2003, the backup session will run under the account under which the Data Protector Inet service is running.

Ensure that this user has been added to the Data Protector admin or operator user group, and has the Informix Server backup rights. This user becomes the backup owner.

- **UNIX systems:** Type `informix` in both **Username** and **Group/Domain name**.
- **Windows Server 2008:** In **Username** and **Group/Domain name**, type the user name and domain (for example, the user name Administrator, domain DP). This account must be set up for the Data Protector Inet service user impersonation. For details, see the *HP Data Protector Help* index: "Inet user impersonation".

Figure 2: 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 systems: Name of the system with the sqlhosts entry in the Windows Registry. Use the UNC notation, for example: \\computer_name.

UNIX systems: 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 systems: *INFORMIXDIR\etc*

UNIX systems: *INFORMIXDIR/etc*

Figure 3: Configuring an Informix instance (Windows)

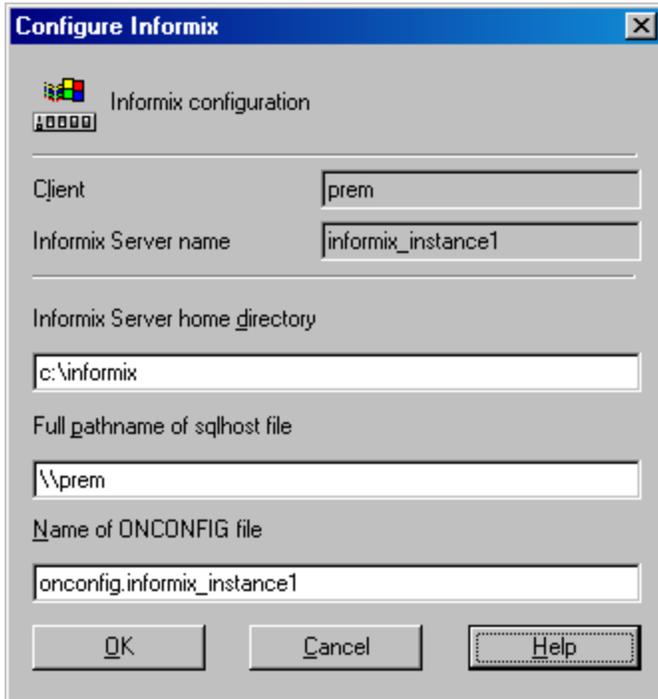
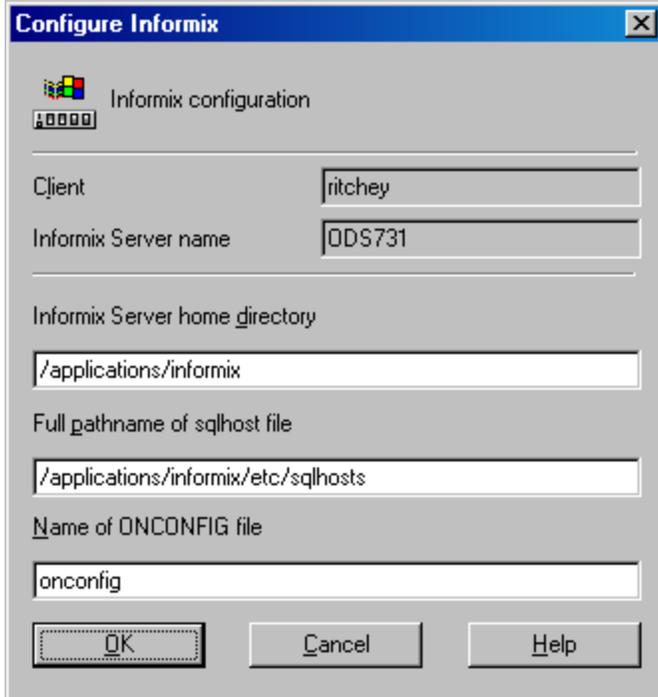


Figure 4: Configuring an Informix instance (UNIX)



Click **OK**.

6. If an error occurs, click **Details** or see [Troubleshooting on page 42](#).
7. The Informix instance is configured. Exit the GUI or proceed with creating the backup specification at [Select the dobjects to be backed up. on page 20](#).

Using the Data Protector CLI

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

Windows systems: `Data_Protector_home\bin`

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

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

execute the following:

Windows systems:

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

UNIX systems:

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

Parameter description

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

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

Figure 5: Configuring an Informix instance (Windows)

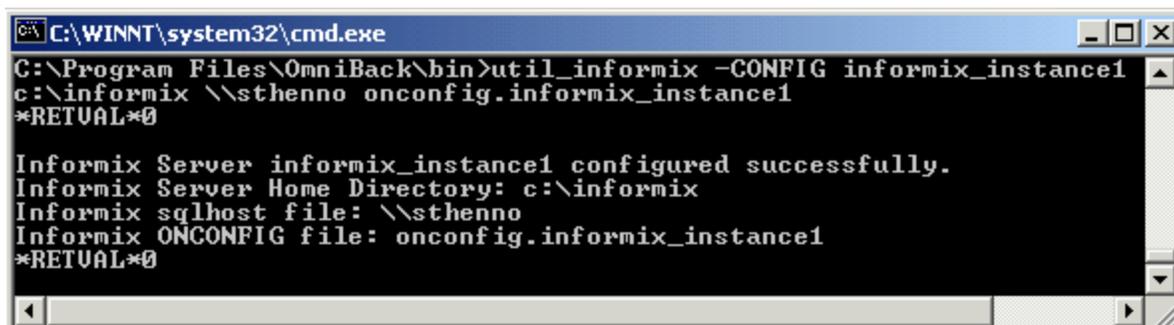
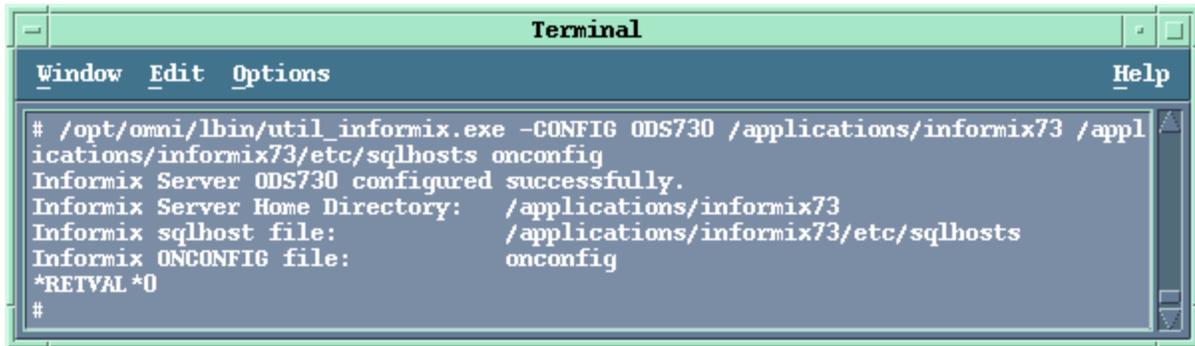


Figure 6: Configuring an Informix instance (HP-UX, Solaris)



```
Terminal
Window Edit Options Help
# /opt/omni/lbin/util_informix.exe -CONFIG ODS730 /applications/informix73 /appl
ications/informix73/etc/sqlhosts onconfig
Informix Server ODS730 configured successfully.
Informix Server Home Directory: /applications/informix73
Informix sqlhost file: /applications/informix73/etc/sqlhosts
Informix ONCONFIG file: onconfig
*RETVAL *0
#
```

Handling errors

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

To get the error description:

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

HP-UX and Solaris systems: Execute:

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

Other UNIX systems: Execute:

```
/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 systems: `Data_Protector_home\bin`

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

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

execute:

Windows systems:

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

UNIX systems:

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

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

Figure 7: Checking configuration (Windows)

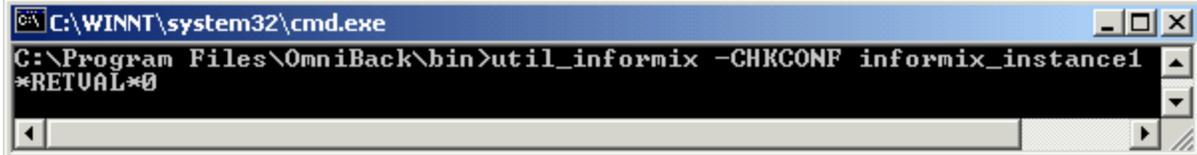


Figure 8: Checking configuration (UNIX)



A successful configuration check displays the message **RETVL*0*.

If an error occurs, the error number is displayed in the form **RETVL*error_number*. For information on how to get the error description, see [Handling errors on the previous page](#).

Backup

The integration provides online database backup 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.

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 4: What needs to be backed up as filesystem

Object	Location
The ONCONFIG file	Windows systems: <i>INFORMIXDIR \etc</i> UNIX systems: <i>INFORMIXDIR /etc</i>
The oncfg_ <i>SERVERNAME</i> . <i>SERVERNUM</i> file	
Emergency boot file, an Informix Server configuration file called <i>ixbar.server_id</i> , where <i>server_id</i> is the value of the <i>SERVERNUM</i> configuration parameter.	
UNIX systems: The <i>sqlhosts</i> 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.

In the **User and group/domain** options, specify the account under which you want the backup session to run. These options are available on UNIX and Windows Server 2008 clients. On Windows Server 2003, the backup session will run under the account under which the Data Protector Inet service is running.

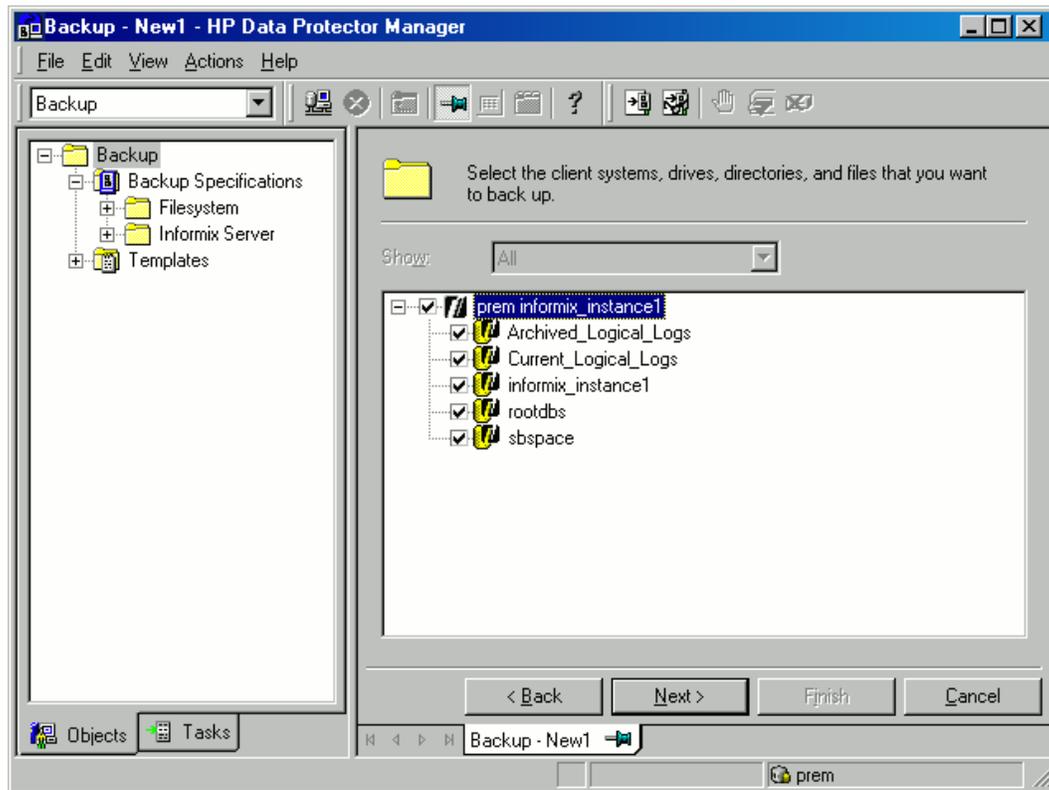
Ensure that this user has been added to the Data Protector admin or operator user group, and has the Informix Server backup rights. This user becomes the backup owner.

- **UNIX systems:** Type `informix` in both **Username** and **Group/Domain name**.
- **Windows Server 2008:** In **Username** and **Group/Domain name**, type the user name and domain (for example, the user name Administrator, domain DP). This account must be set up for the Data Protector Inet service user impersonation. For details, see the *HP Data Protector Help* index: "Inet user impersonation".

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 12](#).
6. Select the dbjects to be backed up.

Figure 9: 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 [Specifying Informix Server resource types below](#).

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 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 10: Specifying Informix Server resource types

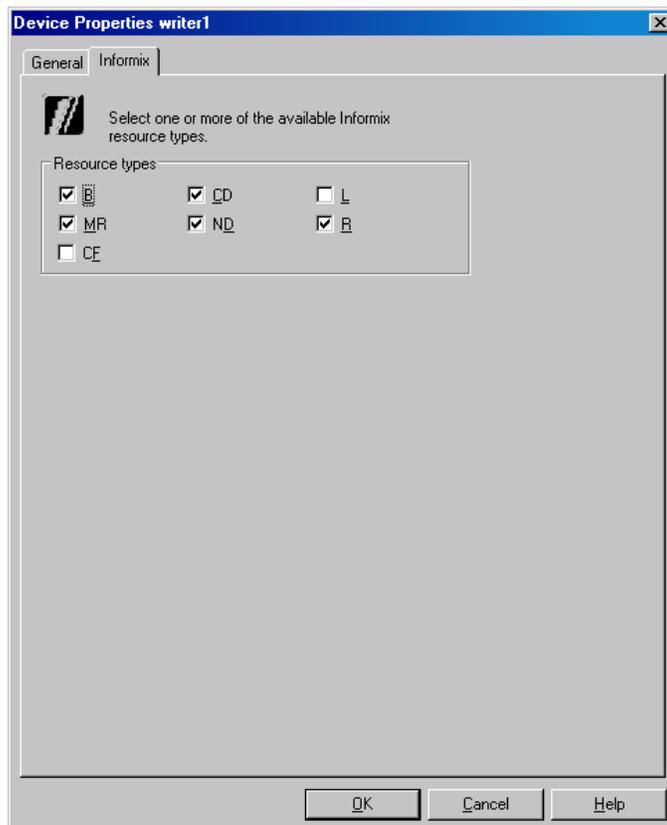


Table 5: Informix Server resource types

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
CF	Critical file

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 ([Informix Server specific backup options \(Windows\)](#) below and [Informix Server specific backup options \(UNIX\)](#) on the next page). For information, see [Informix Server backup options](#) on page 24.

Figure 11: Informix Server specific backup options (Windows)

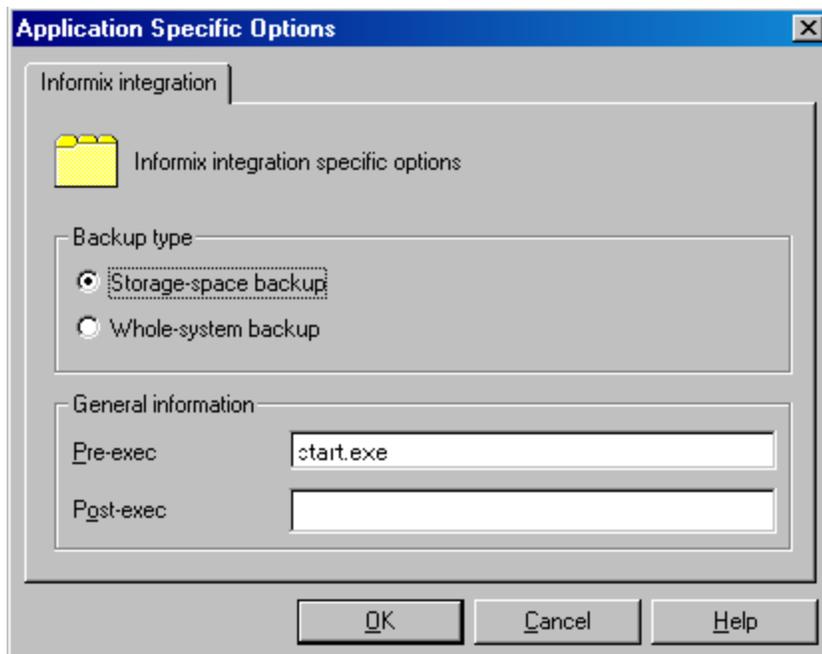
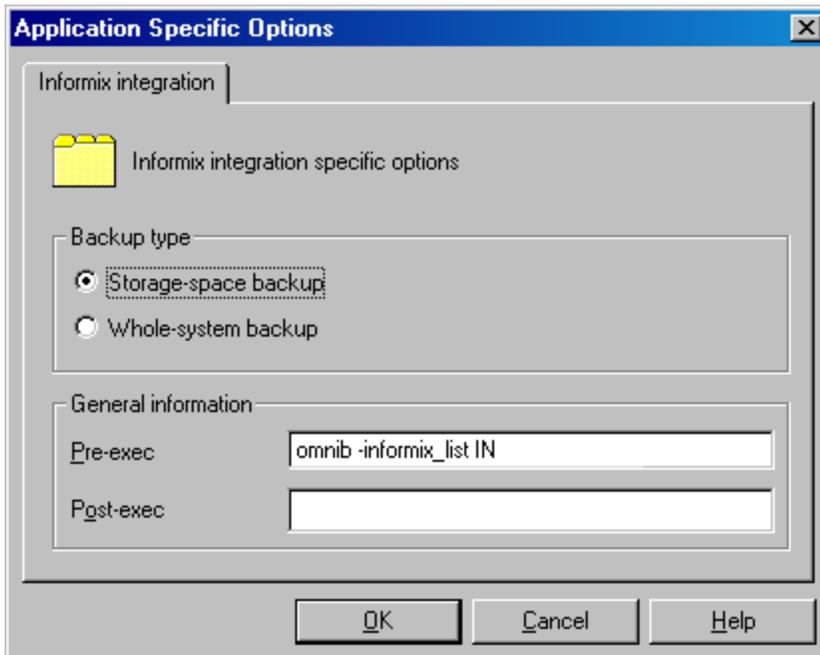


Figure 12: Informix Server specific backup options (UNIX)



Click **Next**.

9. Optionally, schedule the backup. See [Scheduling backup sessions on page 25](#).

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 26](#).

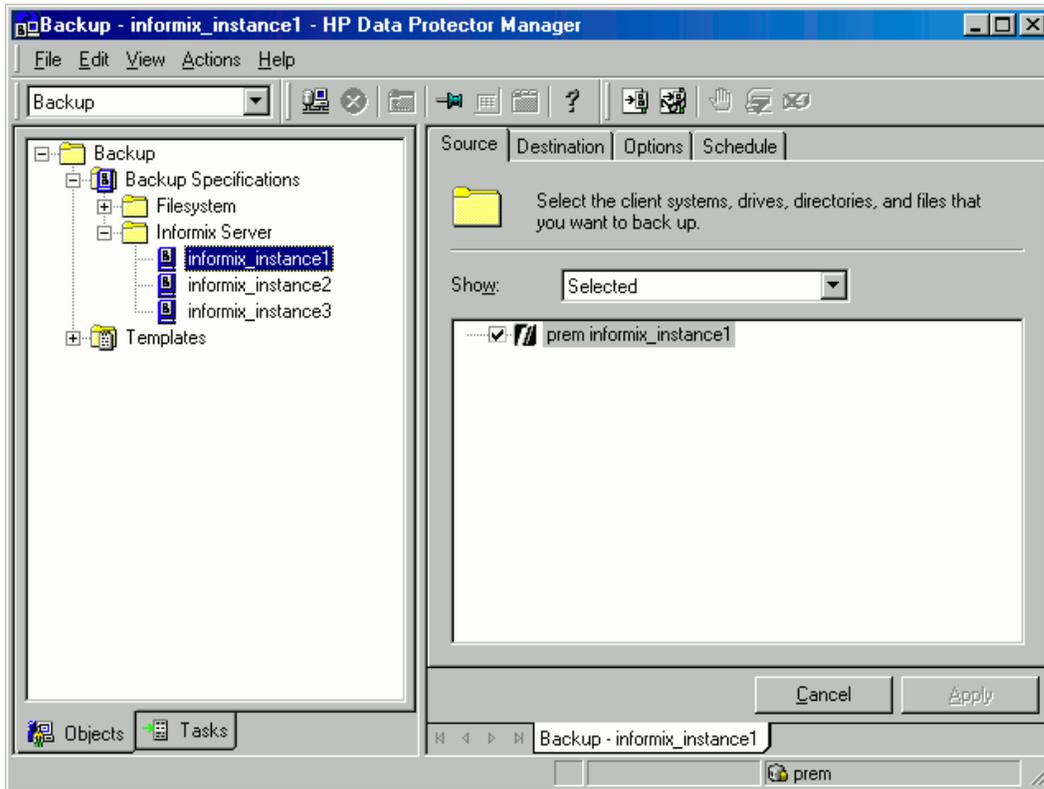
Table 6: Informix Server backup options

Option		Description
Backup type	Storage-space backup (default)	In a storage-space backup, the onbar command backs up the selected storage-spaces and logical logs in parallel. When you restore from a storage-space backup, you also have to restore logical logs to make the data consistent. Storage-space backup is faster than whole-system backup on large databases.
	Whole-system backup	In a whole-system backup, all Informix instance's dobjects 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 disaster recovery, or restore to another client. When you restore from a whole-system backup, you do not need to restore logical logs to make the data consistent.
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, spaces, or special characters. Provide only the name of the command, which must reside in the following directory: Windows systems: <i>Data_Protector_home\bin</i> See Informix Server specific backup options (Windows) on page 22 . HP-UX, Solaris, Linux systems: <i>/opt/omni/lbin</i> See Informix Server specific backup options (UNIX) on the previous page . Other UNIX systems: <i>/usr/omni/bin</i> 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. If the onmode -l command returns a non-zero value, Data Protector interprets this as an error and the backup session does not start.

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 [Modifying a backup specification on the next page](#).

Figure 13: Modifying a backup specification



Scheduling backup sessions

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

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 [Scheduling a backup session on the next page](#).

Click **OK**.

3. Repeat steps 1 and 2 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

Figure 14: Scheduling a backup session

Schedule Backup

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

Recurring

None
 Daily
 Weekly
 Monthly

Time options

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

Recurring options

Every 1 week(s) on

Sun Mon Tue Wed Thu Fri Sat

Session options

Backup type: Full
Network load: High Medium Low
Backup protection: Default

OK Cancel Help

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**.

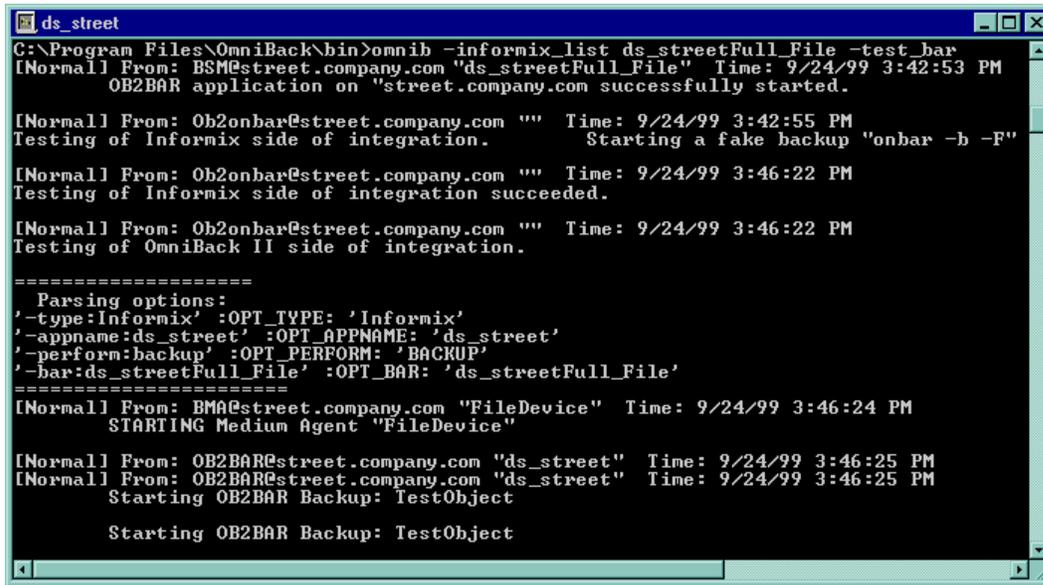
The message `Session completed successfully` is displayed at the end of a successful preview.

Using the Data Protector CLI

Execute the following commands:

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

Figure 15: Previewing a backup with backup specification ds_street (Windows)



```
C:\Program Files\OmniBack\bin>omnib -informix_list ds_streetFull_File -test_bar
[Normal] From: BSM@street.company.com "ds_streetFull_File" Time: 9/24/99 3:42:53 PM
OB2BAR application on "street.company.com" successfully started.

[Normal] From: Ob2onbar@street.company.com "" Time: 9/24/99 3:42:55 PM
Testing of Informix side of integration. Starting a fake backup "onbar -b -F"

[Normal] From: Ob2onbar@street.company.com "" Time: 9/24/99 3:46:22 PM
Testing of Informix side of integration succeeded.

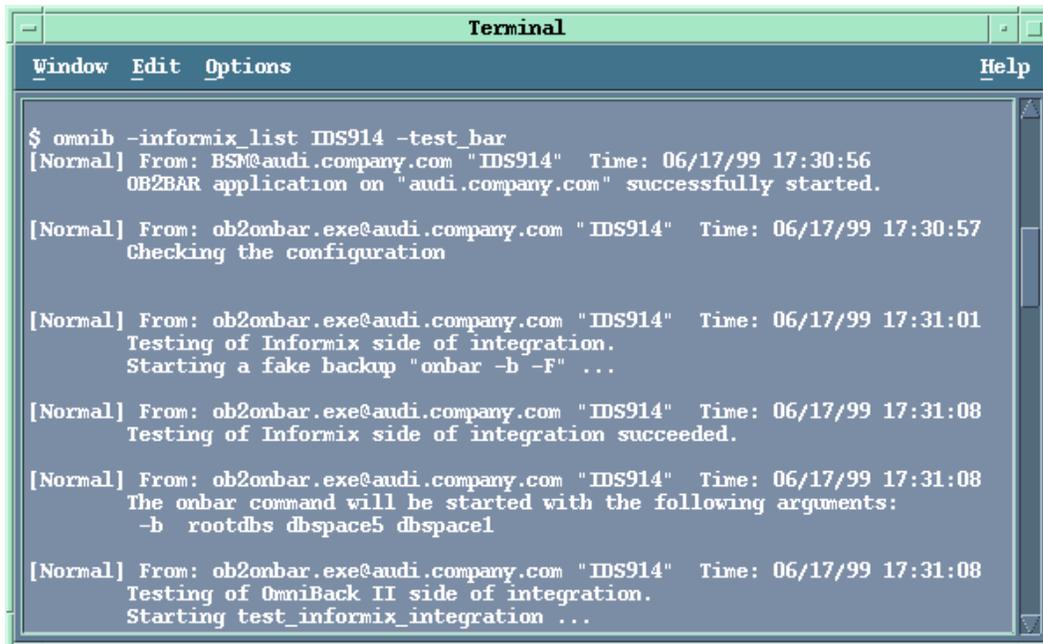
[Normal] From: Ob2onbar@street.company.com "" Time: 9/24/99 3:46:22 PM
Testing of OmniBack II side of integration.

=====
Parsing options:
'-type:Informix' :OPT_TYPE: 'Informix'
'-appname:ds_street' :OPT_APPNAME: 'ds_street'
'-perform:backup' :OPT_PERFORM: 'BACKUP'
'-bar:ds_streetFull_File' :OPT_BAR: 'ds_streetFull_File'
=====
[Normal] From: BMA@street.company.com "FileDevice" Time: 9/24/99 3:46:24 PM
STARTING Medium Agent "FileDevice"

[Normal] From: OB2BAR@street.company.com "ds_street" Time: 9/24/99 3:46:25 PM
[Normal] From: OB2BAR@street.company.com "ds_street" Time: 9/24/99 3:46:25 PM
Starting OB2BAR Backup: TestObject

Starting OB2BAR Backup: TestObject
```

Figure 16: Previewing a backup with backup specification IDS914 (UNIX)



```
Terminal
Window Edit Options Help

$ omnib -informix_list IDS914 -test_bar
[Normal] From: BSM@audi.company.com "IDS914" Time: 06/17/99 17:30:56
OB2BAR application on "audi.company.com" successfully started.

[Normal] From: ob2onbar.exe@audi.company.com "IDS914" Time: 06/17/99 17:30:57
Checking the configuration

[Normal] From: ob2onbar.exe@audi.company.com "IDS914" Time: 06/17/99 17:31:01
Testing of Informix side of integration.
Starting a fake backup "onbar -b -F" ...

[Normal] From: ob2onbar.exe@audi.company.com "IDS914" Time: 06/17/99 17:31:08
Testing of Informix side of integration succeeded.

[Normal] From: ob2onbar.exe@audi.company.com "IDS914" Time: 06/17/99 17:31:08
The onbar command will be started with the following arguments:
-b rootdbs dbspace5 dbspace1

[Normal] From: ob2onbar.exe@audi.company.com "IDS914" Time: 06/17/99 17:31:08
Testing of OmniBack II side of integration.
Starting test_informix_integration ...
```

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 are 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 dobjects in any of the following ways:

- Use the Data Protector GUI. See [Using the Data Protector GUI below](#).
- Use the Data Protector CLI. See [Using the Data Protector CLI on the next page](#).
- Use the Informix Server onbar command. See [Using Informix Server commands on the next page](#).
- **UNIX systems:** Use the Informix Server `log_full.sh` script. See [Using Informix Server log_full.sh on UNIX on page 30](#).

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 use 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

Execute the following command:

```
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`, execute:

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

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

```
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 7: Data Protector and Informix Server variables

ONCONFIG	Name of the Informix instance ONCONFIG file.
INFORMIXSQLHOSTS	Windows systems: System on which the <code>sqlhosts</code> entry in the Windows Registry exists. UNIX systems: Pathname of the <code>sqlhosts</code> file, for example <code>/applications/informix/etc/sqlhosts</code> .
INFORMIXDIR	Pathname of the Informix Server home directory.

INFORMIXSERVER	Name of the Informix instance.
OB2APPNAME	Name of the Informix instance.
OB2BARLIST	For <i>backup</i> , name of the backup specification to be used for the backup. For <i>restore</i> , 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 dobjects, execute:

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

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

Table 8: 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 dobjects.

To back up a list of dbspaces, execute:

```
onbar -b dbspace_list
```

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

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

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 the next page](#).

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 [Data Protector and Informix Server variables on page 29](#) and execute:

```
onbar -l
```

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 9: Informix Server restore types

Complete database restore	Restore from any backup. ON-Bar restores dobjects 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 dobjects in any of the following ways:

- Use the Data Protector GUI. See [Restoring using the Data Protector GUI on page 35](#).
- Use the Data Protector CLI. See [Restoring using the Data Protector CLI on page 38](#).
- Use the Informix Server onbar command. See [Restoring using Informix Server commands on page 39](#).

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 execute:

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

UNIX systems: `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, execute:

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

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

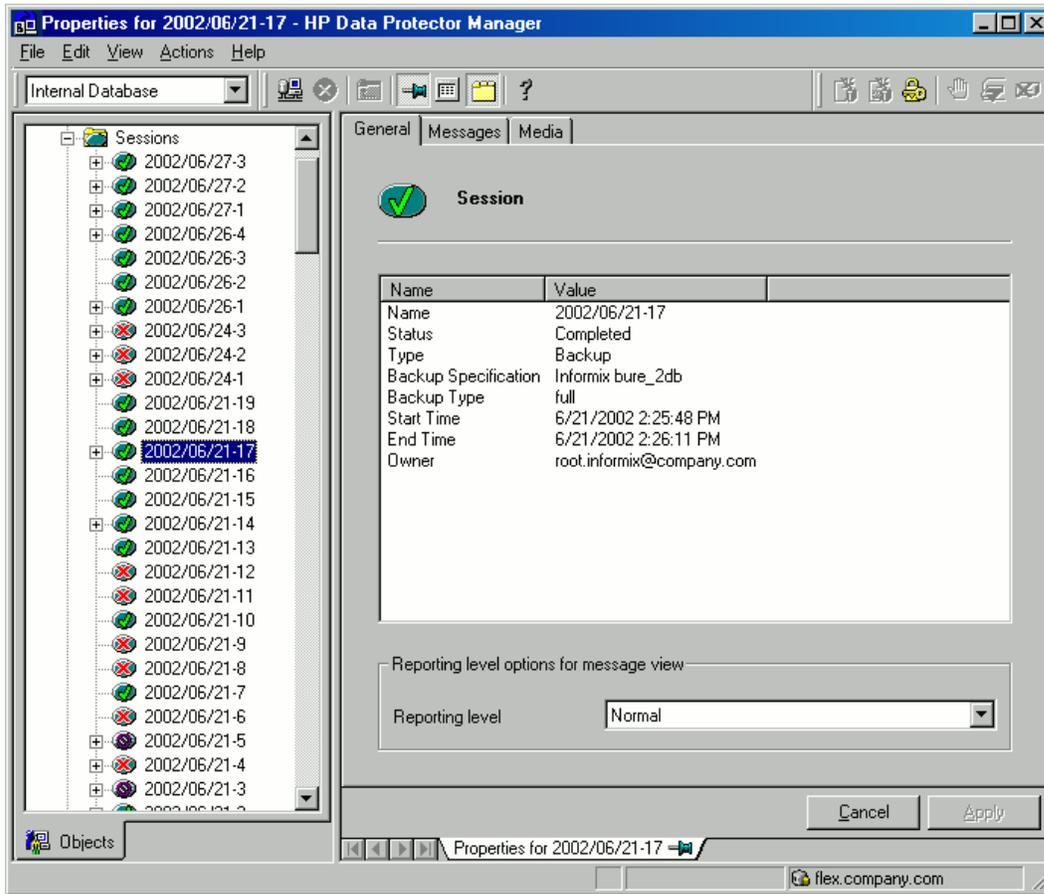
Finding information for restore

To restore dobjects, 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 17: 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:

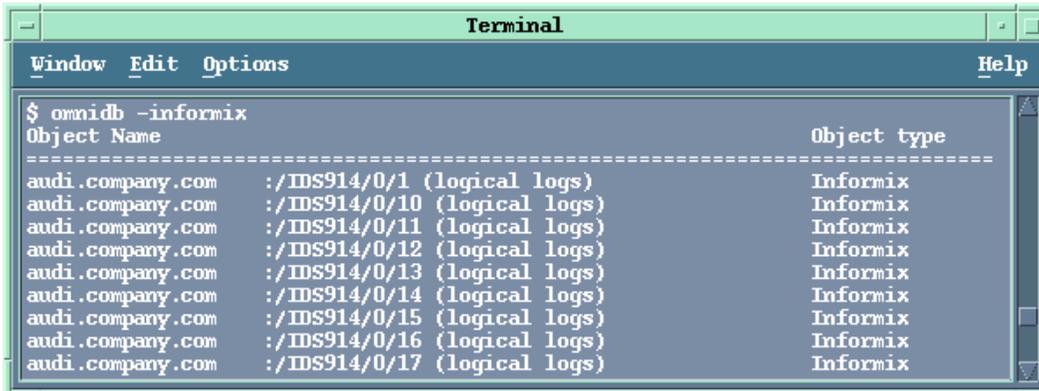
- Set the environment variable `OB2_CLI_UTF8` to 1.
- Set the encoding used on the terminal to UTF-8.

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

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

```
omnidb -informix
```

Figure 18: Example of a list of Informix Server backed up objects

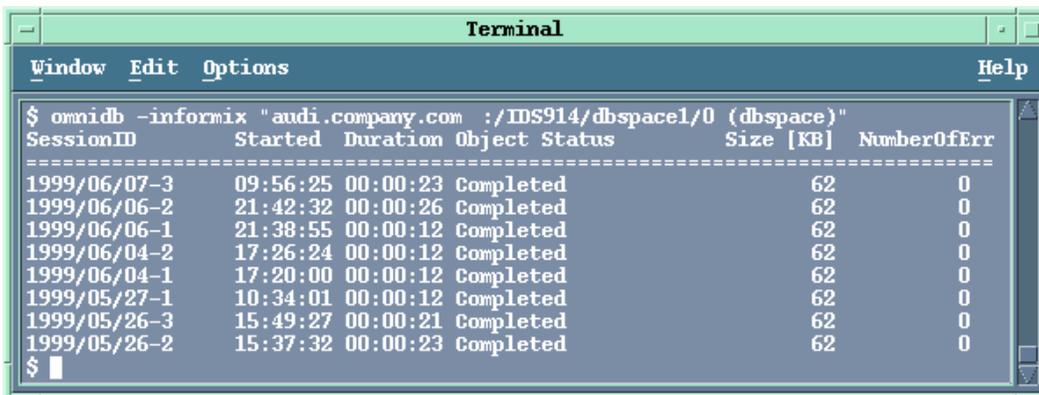


```
Terminal
Window Edit Options Help
$ omnidb -informix
Object Name                                     Object type
-----
audi.company.com :/IDS914/0/1 (logical logs)    Informix
audi.company.com :/IDS914/0/10 (logical logs)   Informix
audi.company.com :/IDS914/0/11 (logical logs)   Informix
audi.company.com :/IDS914/0/12 (logical logs)   Informix
audi.company.com :/IDS914/0/13 (logical logs)   Informix
audi.company.com :/IDS914/0/14 (logical logs)   Informix
audi.company.com :/IDS914/0/15 (logical logs)   Informix
audi.company.com :/IDS914/0/16 (logical logs)   Informix
audi.company.com :/IDS914/0/17 (logical logs)   Informix
```

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

```
omnidb -informix object_name
```

Figure 19: Example of a list of backup sessions for a specific object



```
Terminal
Window Edit Options Help
$ omnidb -informix "audi.company.com :/IDS914/dbspace1/0 (dbspace)"
SessionID      Started  Duration Object Status      Size [KB]  NumberOfErr
-----
1999/06/07-3   09:56:25 00:00:23 Completed 62         0
1999/06/06-2   21:42:32 00:00:26 Completed 62         0
1999/06/06-1   21:38:55 00:00:12 Completed 62         0
1999/06/04-2   17:26:24 00:00:12 Completed 62         0
1999/06/04-1   17:20:00 00:00:12 Completed 62         0
1999/05/27-1   10:34:01 00:00:12 Completed 62         0
1999/05/26-3   15:49:27 00:00:21 Completed 62         0
1999/05/26-2   15:37:32 00:00:23 Completed 62         0
$
```

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

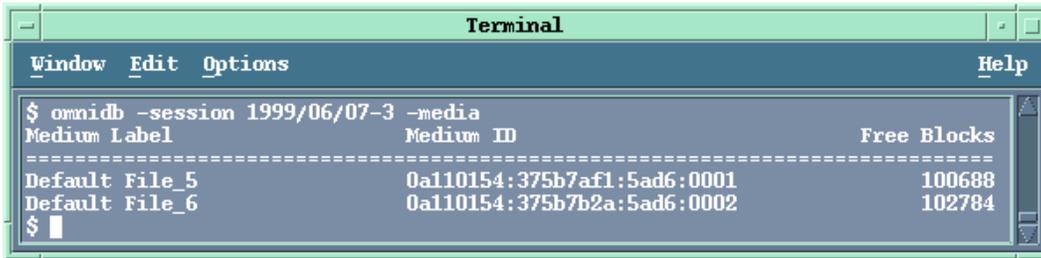
To get information on the object backup ID, execute:

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

3. Get a list of media needed for restore:

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

Figure 20: Example of finding media needed for restore

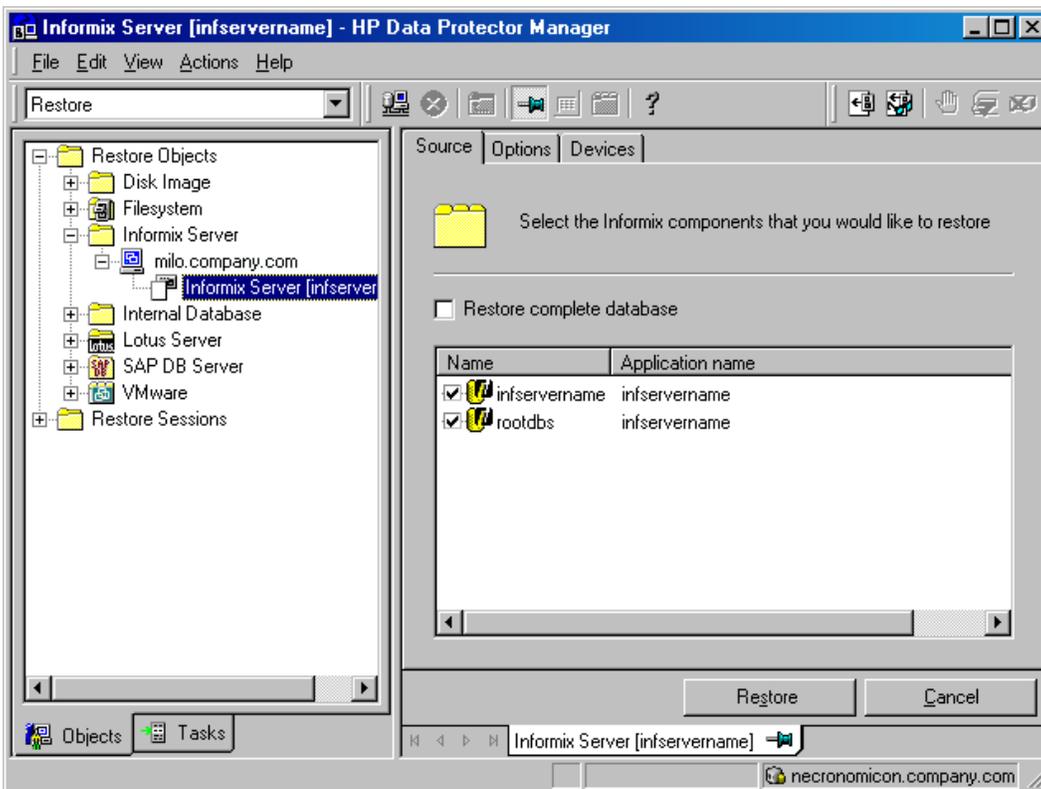


For details on the omnidb command, see the omnidb man page.

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**.

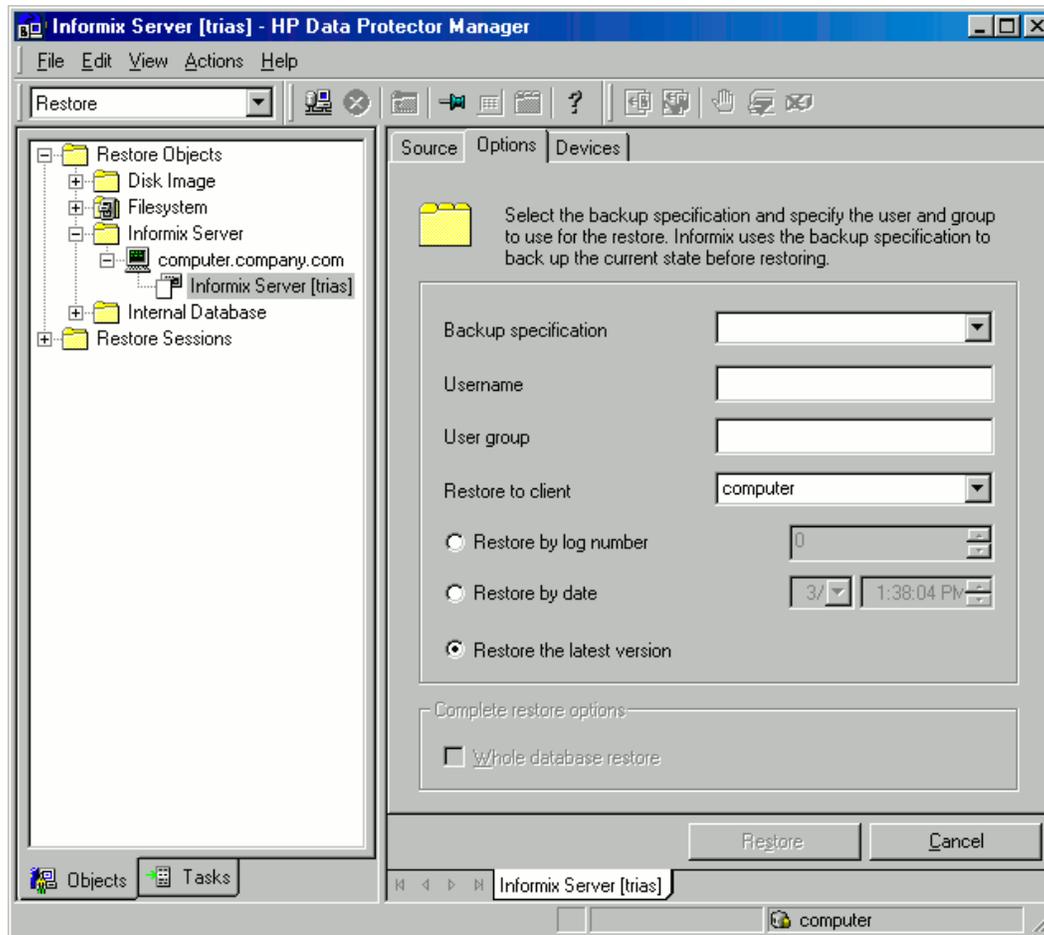
Figure 21: Selecting objects for restore



4. In the **Options** page, set the Informix Server specific restore options. For information, see [Informix](#)

Server restore options on page 37 or press F1.

Figure 22: Informix Server restore options



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

For more information on how to select devices for a restore, see the *HP Data Protector Help* index: “restore, selecting devices for”.

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

```
onmode -ky
```

Click **Restore**.

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

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

output.

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

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

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

```
onmode -m
```

Table 10: 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 systems: User name of the Informix Server backup owner. <code>onbar</code> is started under the account of the specified user.
User group	UNIX systems: 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 <code>onbar -r -n Last_Log_number</code> . For details, see the <i>Backup and restore guide</i> of Informix Server.

Option	Description
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 <code>onbar -r -t time</code> . For details, see the <i>Backup and restore guide</i> 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 <code>onbar -r -w</code> . For details, see the <i>Backup and restore guide</i> of Informix Server.

Important: After the restore, make sure that before you perform the next restore, a full backup has been performed.

Restoring using the Data Protector CLI

Before you begin the restore procedure, set the `OB2BARLIST` environment variable as described in “Data Protector and Informix Server variables” (page 32) [Backup on page 18](#). For example:

```
set OB2BARLIST=dbspace5
```

Run the following command:

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

<i>ClientName</i>	Name of the Informix Server system. In a cluster environment, name of the virtual server.
<i>INFORMIXSERVER</i>	Name of the Informix instance.

<i>User, Group</i>	UNIX systems: The user name and its group name.
<i>OnBarRestoreArguments</i>	ON-Bar restore arguments. Put each argument in double quotes.
<i>INFORMIX_OPTIONS</i>	A subset of general restore options. For information, see the <code>omnir</code> man page.

Important: After the restore, make sure that before you perform the next restore, a full backup has been performed.

Example

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

```
omnir -informix -barhost computer -barcmd ob2onbar.pl -user informix:informix -  
apptime 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 [Data Protector and Informix Server variables on page 29](#).
- If a disk failure occurs, salvage logical log files that are still on the disk by executing:

```
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.

Important: After the restore, make sure that before you perform the next restore, a full backup has been performed.

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, execute:

```
onbar -r -p
```

Restoring a particular dbspace or blobspace

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

```
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 the user `informix` on the target client.
3. Create an Informix instance with the same database name and the same server number as the original Informix instance by using the Informix Server ON-Monitor utility on the target client.

To obtain the database name and the server number, log in as the user `informix` on the original server and execute the following:

- a. To obtain the database name, look up the value of `DBSERVERNAME` in the `onstat -c` output.

On UNIX, you can do this by executing: `onstat -c | grep DBSERVERNAME`

- b. To obtain the server number database name, look up the value of `SERVERNUM` in the `onstat -c` output.

On UNIX, you can do this by executing: `onstat -c | grep SERVERNUM`

4. Ensure that the Informix instance is online.
5. Configure the Informix instance as described in [Configuring Informix instances on page 12](#).
6. Take the Informix instance offline.
7. Copy the following original Informix Server configuration files to the target client:

- ONCONFIG
 - the emergency boot file
 - `oncfg_DBSERVERNAME.SERVERNUM`
8. On UNIX, copy also the `sqlhosts` file to the target client. Change the source client host name in the copied `sqlhosts` file to the target client host name.
 9. On UNIX, add the `service_name` entry from the `sqlhosts` file to the `etc/services` file, together with a unique port number (for example, `1535/tcp`) on the target client to allow the instance to start running properly.
 10. Re-create the database files from the original database on the target client and then alter the files permission and ownership of the file to match the originals.
 11. Start a whole-system restore of `dbojects` as described in [Restoring using the Data Protector GUI on page 35](#).

Restoring using another device

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

Using the Data Protector GUI

For information on how to select another device for a restore using the Data Protector GUI, see the *HP Data Protector 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 systems: `Data_Protector_program_data\Config\Server\cell\restoredev`

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

Use the format:

`"DEV 1" "DEV 2"`

where `DEV 1` is the original device and `DEV 2` 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 *HP Data Protector 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 47](#) 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 *HP Data Protector Help* index: “patches”.
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

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 systems: `Data_Protector_home\log`

HP-UX and Solaris systems: /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 the *HP Data Protector Help*.
- **Windows systems:** Ensure that the Data Protector Inet service is running under the account `informix`.
- **UNIX systems:** 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 *HP Data Protector 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:

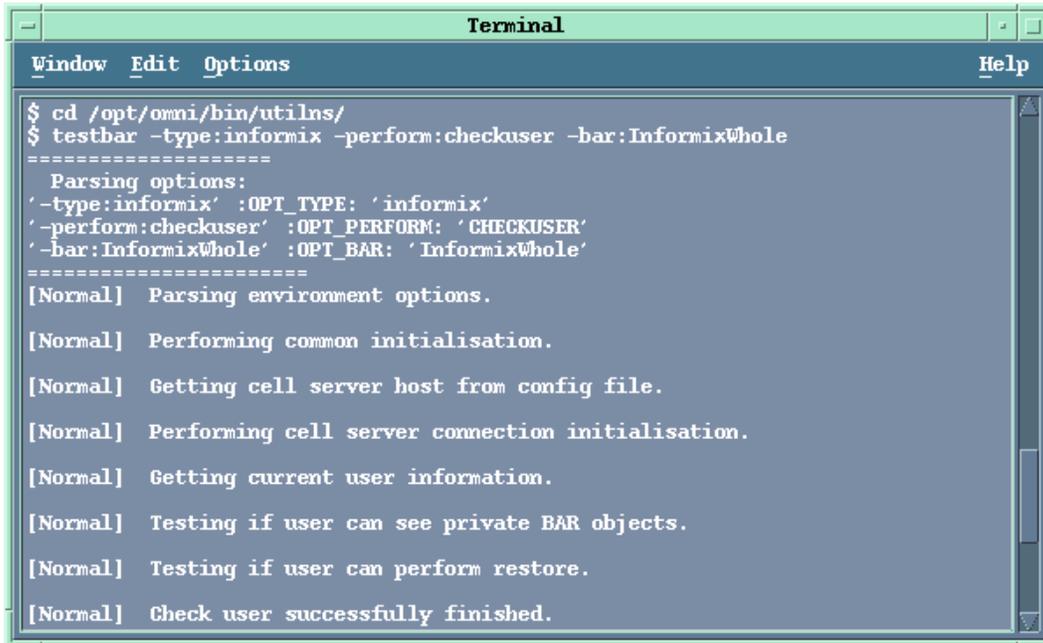
HP-UX and Solaris systems: /opt/omni/bin/utilns

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

execute:

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

Figure 23: Example of checking the Informix Server user



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

In 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: 08/06/2011
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 on page 17](#).
- Test the backup specification as described in [Previewing backup sessions on page 26](#).

- If this fails, check if the Informix Server part of the test failed:

Execute 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 the *HP Data Protector 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 29](#).

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 systems: `Data_Protector_home\bin`

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

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

- if your backup failed, execute:

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

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

- if your restore failed, execute:

```
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:

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

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

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

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

Windows systems: *Data_Protector_home*\log\debug.log

HP-UX and Solaris systems: /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:
 - a. Set the BAR_BSALIB_PATH shell variable to:

Windows systems: *ISMDIR*\bin\libbsa.dll

where *ISMDIR* is the path to the ISM.

UNIX systems: *INFORMIXDIR*/lib/ibsad001.sl

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

- b. 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:
 - a. Log in to the Informix Server system as user `informix`.
 - b. Execute the following:

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

UNIX systems: `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.SERVENUM`, 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 `SERVENUM` configuration parameter). Therefore, the restore of such objects may fail.

Action

Edit the emergency boot file to reflect the changed Data Protector session IDs. List the changed session IDs during the import procedure.

Information about backed-up objects is stored in the emergency boot file in the following format:

```
0DS730 rootdbs R 1 7 0 9 2011008018 2011-08-18 18:10:25 1
```

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

Here, the session ID is `2011/08/18-9`. 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 `SERVENUM` 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*.

Chapter 2: Data Protector DB2 UDB integration

Introduction

This chapter explains how to configure and use the Data Protector 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 11: 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.

Limitations

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 the *HP Data Protector Help*.

Integration concept

Data Protector integrates with the DB2 Server through a set of modules responsible for data backup and restore. [DB2 integration architecture](#) shows the architecture of the Data Protector DB2 integration.

Figure 24: DB2 integration architecture

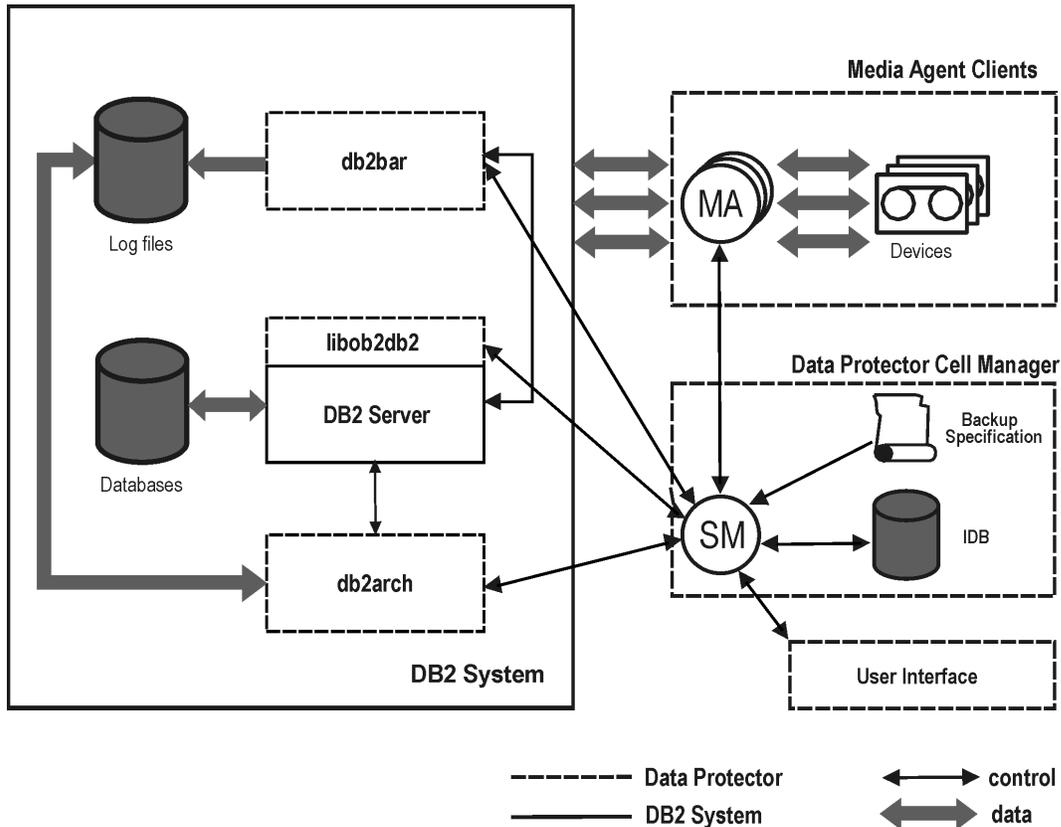


Table 12: Legend

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://support.openview.hp.com/selfsolve/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 `MaxBSession` global option 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 *HP Data Protector 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

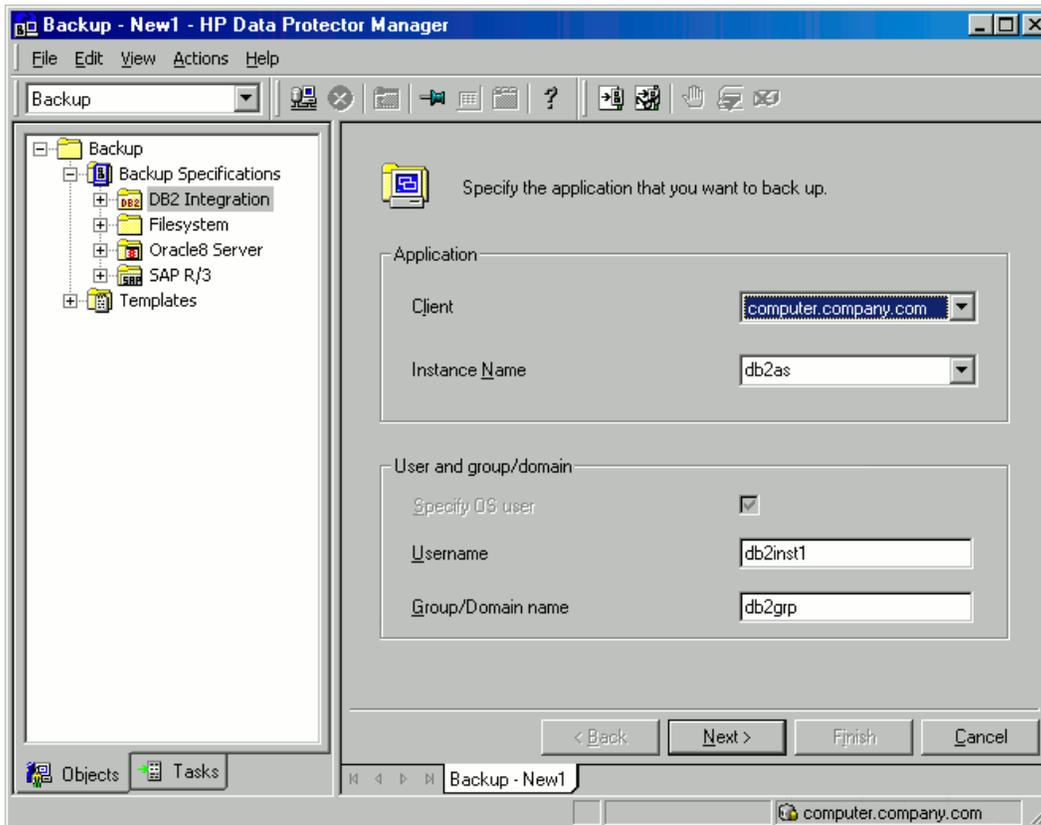
1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **DB2 Integration**, and click **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. In Client, select the DB2 Server system.

In a cluster environment, select the virtual server.

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

For information on the **User and group/domain** options, press **F1**.

Figure 25: 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 50](#).

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 [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 Selecting DB2 objects below. on page 55](#).

Using the Data Protector CLI

Execute the following command:

```
util_db2 -CONFIG DB2_instance username password [DB2 _instance_home]
```

Parameter description

<i>DB2_instance</i>	Name of the DB2 instance.
<i>username</i>	DB2 user.
<i>password</i>	DB2 user password.
<i>DB2_instance_home</i>	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

Execute the following command:

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

Backup

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

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.

Table 14: 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.

Including log files in a backup image

By default, Data Protector does not include log files in a backup image. To include the latest log files in the backup image, set the omnirc variable `OB2_DB2INCLDLOGS` on the DB2 Server system to 1. During a restore session, the included logs are restored to the Data Protector temporary folder and used in a rollforward recovery.

Note that you still need the archived logs if you want to perform a recovery to a point in time which is not covered by the included logs.

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 61](#).

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 15: 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.

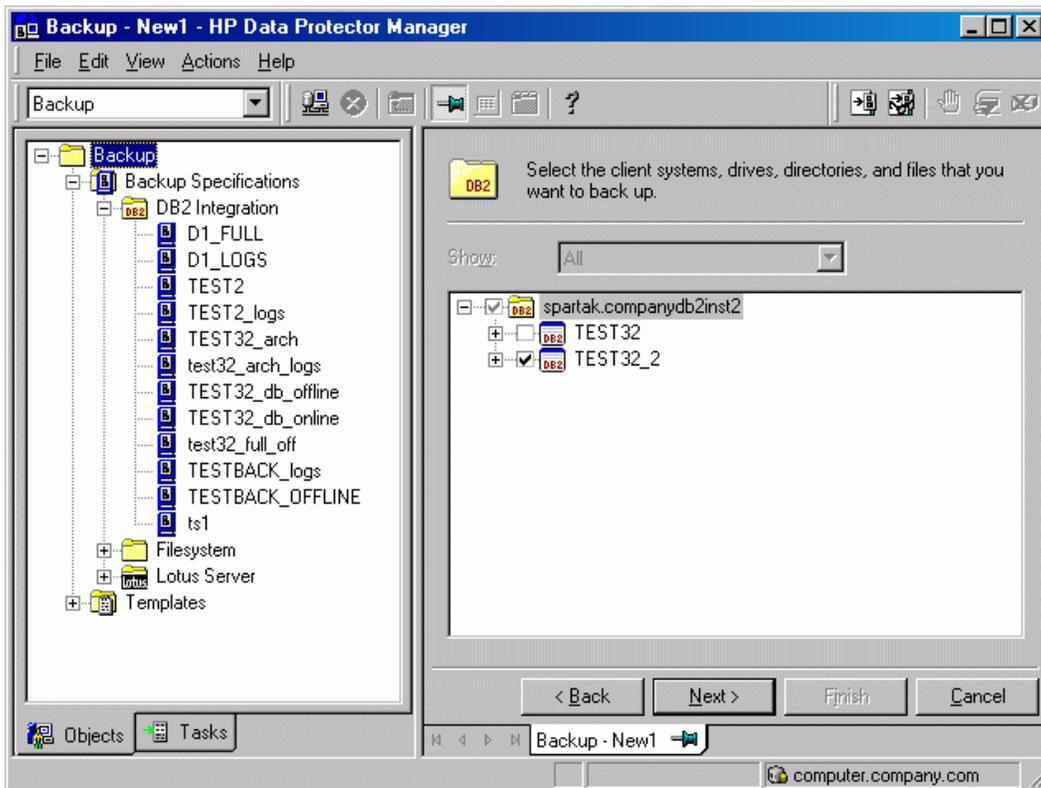
4. In **Client**, select the DB2 Server system; in a cluster environment, select the virtual server.

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

For information on the **User and group/domain** options, press **F1**.

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 51](#).
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 [Selecting DB2 objects below](#).

Figure 26: 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**.

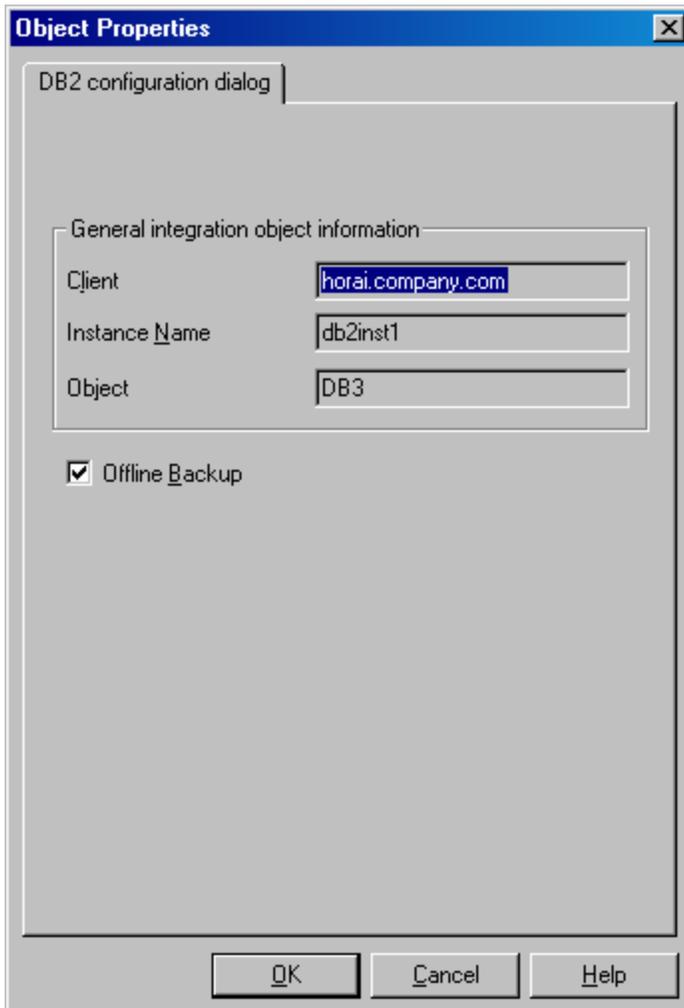
8. Set backup options and click **Next**.

For information on application-specific options, see [DB2 backup options on page 58](#).

9. Optionally, schedule the backup and click **Next**. For more information, see [Scheduling backup sessions on page 58](#).

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 [Selecting offline backup on the next page](#).

Figure 27: 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 26](#).

Table 16: DB2 backup options

Pre-exec Post-exec	<p>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.</p> <p>Type only the name of the command, not the pathname. The command must reside in:</p> <p>Windows systems: <i>Data Protector\bin</i></p> <p>HP-UX systems: <i>/opt/omni/lbin</i></p> <p>Other UNIX systems: <i>/usr/omni/bin</i></p>
Parallelism	<p>Specify the number of data streams for backing up a database from a node.</p> <p>In a partitioned environment, Parallelism must equal the device concurrency.</p> <p>Default: 1.</p>

Modifying backup specifications

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

Scheduling backup sessions

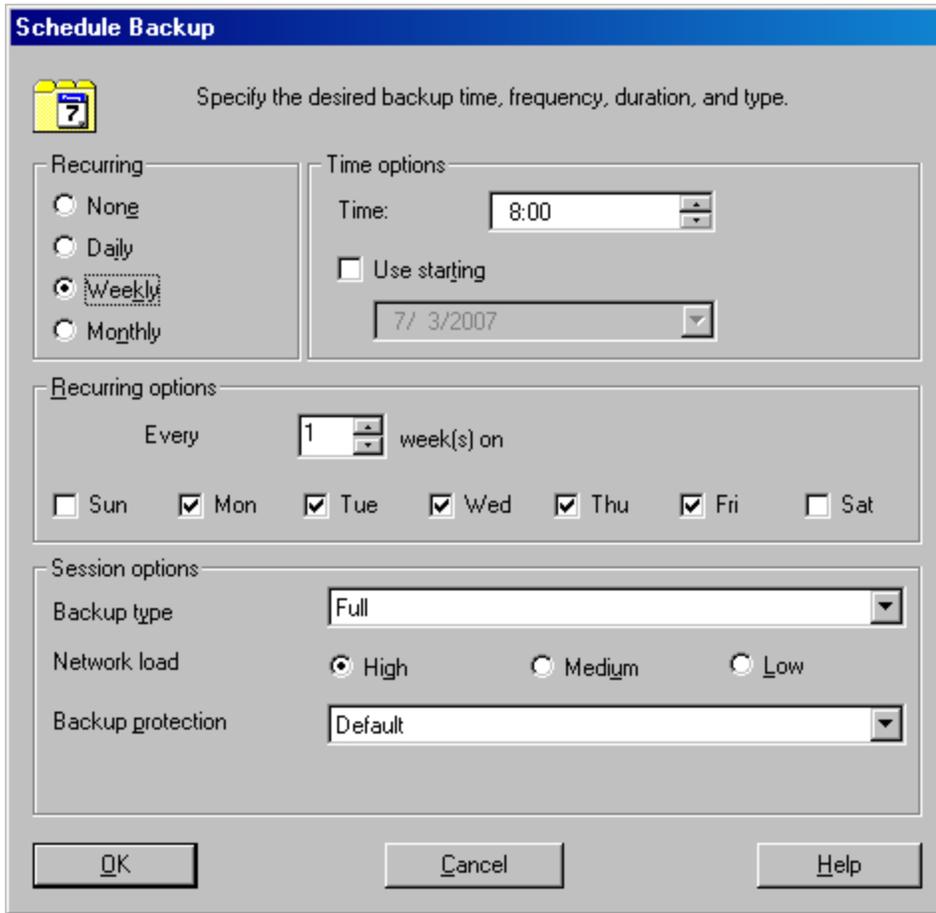
You can run unattended backups at specific times or periodically. For details on scheduling, see the *HP Data Protector Help* index: "scheduled backups".

Example

To 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 [Scheduling a backup session on the next page](#). Click **OK**.
3. Repeat [In the Schedule property page, select the starting date in the calendar and click Add to open the Schedule Backup dialog box.](#) above and Under Recurring, select **Weekly**. Under Time options, select **8:00**. Under Recurring Options, select **Mon, Tue, Wed, Thu, and Fri**. See [Scheduling a backup session on the next page](#). Click **OK.** above to schedule another backup at 13:00, and another one at 18:00.
4. Click **Apply** to save the changes.

Figure 28: Scheduling a backup session



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

Execute the following command:

```
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:

a. 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.

b. Restart the database.

c. 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 use 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

Execute the following command:

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

Table 17: Parameter description

<code>db2_mode</code>	Backup type: {-full -incr -delta}
<code>options</code>	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, execute:

```
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 *HP Data Protector 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.

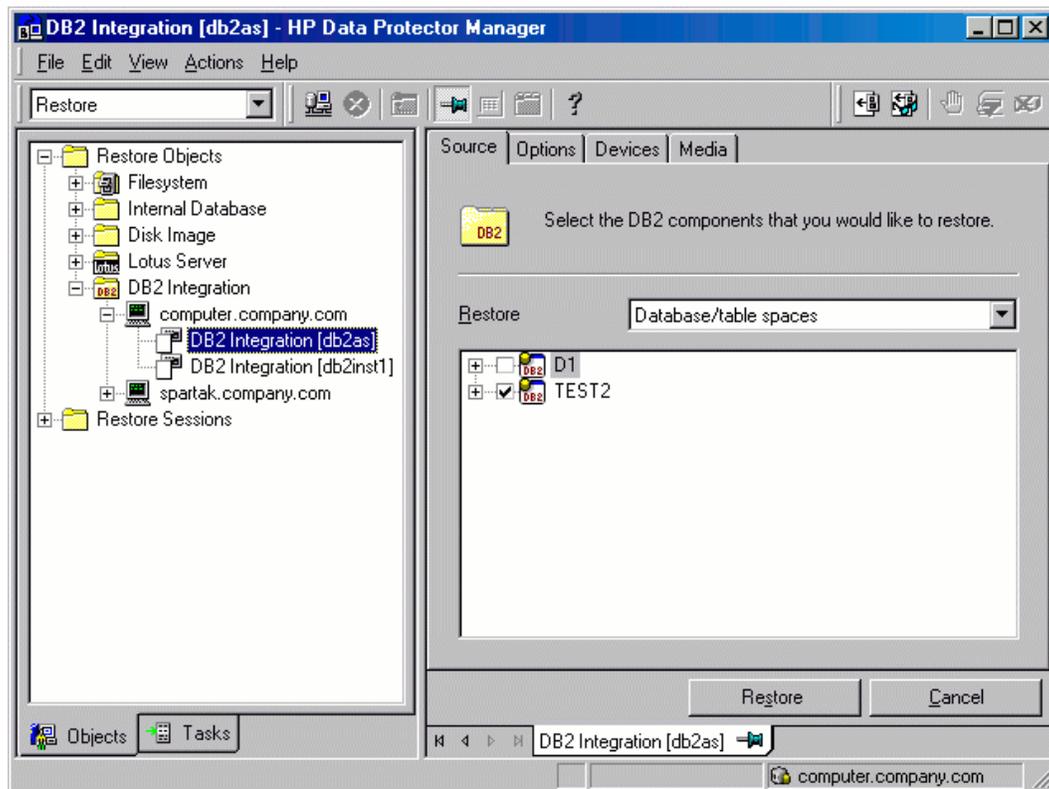
For information on how to restore a DB2 database to a new database, see [Restoring to a new database or another DB2 instance on page 67](#).

For information on how to restore partitioned DB2 objects, see [Restore in a partitioned environment on page 70](#).

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 [Selecting objects for restore below](#).

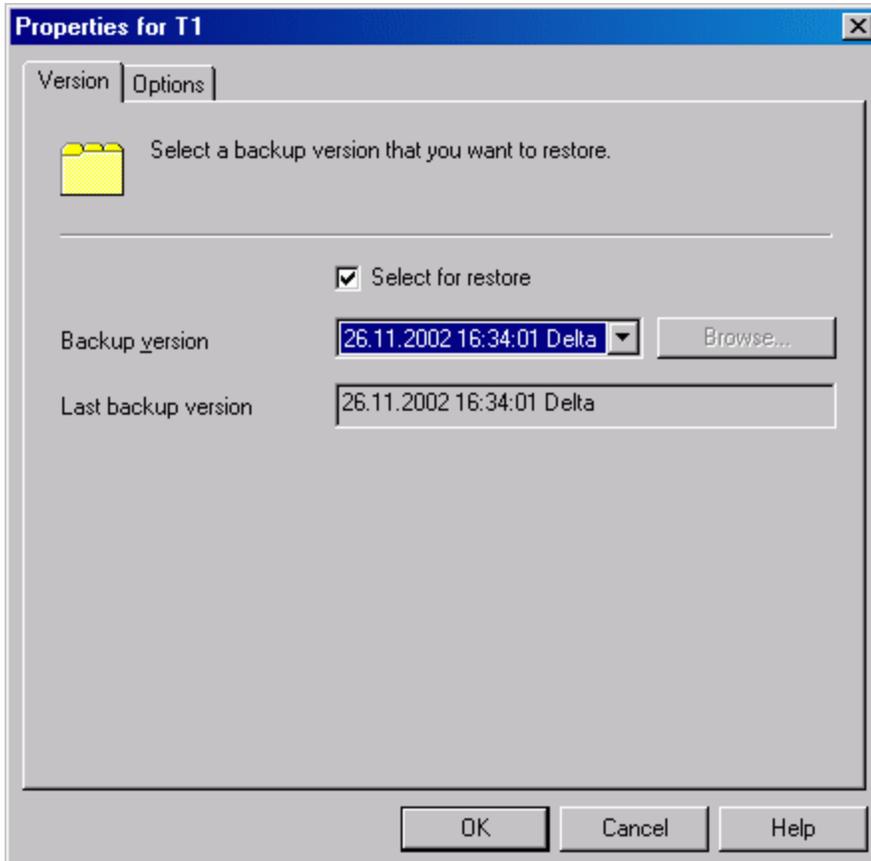
Figure 29: Selecting objects for restore



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

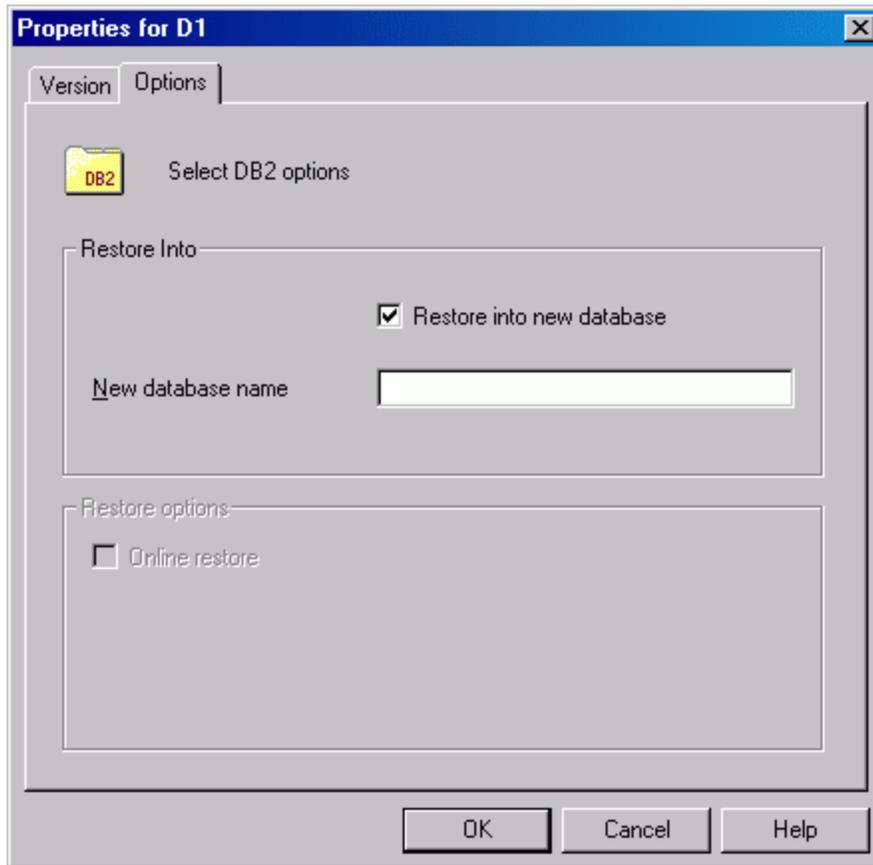
By default, the latest backup version is restored. To restore a DB2 object from a specific backup version, right-click the object, click **Properties**, and specify the backup version in the Properties for *DB2_object* dialog box. See [Selecting a version](#) .

Figure 30: 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 [Restoring to a new database on the next page](#).

Figure 31: Restoring to a new database



4. In the Options page, set the DB2 restore options. For information, see [DB2 restore options](#) 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.

For more information of how to select devices for a restore, see the *HP Data Protector Help* index: "restore, selecting devices for".

Click **Restore**.

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

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

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

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

Table 18: 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 12 . For details, see Restoring to a new database or another DB2 instance on page 67 .

<p>Rollforward</p>	<p>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. If log files are included in the backup and the omnirc variable <code>OB2_DB2INCLDLOGS</code> is set to 1, the included logs are restored to the Data Protector temporary folder. Specify the rollforward recovery by selecting <code>Rollforward</code> to the end of the logs or <code>Rollforward to date</code>. When specifying <code>Rollforward to date</code>, use the coordinated universal time (UTC).</p> <p>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.</p> <p>To perform a rollforward recovery in a physically partitioned environment, restore all the parts with <code>Rollforward</code> cleared (see Restore in a partitioned environment on page 70), connect to the catalog node, and then start a rollforward recovery using the DB2 Command Line Processor.</p> <p>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 <code>Rollforward</code> 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).</p>
---------------------------	--

Restoring using the Data Protector CLI

Execute the following command:

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

Parameter description

<i>source_client</i>	The DB2 Server system from which DB2 objects were backed up. In a cluster environment, the name of the virtual server.
<i>target_client</i>	The target DB2 Server system (only if you are not restoring to the source client).
<i>source_instance</i>	The DB2 instance whose DB2 objects were backed up.

<i>target_instance</i>	The target DB2 instance.
<i>source_db</i>	The database you want to restore.
<i>new_db</i>	The target database (specify only if not the source database).
<i>table_space</i>	The table space you want to restore.
<i>Log_file</i>	The log file you want to restore.
<i>BackupID</i>	<p>Specifies from which backup data to restore, for example, 2011/10/09-2.</p> <p>A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.</p> <p>Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original backup session (that is, the backup ID) and not the session ID of the object copy session.</p> <p>The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.</p>

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 2011, 9:15 a.m., execute:

```
omnir -db2 -barhost degas -instance DB2Inst -dbname TEMP -rollforward time: 2011-01-10.09.15.00
```

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. a. Find the containers of the source database:
 - o 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
```

- o To list the containers for a particular table space, run:

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

Define new table space containers for the non-system table spaces by adding options for redirection to the DB2 configuration file. Execute the following command for every pair of table space containers:

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

The DB2 user of the target instance must have read and write permissions for the new containers.

- b. If you are using an automatic storage database, in which table spaces can be created and whose container and space management characteristics are completely determined by the DB2 database manager, define a new storage path. To do this, execute the following command for each storage path:

```
util_cmd -putopt DB2 target_instance "index_number" "new_storage_path" -  
sublist Autostore/source_db
```

Parameter description

<i>target_instance</i>	The target instance.
<i>source_db</i>	The backed up database.

2. In a physically partitioned environment, repeat [Find the containers of the source database: on the previous page](#) on every system.
3. 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.

4. 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 database `db2db_old` to the database `db2db_new` from an online backup (both databases reside in the instance `db2inst`, 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 database `db2db_old` to the database `db2db_new`, 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 database db2db from the instance inst1 to the database db2db in the instance inst2:

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 database db2db to the instance inst1, using the Data Protector CLI:

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

Example 3

To restore the automatic storage database db2db_old, which has two associated storage paths, to the database db2db_new from an online backup:

1. Check if the paths exist and specify new storage paths, using the Data Protector CLI:

```
util_cmd -putopt DB2 inst2 "1" "c:\db2\db2db_new\1" -sublist Autostore/db2db_old
```

```
util_cmd -putopt DB2 inst2 "1" "c:\db2\db2db_new\2" -sublist Autostore/db2db_old
```

2. Restore the database db2db_old to the database db2db_new, using the Data Protector CLI or GUI.
3. Restore all archived logs needed for rollforward recovery using the Data Protector GUI.
4. Copy the archived and the redo logs of the source database to the corresponding log directories of the new database.
5. Perform a rollforward recovery to the end of logs, using the DB2 CLI.

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 other way round) 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 67](#).

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:
 - a. Restore the object, without specifying rollforward recovery. Use the Data Protector GUI or CLI.
 - b. 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 67](#).

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 67](#).

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 *HP Data Protector 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 47](#) 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 *HP Data Protector Help* index: “patches” for more information of how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <http://support.openview.hp.com/selfsolve/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

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

Windows systems: `Data Protector\log`

HP-UX and Solaris systems: `/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 26](#).
 - 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 the *HP Data Protector Help*.

Additionally, if your backup or restore failed:

- Try performing:
 - A Data Protector filesystem backup and restore. For information, see the *HP Data Protector 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 fails

When performing an offline backup, the session fails with an error similar to the following:

```
[Major] From: OB2BAR_DB2BAR@ DB2ClientName InstanceName Time: DateTime  
DB2 returned error:  
SQL1035N The database is currently in use. SQLSTATE=57019
```

Action

The error implies that there are still some existing connections to the database. Execute:

```
db2 force application all
```

Note: This command disconnects all users and aborts all existing transactions.

Alternatively, you can disconnect individual users or applications. To get a list of all applications, execute:

```
db2 list applications
```

To disconnect the user that is identified with the application handle number *AppNum*, execute:

```
db2 force application AppNum
```

Tip: You can create a pre-exec script to execute the command. Note that the command returns the status `completed` before the operation is actually finished. Therefore, provide extra time for the operation to complete by adding the `sleep` command to the script.

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 session for restoring data from an object copy gets blocked

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:
 - a. 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.

- b. 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.
- c. 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.

Chapter 3: Data Protector Lotus Notes/Domino Server integration

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.

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

Table 19: Lotus Notes/Domino Server backup types

Full	<p>Backs up all the selected Lotus Notes/Domino Server databases.</p> <p>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.</p>
Incremental	<p>Backs up the selected Lotus Notes/Domino Server databases that meet at least one of the following two conditions:</p> <ul style="list-style-type: none">• The size of the changes made to a database since it was last backed up exceeds the size set by the Amount of log changes (KB) option.• The Lotus Notes/Domino Server DBIID for a database has changed. <p>Databases that do not meet at least one of the two conditions are not backed up.</p> <p>If archived logs are selected, it also backs up the archived logs that have not been backed up yet.</p>

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 the *HP Data Protector 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. [Data Protector Lotus Notes/Domino Server integration architecture below](#) shows the architecture of the integration.

Figure 32: Data Protector Lotus Notes/Domino Server integration architecture

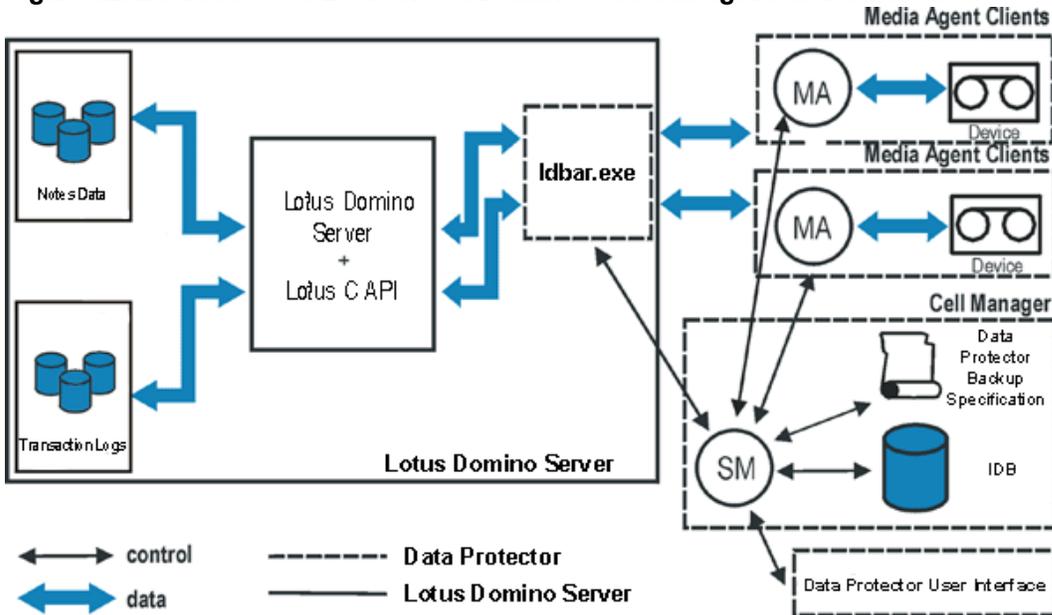


Table 20: Legend

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, HP 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 Enterprise Server or the Lotus Domino Utility Server.

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.

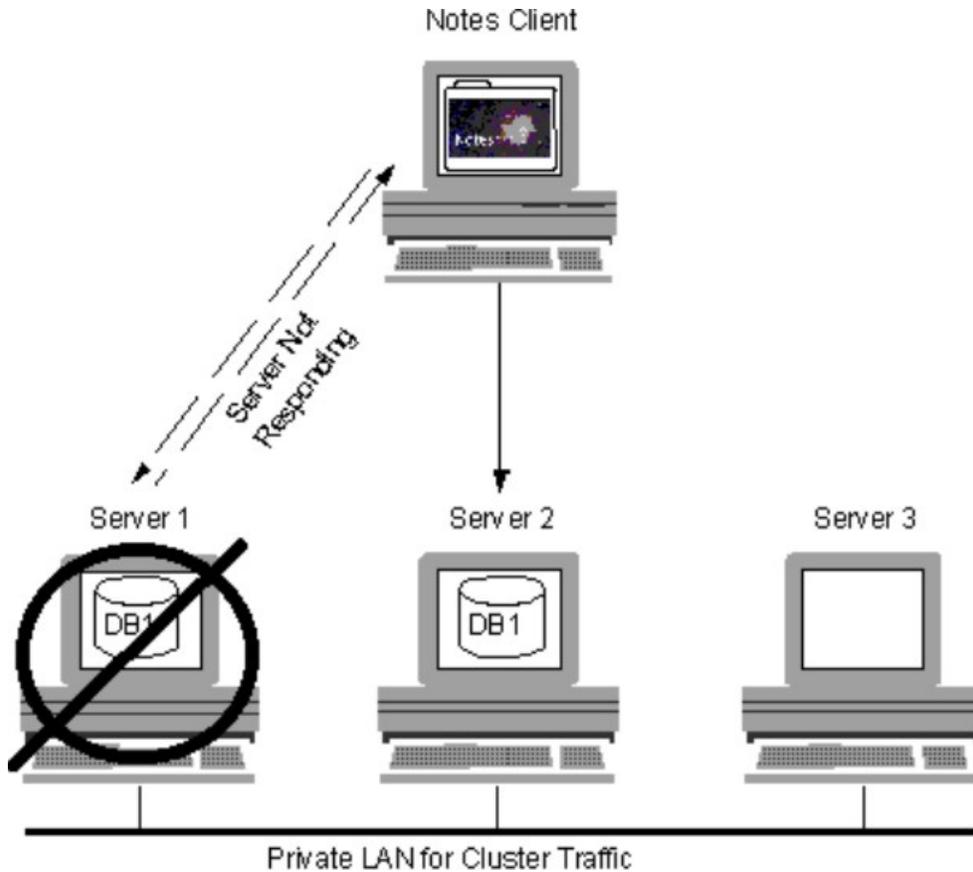
If the source server shuts down before replication completes, the replication events in memory are lost. For this reason, you should use standard replication (the REPLICA task) to perform immediate replication with all members of the cluster whenever you restart a cluster server. It is also a good idea 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://support.openview.hp.com/selfsolve/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.
- **Windows systems:**
 - On Windows Server 2003 system, you need to restart the Data Protector Inet service under a Windows domain user account that has the appropriate Lotus Notes/Domino Server permissions for running backups and restores. Stop the service and restart it as user lotus.

For information on changing the user account under which the Data Protector Inet service is running, see the *HP Data Protector Help* index: "Inet, changing account".

- On other Windows operating systems, configure the Data Protector Inet service user impersonation for the user that has the appropriate Lotus Notes/Domino Server permissions for running backups and restores.

For details, see the *HP Data Protector Help* index: "Inet user impersonation".

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 21: 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

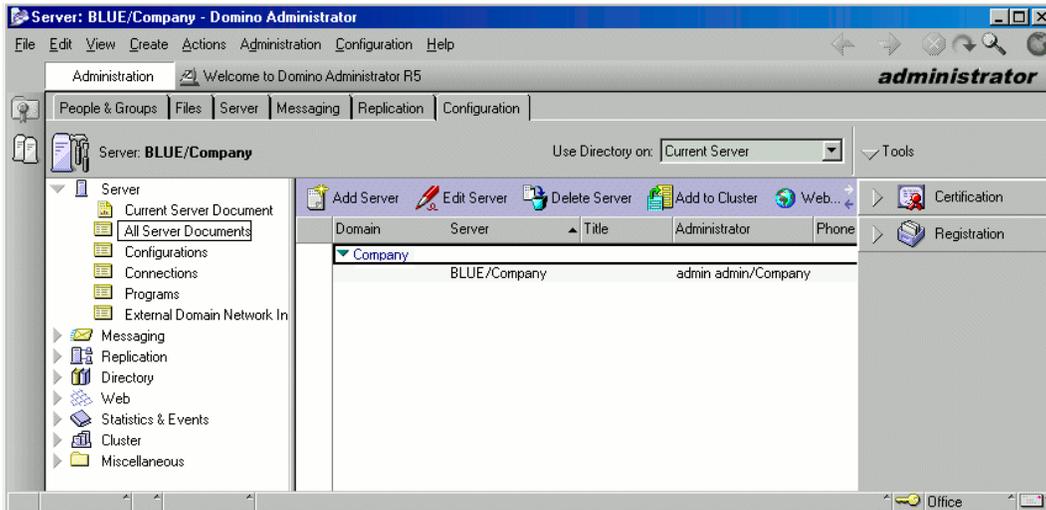
Use Lotus Domino Administrator on the Lotus Notes/Domino Server system. Alternatively, use Web Administrator or edit the `notes.ini` file.

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

To enable transaction logging and set archived logging:

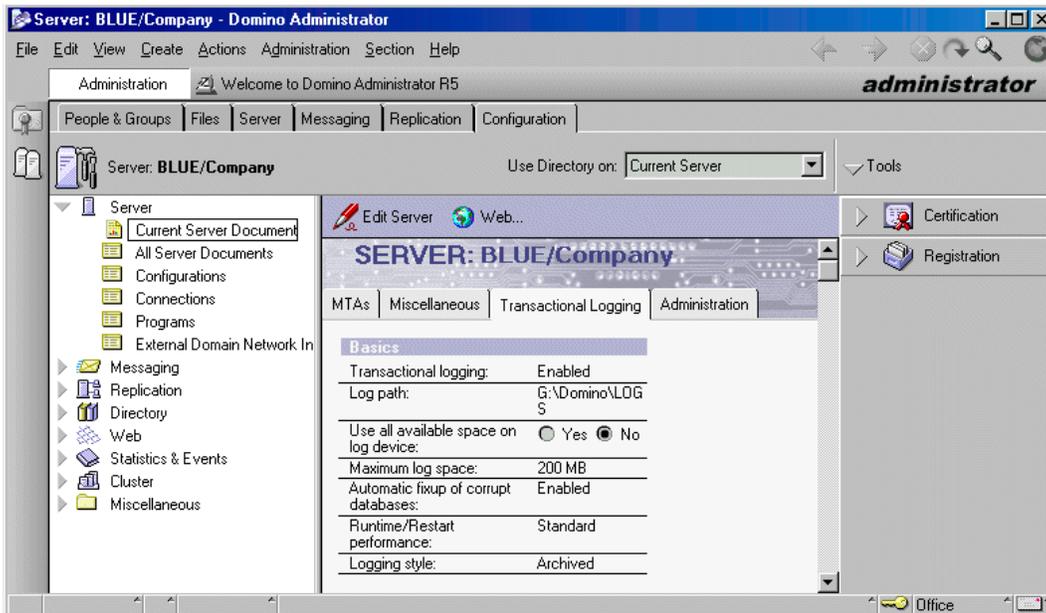
1. Start Lotus Domino Administrator.
2. Log on to Lotus Notes/Domino Server and select the **Configuration** tab.
3. Expand **Server**, select **All Server Documents**, and select the desired Lotus Notes/Domino Server. See [Browsing Lotus Notes/Domino Server on the next page](#).

Figure 33: Browsing Lotus Notes/Domino Server



4. Select the **Transactional Logging** tab and set appropriate values. See [Enabling archived transactional logging below](#).

Figure 34: 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 *HP Data Protector Help* index: “adding users”.

Configuring Lotus Notes/Domino Server systems

Using the Data Protector GUI

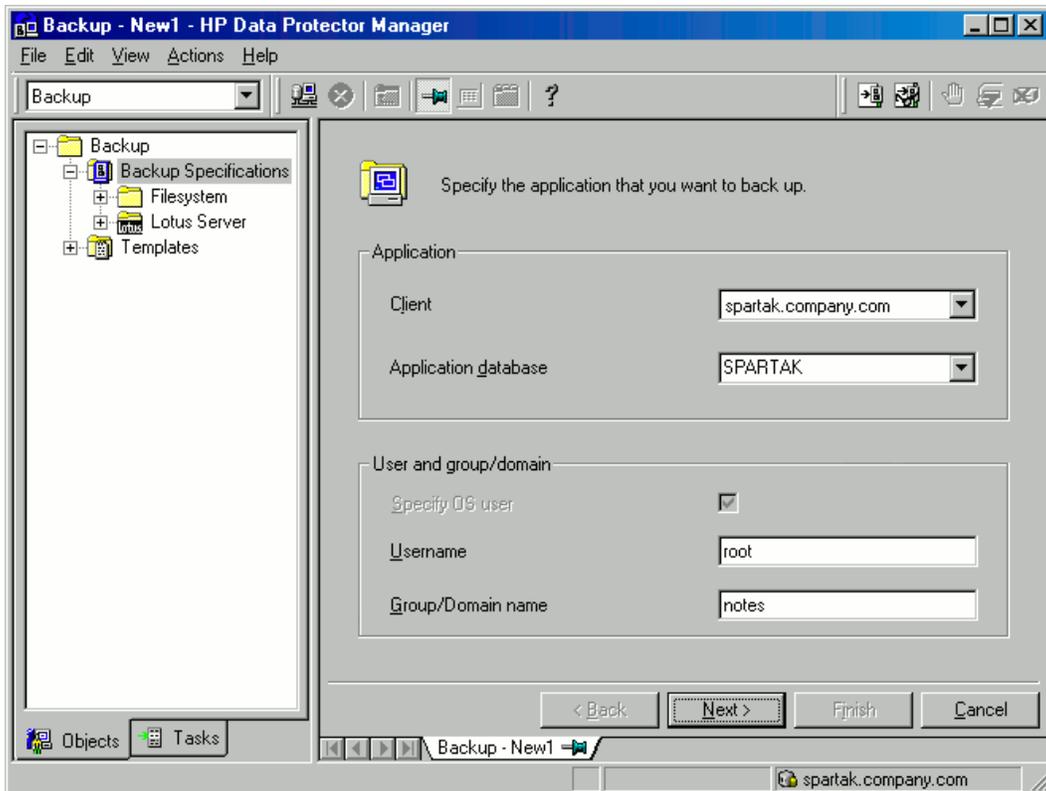
1. In the Context list, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Lotus Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, click **OK**.
4. In Client, select the **Lotus Notes/Domino Server system**. In a cluster environment, select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

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

For information on the **User and group/domain** options, press **F1**.

See [Specifying the Lotus Notes/Domino Server system below](#).

Figure 35: Specifying the Lotus Notes/Domino Server system



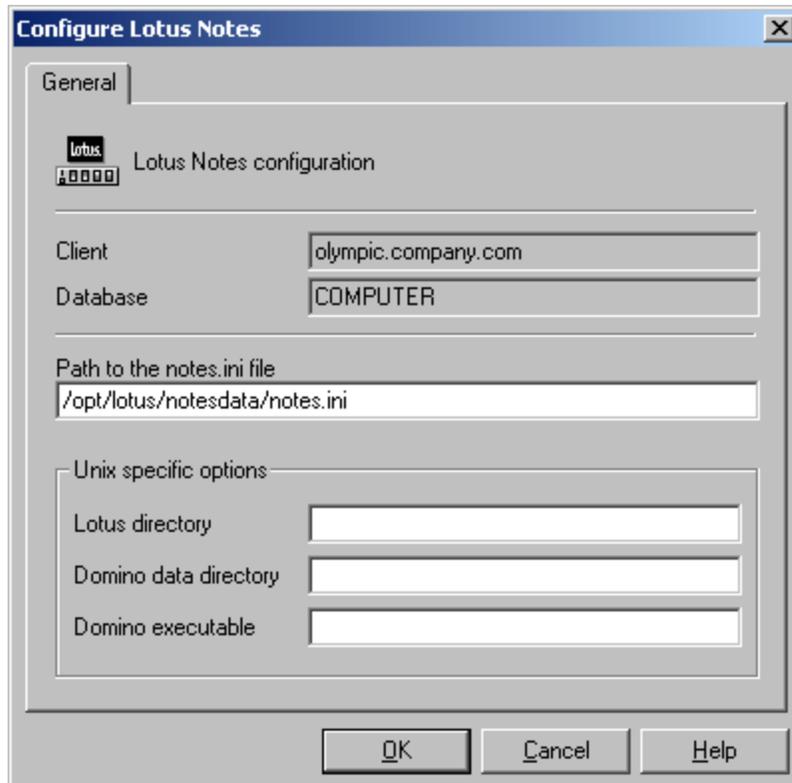
Click **Next**.

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 [Specifying Lotus Notes/Domino Server data below](#).

Figure 36: Specifying Lotus Notes/Domino Server data



Click **OK**.

If an error occurs, click **Details** or see [Troubleshooting on page 106](#).

6. The integration is configured. Exit the GUI or proceed with creating the backup specification at [Step 6](#).

Using the Data Protector CLI

On the Lotus Notes/Domino Server system, execute:

Windows systems:

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

Solaris systems:

```
/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 systems:

```
/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]
```

Table 22: Parameter description

<i>SRV_NAME</i>	Lotus Notes/Domino Server name.
<i>notes.ini_file</i>	Pathname of the Lotus Notes/Domino Server notes.ini file.
<i>Lotus_home_directory</i>	Pathname of the Lotus Notes/Domino Server home directory.
<i>Domino_data_directory</i>	Pathname of the Lotus Notes/Domino Server data directory.
<i>Domino_executables_ directory</i>	Pathname of the Lotus Notes/Domino Server executables directory.

Note: UNIX systems: 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 systems:

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

Solaris systems:

```
/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 systems: `Data_Protector_home\bin`

Solaris systems: `/opt/omni/lbin`

AIX systems: `/usr/omni/bin`

execute:

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

Solaris systems: Execute:

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

AIX systems: Execute:

```
/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	<p>Backs up all the selected Lotus Notes/Domino Server databases.</p> <p>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.</p>
------	--

Incremental	<p>Backs up the selected Lotus Notes/Domino Server databases that meet at least one of the following two conditions:</p> <ul style="list-style-type: none">• The size of the changes made to a database since it was last backed up exceeds the size set by the Amount of log changes (KB) option.• The Lotus Notes/Domino Server DBIID for a database has changed. <p>Databases that do not meet at least one of the two conditions are not backed up.</p> <p>If archived logs are selected, it also backs up the archived logs that have not been backed up yet.</p>
-------------	--

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.

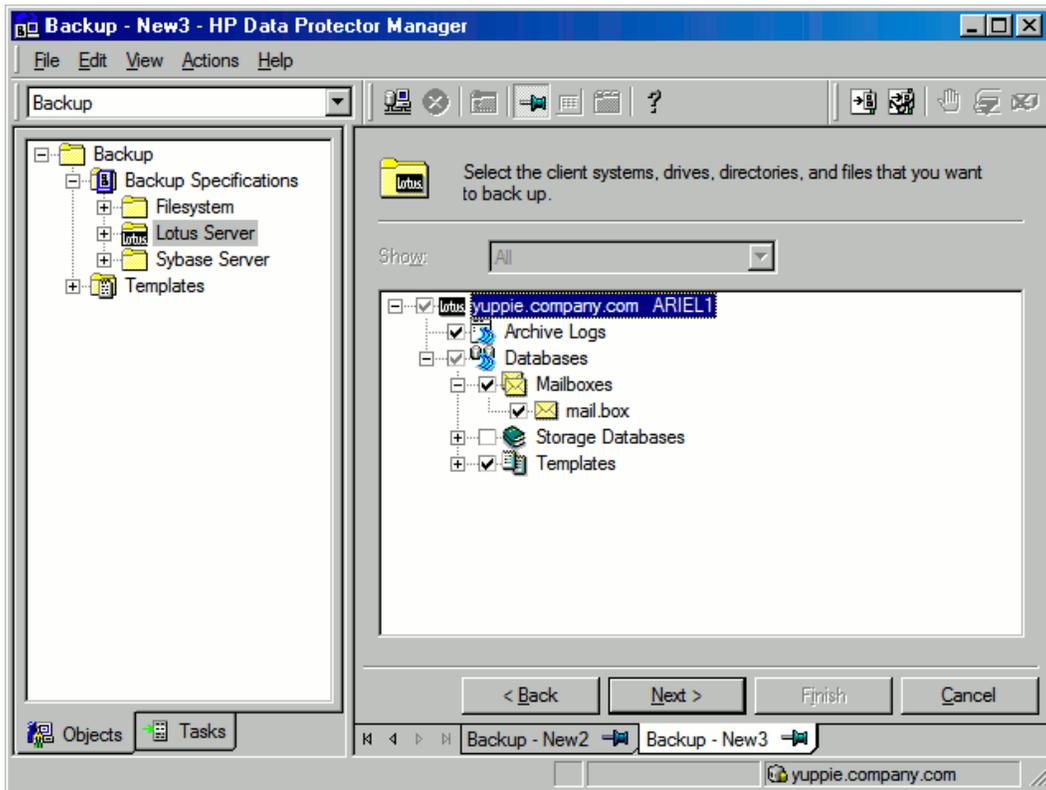
In **Application database**, select the Lotus Notes/Domino Server to be backed up.

For information on the **User and group/domain** options, press **F1**.

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 86](#).
6. Select the Lotus Notes/Domino Server objects to be backed up. See [Selecting backup objects](#).

Figure 37: Selecting backup objects



Click **Next**.

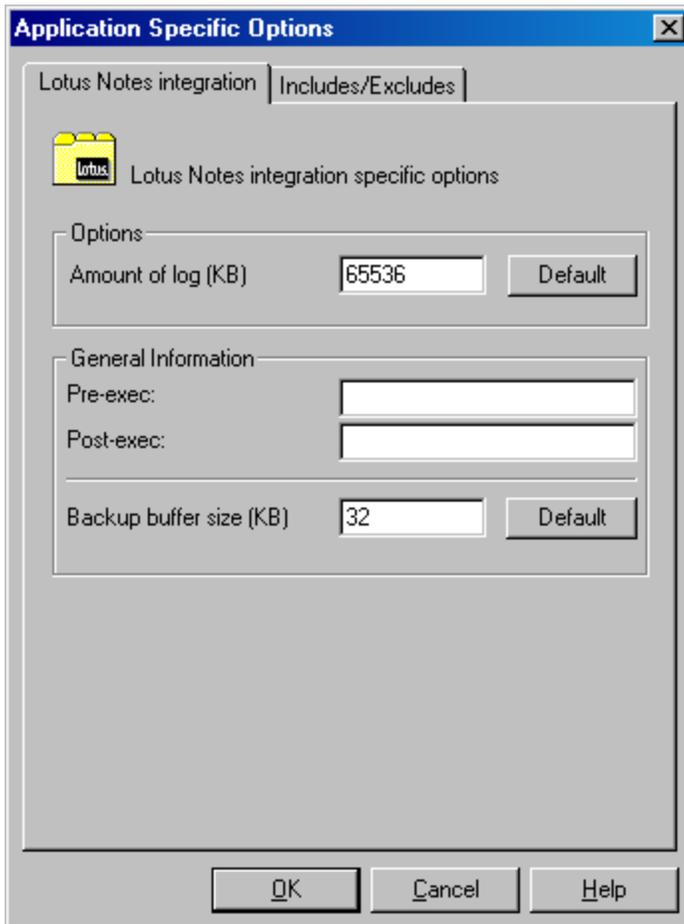
7. Select devices to use for the backup.

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

Click **Next**.

8. Set backup options. For information on the application-specific options ([Application-specific options](#)), see [Lotus Notes/Domino Server backup options on page 94](#) or press **F1**.

Figure 38: Application-specific options



Click **Next**.

9. Optionally, schedule the backup and click **Next**. See [Scheduling backup sessions on the next page](#).
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 95](#).

Table 24: Lotus Notes/Domino Server backup options

Amount of log changes (KB)	<p>Applies to incremental backups. The backup skips the database if at least one of the following two conditions is met:</p> <ul style="list-style-type: none">• The database to be backed up has a smaller log amount than specified by the option.• The Lotus Notes/Domino Server DBIID for the database has not changed. <p>If the database exceeds the specified log amount or if the Lotus Notes/Domino Server DBIID for the database has not changed, the database is backed up.</p>
Pre-exec, Post-exec	<p>Specify a command that will be started by <code>ldbar.exe</code> on the Lotus Notes/Domino Server system before the backup (<i>pre-exec</i>) or after it (<i>post-exec</i>). The command must reside in the directory:</p> <p>Windows systems: <code>Data_Protector_home\bin</code></p> <p>Solaris and Linux systems: <code>/opt/omni/bin</code></p> <p>AIX systems: <code>/usr/omni/bin</code></p> <p>In the backup specification, provide only the filename.</p>
Backup buffer size	<p>The size of the buffer used for reading and writing data during the backup.</p>

Modifying backup specifications

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

Scheduling backup sessions

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

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 [Scheduling backup sessions on the next page](#).

Click **OK**.

3. Repeat [Step 1](#) and [Step 2](#) to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

Figure 39: Scheduling backup sessions

Schedule Backup

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

Recurring

None
 Daily
 Weekly
 Monthly

Time options

Time: 9:00
 Use starting
7/ 3/2007

Recurring options

Every 1 week(s) on

Sun Mon Tue Wed Thu Fri Sat

Session options

Backup type: Full
Network load: Full
Backup protection: Default

OK Cancel Help

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.

Execute the following command:

```
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 are 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 use 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 taken 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 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.

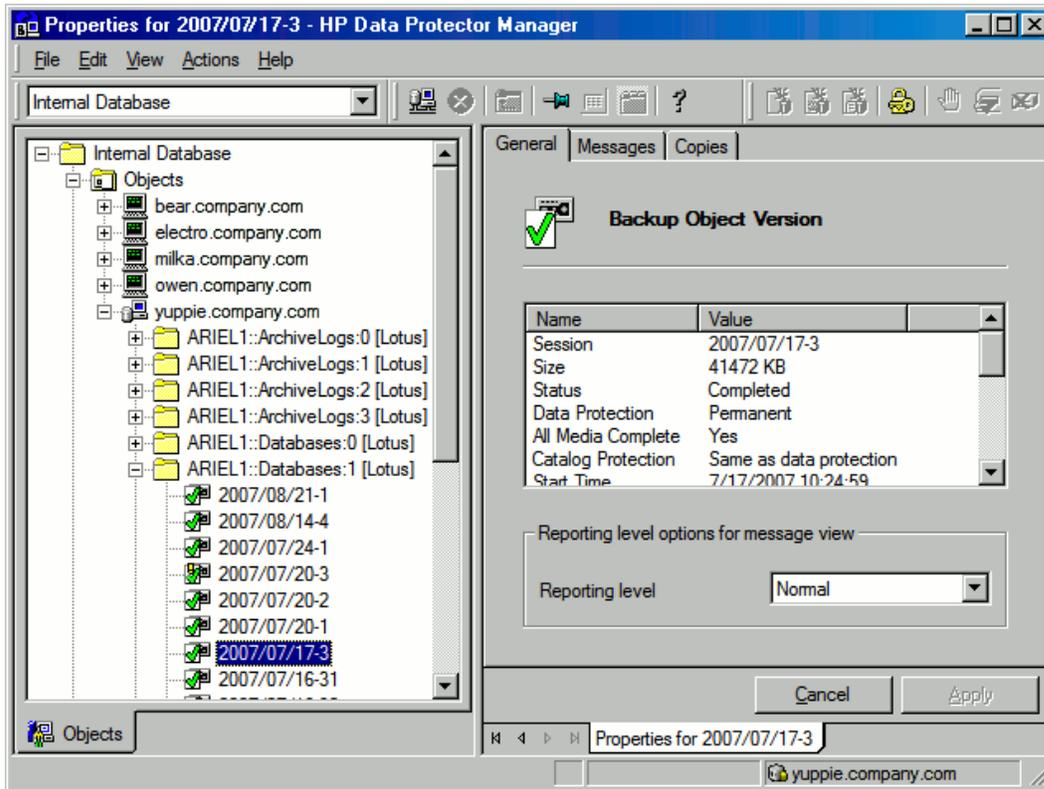
Finding information for restore

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

```
omnidb -session session_id
```

Figure 41: Lotus Notes/Domino Server objects from a particular session

```

G:\WINDOWS\system32\cmd.exe
G:\Program Files\OmniBack\bin>omnidb -session 2007/10/26-2
Object Name                               Object Type   Object Status  CopyID
-----
behar.company.com:BAMBI::Databases:5      Lotus        Completed     149 (0)
behar.company.com:BAMBI::Databases:4      Lotus        Completed     150 (0)
behar.company.com:BAMBI::Databases:6      Lotus        Completed     151 (0)
behar.company.com:BAMBI::Databases:7      Lotus        Completed     152 (0)
behar.company.com:BAMBI::Databases:8      Lotus        Completed     153 (0)
behar.company.com:BAMBI::Databases:9      Lotus        Completed     154 (0)
behar.company.com:BAMBI::Databases:10     Lotus        Completed     155 (0)
behar.company.com:BAMBI::Databases:11     Lotus        Completed     156 (0)
behar.company.com:BAMBI::Databases:12     Lotus        Completed     157 (0)
behar.company.com:BAMBI::Databases:13     Lotus        Completed     158 (0)
behar.company.com:BAMBI::Databases:0      Lotus        Completed     159 (0)
behar.company.com:BAMBI::Databases:1      Lotus        Completed     160 (0)
behar.company.com:BAMBI::Databases:3      Lotus        Completed     161 (0)
behar.company.com:BAMBI::Databases:2      Lotus        Completed     162 (0)
behar.company.com:BAMBI::ArchiveLogs:0    Lotus        Completed     163 (0)

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

2. 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 42: Lotus Notes/Domino Server databases of a particular object

```

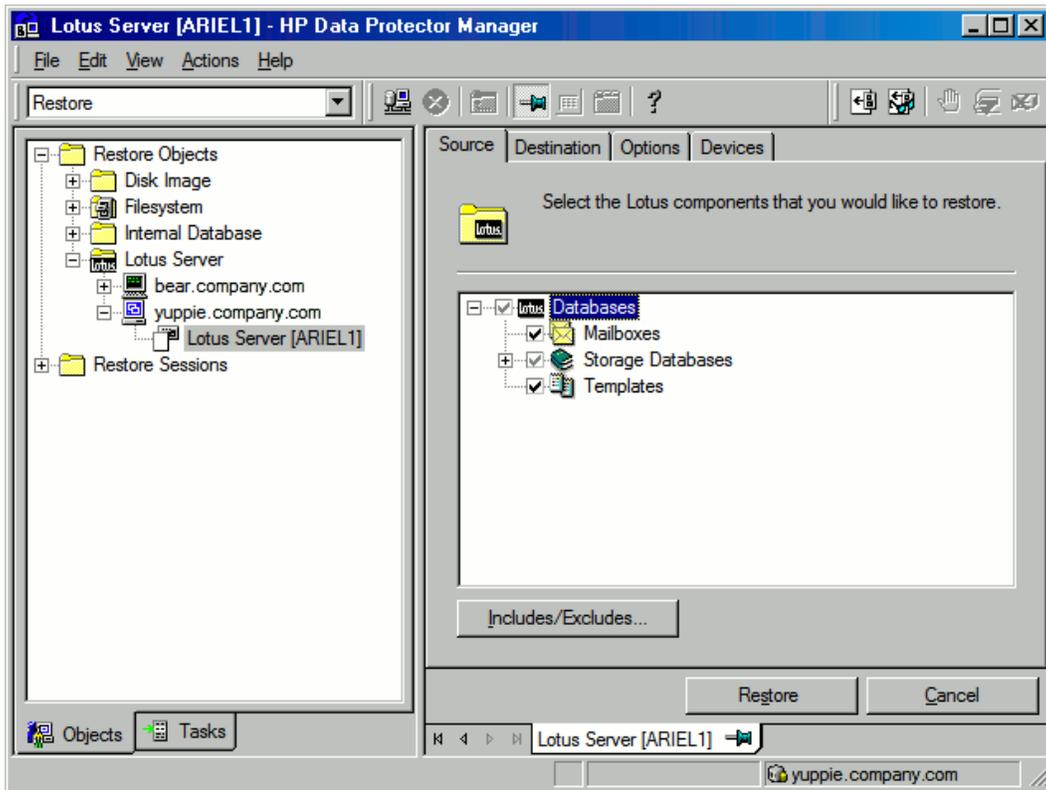
G:\WINDOWS\system32\cmd.exe
G:\Program Files\OmniBack\bin>omnidb -lotus behar.company.com:BAMBI::Databases:5-
session 2007/10/26-2 -catalog
Protection  Owner      Group      Size      Time      Path
-----
-2          nogroup   1613824    26.10.2007 5:35:03  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/admin4.ntf
-2          nogroup   1384448    26.10.2007 5:35:03  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/certreq.ntf
-2          nogroup   241152     26.10.2007 5:35:03  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/dbdirman.ntf
-2          nogroup   327680     26.10.2007 5:35:03  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/decomsrv.ntf
-2          nogroup   835584     26.10.2007 5:35:04  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/nntpd150.ntf
-2          nogroup   294912     26.10.2007 5:35:04  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/ntsunc45.ntf
-2          nogroup   10747904   26.10.2007 5:35:04  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/pubnames.ntf
-2          nogroup   1018880    26.10.2007 5:35:03  /0b27c7dd-ac2d-47c1
-9e8b-5ccd22bd4878/BAMBI/Databases/dspd.ntf
    
```

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 [Selecting objects for restore on the next page](#).

Figure 43: 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.

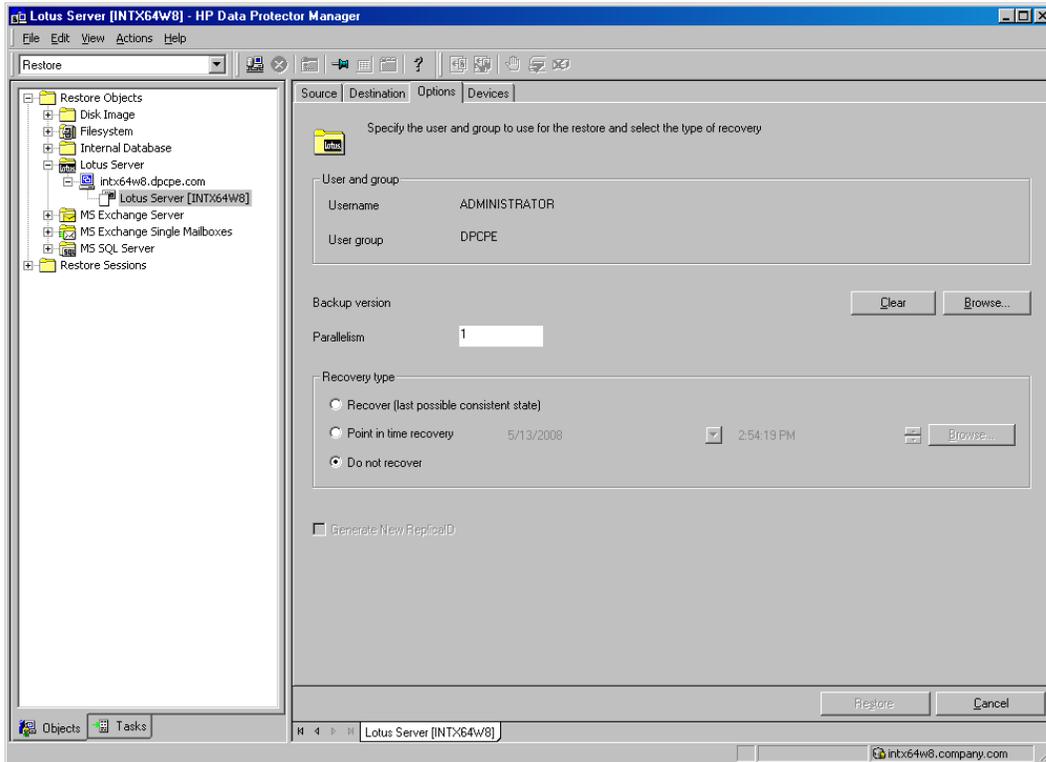
You can select the backup version in the Options page ([Lotus Notes/Domino Server restore options on the next page](#)). Click **Browse** to select a different version of backup.

4. In the Destination page, set the destination options. For information, see [Destination options on page 102](#) 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 ([Lotus Notes/Domino Server restore options on the next page](#)). For information on the application-specific options, see [Restore options on page 103](#) or press **F1**.

Figure 44: Lotus Notes/Domino Server restore options



6. In the Devices page, select the devices to be used for the restore.

For more information on how to select devices for a restore, see the *HP Data Protector 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**.

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

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

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

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 systems:** 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 systems: `C:\Lotus\Domino\BLUE\`

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

To restore a database to the directory:

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

UNIX systems: `/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 systems: Username of the Lotus Notes/Domino Server backup owner, for example, notes.
	User group	UNIX systems: 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.
Generate New ReplicaID		This option is only available with the recovery type Recover (last possible consistent state) . If this option is selected, each restored storage database (NSF database) is assigned a new replica ID. Default: not selected.

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](#), [Restoring with recovery to the latest possible state](#), and [Point-in-time recovery](#).

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 *HP Data Protector 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 *HP Data Protector 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 109](#) 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 *HP Data Protector Help* index: “patches”.
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

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 systems:

- Check if the nNotes.dll library is linked. Execute:

```
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 systems: `Data_Protector_home\log`

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

AIX systems: `/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 systems:** 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 *HP Data Protector 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 takes 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 88](#).
- 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.

- a. From the directory:

Windows systems: `Data_Protector_home\bin`

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

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

execute:

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

- b. 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 systems: `Data_Protector_home\bin`

Solaris systems: `/opt/omni/bin`

AIX systems: `/usr/omni/bin`

execute:

Windows systems:

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

UNIX systems:

```
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, execute `ldbar.exe -help`.

- **Windows systems:** 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 systems: `Data_Protector_home\bin`

Solaris systems: `/opt/omni/bin`

AIX systems: /usr/omni/bin

execute:

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

For other ldbar.exe parameters, execute 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

```
2011/08/26.18:15:00
```

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 106](#).

Problem

Slow incremental backup with large number of databases to be backed up

When backing up large numbers of databases with the set transaction logging and using the incremental backup type, the backup is slow.

Action

Set the omnirc file option OB2_LOTUS_NODBIID to 1. For information on how to use the omnirc file and its location, see the *HP Data Protector Troubleshooting Guide*. For information on the OB2_LOTUS_NODBIID option, see the file itself.

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 systems:** 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 systems:** Log in to the Lotus Notes/Domino Server system as user `root`.
2. **Windows systems:** Kill all the `ldbar.exe` processes using Task Manager.
3. **UNIX systems:** 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 systems: `Data_Protector_home\bin`

Solaris systems: `/opt/omni/lbin`

AIX systems: `/usr/omni/bin`

execute: `util_notes.exe -box -ini:path_to_notes.ini`

If everything is working properly, the `*RETVL*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.company.com "BLUE" Time: 8/22/2011 4:07:09 PM
      Lotus Notes C API 'NSFTakeDatabaseOffline' returned error 5098:
      The database is in use and cannot be taken offline.
```

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.company.com "BLUE" Time: 19.10.11 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

A

access rights

See user rights.

ACSL (StorageTek specific term)

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

Active Directory (Windows specific term)

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

AES 256-bit encryption

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

AML (ADIC/GRAU specific term)

Automated Mixed-Media library.

AMU (ADIC/GRAU specific term)

Archive Management Unit.

application agent

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

application system (ZDB specific term)

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

archive logging (Lotus Domino Server specific term)

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

archived log files (Data Protector specific term)

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

archived redo log (Oracle specific term)

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

ASR set

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

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

audit logs

Data files to which auditing information is stored.

audit report

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

auditing information

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

autochanger

See library.

autoloader

See library.

Automatic Storage Management (ASM) (Oracle specific term)

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

auxiliary disk

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

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

B

BACKINT (SAP R/3 specific term)

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

backup API (Oracle specific term)

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

backup chain

See restore chain.

backup device

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

backup generation

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

backup ID

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

backup object

A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk). A backup object is defined by: client name (hostname of the Data Protector client where the backup object resides), mount point (for filesystem objects - the access point in a directory structure on the client where the backup object is located (drive on Windows systems and mount point on UNIX systems), for integration objects - backup stream identification, indicating the backed up database/application items), description (for filesystem objects - uniquely defines objects with identical client name and mount point, for integration objects - displays the integration type), and type (for filesystem objects - filesystem type, for integration objects - "Bar").

backup owner

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

backup session

A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification, incremental backup, and full backup.

backup set

A complete set of integration objects associated with a backup.

backup set (Oracle specific term)

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

backup specification

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

backup system (ZDB specific term)

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

backup types

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

backup view

Data Protector provides different views of your backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of

backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC (EMC Symmetrix specific term)

Business Continuances are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

BC Process (EMC Symmetrix specific term)

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

BCV (EMC Symmetrix specific term)

Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected. See also BC and BC Process.

Boolean operators

The Boolean operators for the full text search functionality of the Data Protector Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

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

BRARCHIVE (SAP R/3 specific term)

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

BRBACKUP (SAP R/3 specific term)

An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files. See also BRARCHIVE and BRRESTORE.

BRRESTORE (SAP R/3 specific term)

An SAP R/3 tool that can be used to restore files of the following type: Database data files, control files, and online redo log files saved with BRBACKUP; Redo log files archived with BRARCHIVE; Non-database files saved with BRBACKUP. You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup. See also BRARCHIVE and BRBACKUP.

BSM

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

C

CAP (StorageTek specific term)

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

Catalog Database (CDB)

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

catalog protection

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

CDB

See Catalog Database (CDB).

CDF file (UNIX systems specific term)

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

cell

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

Cell Manager

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

centralized licensing

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

Centralized Media Management Database (CMMDB)

See CMMDB.

Certificate Server

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

Change Journal (Windows specific term)

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

Change Log Provider

A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

channel (Oracle specific term)

An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used: type 'disk' or type 'sbt_tape'. If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

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

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

client backup

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

backup specification was created are not backed up.

client or client system

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

cluster continuous replication (Microsoft Exchange Server specific term)

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

cluster-aware application

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

CMD script for Informix Server (Informix Server specific term)

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

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

CMMDB

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

COM+ Class Registration Database (Windows specific term)

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

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

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

command-line interface (CLI)

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

concurrency

See Disk Agent concurrency.

container (HP P6000 EVA Disk Array Family specific term)

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

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

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

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

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

CRS

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

CSM

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

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

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

data protection

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

Data Protector user account

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

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

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

data stream

Sequence of data transferred over the communication channel.

Data_Protector_home

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

Data_Protector_program_data

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

database library

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

database parallelism

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

database server

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

Dbobject (Informix Server specific term)

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

DC directory

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

DCBF

See Detail Catalog Binary Files (DCBF).

delta backup

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

Detail Catalog Binary Files (DCBF)

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

device

See backup device.

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (EMC Symmetrix specific term)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to: -Identify and work with a subset of the available

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

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes writing to the tape, and so on. If data is written to the tape slower than it is delivered to the device then the device is streaming. Streaming significantly improves the use of space and the performance of the device.

DHCP server

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

differential backup

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

differential backup (Microsoft SQL Server specific term)

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

differential database backup

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

directory junction (Windows specific term)

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

disaster recovery

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

disaster recovery operating system

See DR OS.

Disk Agent

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

Disk Agent concurrency

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

disk group (Veritas Volume Manager specific term)

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

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory

structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging

The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format

A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. See also virtual full backup.

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

DR image

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

DR OS

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

drive

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

drive index

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

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

drive-based encryption

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

E

EMC Symmetrix Agent

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

emergency boot file (Informix Server specific term)

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

encrypted control communication

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

encryption key

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

encryption KeyID-StoreID

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

enhanced incremental backup

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

enterprise backup environment

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

Event Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Event Logs (Windows specific term)

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

Exchange Replication Service (Microsoft Exchange Server specific term)

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

exchanger

See library.

exporting media

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

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

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

F**failover**

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

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

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

FC bridge

See Fibre Channel bridge.

Fibre Channel

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

Fibre Channel bridge

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

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

file depot

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

file jukebox device

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

file library device

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

File Replication Service (FRS)

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

file tree walk

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

file version

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

filesystem

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

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

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

flash recovery area (Oracle specific term)

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

formatting

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

free pool

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

full backup

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

full database backup

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

full mailbox backup

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

full ZDB

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

G

global options

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

group (Microsoft Cluster Server specific term)

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

GUI

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

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

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

heartbeat

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

Hierarchical Storage Management (HSM)

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

Holidays file

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

hosting system

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

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

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

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

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

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

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

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

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

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

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

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

HP P6000 / HP 3PAR SMI-S Agent

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

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

HP P9000 XP Agent

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

HP SMI-S P6000 EVA Array provider

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

I

ICDA (EMC Symmetrix specific term)

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

IDB

See Internal Database (IDB).

IDB recovery file

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

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

importing media

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

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

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

incremental backup

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

incremental backup (Microsoft Exchange Server specific term)

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

incremental mailbox backup

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

incremental restore (EMC Symmetrix specific term)

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

incremental ZDB

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

Incremental1 Mailbox Backup

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

Inet

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

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

Information Store (Microsoft Exchange Server specific term)

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

Informix Server (Informix Server specific term)

Refers to Informix Dynamic Server.

initializing

See formatting.

Installation Server

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

instant recovery (ZDB specific term)

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

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

integration object

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

Internal Database (IDB)

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

Internet Information Server (IIS) (Windows specific term)

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

ISQL (Sybase specific term)

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

J

jukebox

See library.

jukebox device

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

K

Key Management Service (Microsoft Exchange Server specific term)

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

keychain

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

keystore

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

KMS

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

L

LBO (Symmetric specific term)

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

one entity and can only be restored as a whole.

LDEV (HP P9000 XP Disk Array Family specific term)

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

library

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

lights-out operation or unattended operation

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

LISTENER.ORA (Oracle specific term)

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

load balancing

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

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

local and remote recovery

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

Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication (Microsoft Exchange Server specific term)

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

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

lock name

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

log_full shell script (Informix Server UNIX systems specific term)

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

logging level

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

logical-log files

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

login ID (Microsoft SQL Server specific term)

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

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

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

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

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

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

Lotus C API (Lotus Domino Server specific term)

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

LVM

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

M**Magic Packet**

See Wake ONLAN.

mailbox (Microsoft Exchange Server specific term)

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

mailbox store (Microsoft Exchange Server specific term)

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

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

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

maintenance mode

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

make_net_recovery

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

make_tape_recovery

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

Manager-of-Managers

See MoM.

MAPI (Microsoft Exchange specific term)

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

MCU

See Main Control Unit (MCU).

Media Agent

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

media allocation policy

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

media condition

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

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

media condition factors

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

media label

A user-defined identifier used to describe a medium.

media location

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

media management session

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

media pool

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

media set

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

media type

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

media usage policy

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

medium ID

A unique identifier assigned to a medium by Data Protector.

merging

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

Microsoft Exchange Server

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

**Microsoft Management Console (MMC)
(Windows specific term)**

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

Microsoft SQL Server

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

Microsoft Volume Shadow Copy Service (VSS)

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

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

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

See target volume.

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

See replica set rotation.

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

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

mirrorclone (HP P6000 EVA Disk Array Family specific term)

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

MMD

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

MMDB

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

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

MoM

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

mount point

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

mount request

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

MSM

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

multisnapping (HP P6000 EVA Disk Array Family specific term)

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



OBDR capable device

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

obdrindex.dat

See IDB recovery file.

object

See backup object.

object consolidation

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

object consolidation session

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

object copy

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

object copy session

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

object copying

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

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

object ID (Windows specific term)

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

object mirror

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

object mirroring

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

object verification

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

object verification session

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

offline backup

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

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

offline recovery

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

offline redo log

See archived redo log.

ON-Bar (Informix Server specific term)

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

ONCONFIG (Informix Server specific term)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file located in

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

online backup

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

online recovery

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

online redo log (Oracle specific term)

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

OpenSSH

A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard (Oracle specific term)

Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance (Oracle specific term)

Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID (Oracle specific term)

A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <ORACLE_SID>. The <ORACLE_SID> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system

The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files. See also merging.

ownership

The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner. When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

P1S file

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

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

package (HP ServiceGuard and Veritas Cluster Specific Term)

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

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

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

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

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

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

parallel restore

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

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

parallelism

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

phase 0 of disaster recovery

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

phase 1 of disaster recovery

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

phase 2 of disaster recovery

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

phase 3 of disaster recovery

Restoration of user and application data.

physical device

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

post-exec

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

pre- and post-exec commands

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

prealloc list

A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec

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

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

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

protection

See data protection and catalog protection.

public folder store (Microsoft Exchange Server specific term)

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

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

public/private backed up data

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

R

RAID

Redundant Array of Independent Disks.

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

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

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

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

rawdisk backup

See disk image backup.

RCU

See Remote Control Unit (RCU).

RDBMS

Relational Database Management System.

RDF1/RDF2 (EMC Symmetrix specific term)

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

Recovery Catalog (Oracle specific term)

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

Recovery Catalog Database (Oracle specific term)

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

recovery files (Oracle specific term)

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

Recovery Manager (RMAN) (Oracle specific term)

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

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

RecoveryInfo

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

recycle or unprotect

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

redo log (Oracle specific term)

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

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

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

Removable Storage Management Database (Windows specific term)

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

to access and share the same media resources.

reparse point (Windows specific term)

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

replica (ZDB specific term)

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

replica set (ZDB specific term)

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

replica set rotation (ZDB specific term)

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

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

restore chain

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

restore session

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

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

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

RMAN (Oracle specific term)

See Recovery Manager.

RSM

The Data Protector Restore Session Manager controls restore and object

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

RSM (Windows specific term)

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

S

scanning

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

Scheduler

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

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

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

session

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

session ID

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

session key

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

shadow copy (Microsoft VSS specific term)

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

shadow copy provider (Microsoft VSS specific term)

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

shadow copy set (Microsoft VSS specific term)

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

shared disks

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

Site Replication Service (Microsoft Exchange Server specific term)

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

slot

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

SMB

See split mirror backup.

SMBF

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

SMI-S Agent (SMISA)

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

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

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

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

snapshot backup

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

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

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

source (R1) device (EMC Symmetrix specific term)

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

source volume (ZDB specific term)

A storage volume containing data to be replicated.

sparse file

A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

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

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

split mirror backup (EMC Symmetrix specific term)

See ZDB to tape.

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

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

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

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

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

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

sqlhosts file or registry (Informix Server specific term)

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

SRD file

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

SRDF (EMC Symmetrix specific term)

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

SSE Agent (SSEA)

See HP P9000 XP Agent.

sst.conf file

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

st.conf file

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

stackers

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

standalone file device

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

Storage Group (Microsoft Exchange Server specific term)

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

storage volume (ZDB specific term)

A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist.

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

StorageTek ACS library (StorageTek specific term)

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

switchover

See failover.

Sybase Backup Server API (Sybase specific term)

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

Sybase SQL Server (Sybase specific term)

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

SYMA

See EMC Symmetrix Agent.

synthetic backup

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

synthetic full backup

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

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

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

system databases (Sybase specific term)

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

System Recovery Data file

See SRD file.

System State (Windows specific term)

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

system volume/disk/partition

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

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

SysVol (Windows specific term)

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

T

tablespace

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

tapeless backup (ZDB specific term)

See ZDB to disk.

target (R2) device (EMC Symmetrix specific term)

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

target database (Oracle specific term)

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

target system (disaster recovery specific term)

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

target volume (ZDB specific term)

A storage volume to which data is replicated.

Terminal Services (Windows specific term)

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

thread (Microsoft SQL Server specific term)

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

TimeFinder (EMC Symmetrix specific term)

A business continuation process that creates an instant copy of single or multiple EMC Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system (s).

TLU

See Tape Library Unit.

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

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

transaction

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

transaction backup

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

transaction backup (Sybase and SQL specific term)

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

transaction log backup

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

transaction log files

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

transaction log table (Sybase specific term)

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

transportable snapshot (Microsoft VSS specific term)

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

U**unattended operation**

See lights-out operation.

user account (Data Protector user account)

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

User Account Control (UAC)

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

user disk quotas

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

user group

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

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

user profile (Windows specific term)

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

user rights

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

user_restrictions file

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

V**vaulting media**

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

verify

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

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

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

Virtual Device Interface (Microsoft SQL Server specific term)

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

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

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

virtual full backup

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

Virtual Library System (VLS)

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

virtual server

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

virtual tape

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

virtual tape library (VTL)

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

volser

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

volume group

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

volume mountpoint (Windows specific term)

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

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

Volume Shadow Copy Service

See Microsoft Volume Shadow Copy Service (VSS).

VSS

See Microsoft Volume Shadow Copy Service (VSS).

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

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

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

W

Wake ONLAN

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

Web reporting

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

wildcard character

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

Windows configuration backup

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

Windows Registry

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

WINS server

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

writer (Microsoft VSS specific term)

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

X**XBSA Interface (Informix Server specific term)**

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

Z**ZDB**

See zero downtime backup.

ZDB database (ZDB specific term)

A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore. See also zero downtime backup (ZDB).

ZDB to disk (ZDB specific term)

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process. See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape (ZDB specific term)

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore. See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape (ZDB specific term)

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

zero downtime backup (ZDB)

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

Index

A

architecture

- DB2 integration 49
- Informix integration 10
- Lotus integration 79

archived logs backups

- DB2 integration 54

B

backing up DB2 53

- archived logs backups 54
- backup modes 53
- backup options 58
- backup specification, modifying 58
- backup specifications, creating 54
- backup templates 55
- backup types 48
- database objects backups 54
- full backups 53
- incremental backups 53, 60
- incremental delta backups 60
- previewing backups 59
- scheduling backups 58
- scheduling backups, example 58
- starting backups 60

backing up Informix 18

- backup modes 30
- backup options 24
- backup specification, modifying 24
- backup specifications, creating 19

backup types 9, 18

- continuous backups 31
- full backups 9, 18
- incremental backups 9, 18
- manual backups 31
- onbar utility 30
- online mode 30
- previewing backups 26
- quiescent mode 30
- scheduling backups 25
- scheduling backups, example 25
- starting backups 28

backing up Lotus 89

- backup options 94
- backup specification, modifying 94
- backup specifications, creating 91
- backup types 78
- BOX files 90
- full backups 78, 89
- incremental backups 78, 90
- Lotus Domino Cluster 91
- Notes Storage Facility files 90
- Notes Template Facility files 90
- performance tuning 105
- previewing backups 95
- scheduling backups 94
- scheduling backups, example 94
- starting backups 96
- transaction log files 90

backup modes

- DB2 integration 53

- Informix integration 30
- backup options
 - DB2 integration 58
 - Informix integration 24
 - Lotus integration 94
- backup specifications, creating
 - DB2 integration 54
 - Informix integration 19
 - Lotus integration 91
- backup specifications, modifying
 - DB2 integration 58
 - Informix integration 24
 - Lotus integration 94
- backup templates
 - DB2 integration 55
- backup types
 - DB2 integration 48
 - Informix integration 9, 18
 - Lotus integration 78
- BOX files
 - Lotus integration 90
- C**
- checking configuration
 - DB2 integration 53
 - Informix integration 17
 - Lotus integration 88
- complete database restore, Informix integration 31
- concepts
 - DB2 integration 49
 - Informix integration 9
 - Lotus integration 79
- configuring DB2 50
 - checking configuration 53
- configuring Informix 11
 - checking configuration 17
- configuring Lotus 82
 - checking configuration 88
 - enabling transaction logging 84
- creating backup specifications
 - DB2 integration 54
 - Informix integration 19
 - Lotus integration 91
- D**
- DB2 backup 53
 - archived logs backups 54
 - backup modes 53
 - backup options 58
 - backup specification, modifying 58
 - backup specifications, creating 54
 - backup templates 55
 - backup types 48
 - database objects backups 54
 - delta backups 53
 - full backups 53
 - incremental backups 53, 60
 - incremental delta backups 60
 - modification tracking, enabling 60
 - previewing backups 59
 - scheduling backups 58
 - scheduling backups, example 58
 - starting backups 60

- DB2 configuration 50
 - checking configuration 53
- DB2 integration
 - architecture 49
 - backup 53
 - concepts 49
 - configuration 50
 - introduction 48
 - monitoring sessions 72
 - restore 61
 - troubleshooting 73
- DB2 restore 61
 - examples 67, 69
 - partitioned environment 70
 - restore options 65
 - to a new database 67
 - to another DB2 instance 67
 - using CLI 66
 - using GUI 62
- DB2 troubleshooting 73
- delta backups
 - DB2 integration 53
- E**
 - enabling modification tracking
 - DB2 integration 60
 - enabling transaction logging
 - Lotus integration 84
 - examples
 - DB2 integration, restore 67, 69
 - DB2 integration, scheduling backups 58
 - Informix integration, restore using onbar 39
 - Informix integration, scheduling backups 25
 - Informix integration, starting interactive backups 29
 - Lotus integration, restore 102
 - Lotus integration, scheduling backups 94
- F**
 - full backups
 - DB2 integration 53
 - Informix integration 9, 18
 - Lotus integration 78, 89
- I**
 - incremental backups
 - DB2 integration 53, 60
 - Informix integration 9, 18
 - Lotus integration 78, 90
 - incremental delta backups
 - DB2 integration 60
 - Informix backup 18
 - backup modes 30
 - backup options 24
 - backup specification, modifying 24
 - backup specifications, creating 19
 - backup types 9, 18
 - continuous backups 31
 - full backups 9, 18
 - incremental backups 9, 18
 - manual backups 31
 - onbar utility 30
 - online mode 30

- previewing backups 26
- quiescent mode 30
- scheduling backups 25
- scheduling backups, example 25
- starting backups 28
- Informix configuration 11
 - checking configuration 17
- Informix integration
 - architecture 10
 - backup 18
 - concepts 9
 - configuration 11
 - introduction 9
 - monitoring sessions 42
 - onbar utility 9
 - restore 31
 - troubleshooting 42
- Informix restore 31
 - complete database restore 31
 - finding information for restore 32
 - restore options 37
 - to another Informix Server 40
 - using another device 41
 - using CLI 38
 - using GUI 35
 - using Informix commands 39
 - using Informix commands, examples 39
 - whole-system restore 31
- Informix troubleshooting 42
- interactive backups
 - DB2 integration 60
 - Informix integration 28
 - Lotus integration 96
- introduction
 - DB2 integration 48
 - Informix integration 9
 - Lotus integration 78
- L**
- Lotus backup 89
 - backup options 94
 - backup specifications, creating 91
 - backup specifications, modifying 94
 - backup types 78
 - BOX files 90
 - full backups 78, 89
 - incremental backups 78, 90
 - Lotus Domino Cluster 91
 - Notes Storage Facility files 90
 - Notes Template Facility files 90
 - performance tuning 105
 - previewing backups 95
 - scheduling backups 94
 - scheduling backups, example 94
 - starting backups 96
 - transaction log files 90
- Lotus configuration 82
 - checking configuration 88
 - enabling transaction logging 84
- Lotus integration
 - architecture 79
 - backup 89
 - concepts 79

- configuration 82
 - introduction 78
 - monitoring sessions 105
 - restore 96
 - troubleshooting 106
- Lotus restore 96
- examples 102
 - finding information 97
 - restore options 102
 - using GUI 99
- Lotus troubleshooting 106
- M**
- modification tracking, enabling
- DB2 integration 60
- modifying backup specifications
- DB2 integration 58
 - Informix integration 24
 - Lotus integration 94
- monitoring sessions
- DB2 integration 72
 - Informix integration 42
 - Lotus integration 105
- N**
- Notes Storage Facility files
- Lotus integration 90
- Notes Template Facility files
- Lotus integration 90
- NSF 90
- NTF 90
- O**
- onbar utility
- Informix integration 9, 30
- online backups
- DB2 integration 48, 53
 - Informix integration 9, 30
 - Lotus integration 78
- P**
- performance tuning
- Lotus integration 105
- previewing backups
- DB2 integration 59
 - Informix integration 26
 - Lotus integration 95
- Q**
- quiescent backups
- Informix integration 30
- R**
- restore methods
- Informix integration 32
- restore options
- DB2 integration 65
 - Informix integration 37
 - Lotus integration 102
- restoring DB2 61
- examples 67, 69
 - partitioned environment 70
 - restore options 65
 - to a new database 67
 - to another DB2 instance 67
 - using CLI 66

- using GUI 62

restoring Informix 31

- complete database restore 31

- finding information for restore 32

- restore options 37

- to another Informix Server 40

- using another device 41

- using CLI 38

- using GUI 35

- using Informix commands 39

- using Informix commands, examples 39

- whole-system restore 31

restoring Lotus 96

- examples 102

- finding information 97

- restore options 102

- using GUI 99

running backups 60

S

scheduling backups

- DB2 integration 58

- Informix integration 25

- Lotus integration 94

starting backups

- DB2 integration 60

- Informix integration 28

- Lotus integration 96

T

transaction log files

- Lotus integration 90

transaction logging, enabling

- Lotus integration 84

troubleshooting DB2 73

troubleshooting Informix 42

troubleshooting Lotus 106

W

whole-system restore

- Informix integration 31

We appreciate your feedback!

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

Feedback on Integration guide for IBM Applications (Data Protector 8.10)

Just add your feedback to the email and click send.

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

