HP Data Protector 8.1 integration with HP 3PAR Storage System

(Zero Downtime Backup Solution)

Table of Contents

Executive summary

The growing requirement for data protection for mission critical applications, together with the increasing sophistication of Storage Area Network (SAN) environments, has resulted in a rapid expansion in the use of large disk arrays employing the RAID technology. These can hold large application databases containing vast amounts of data. By using storage virtualization techniques, disk arrays can be divided into many virtual disks. These can easily be copied within a disk array, perhaps many times dependent on the disk array technology and the available storage space. This makes it possible to perform operations on copies of data without any risk to the original data. In particular, it enables effective backup solutions for applications in high-availability and mission-critical areas.

If the Oracle database availability is one of the highest priorities, your backup strategy should include online backups that are frequently performed to minimize the recovery time. The HP Data Protector Zero Downtime Backup (ZDB) functionality offers online backup capabilities with minimal degradation of the application system performance.

Data Protector ZDB is designed to improve backup strategies for high-availability (HA) systems and non-HA systems. ZDB delivers best-in-class backup and recovery solutions for the Oracle database. ZDB and instant recovery (IR) have following advantages over other backup and restore techniques:

- Minimal downtime or impact on the application during backup
- Short restore times (minutes instead of hours), which are required to meet strict RTOs

Data Protector provides the ZDB integration module for HP 3PAR Storage Systems, which can be used in conjunction with the Data Protector Oracle integration to achieve a true zero impact backup.

Audience

This white paper is intended for solution architects, project managers, engineers, and support personnel involved in planning, designing, and configuring the Data Protector ZDB solution with HP 3PAR Storage Systems for the Oracle database.

You need to be familiar with:

- HP 3PAR Storage architecture
- Data Protector and its ZDB solution
- Oracle Real Application Clusters (RAC)

This white paper does not replace the standard product documentation.

Concepts

This white paper provides details on using the HP Data Protector 8.1 and HP 3PAR Storage System integration to perform zero downtime backup and instant recovery for the Oracle database. Data Protector provides this support using a native Storage Management Initiative-Specification (SMI-S) agent—HP P6000 / HP 3PAR SMI-S Agent.

The key concepts for this solution are described in the following sections.

Zero downtime backup

Conventional methods of backing up to tape are not well suited for large database applications; either the database has to be taken offline or, if the application allows it, put into “hot-backup mode” while data in it is streamed to tape.

The first can cause major disruption to the application's operation. The second can produce many large transaction log files, putting extra load on the application system.

Zero downtime backup (ZDB) uses disk array technology to minimize the disruption. In very general terms, a copy or **replica** of the data is created or maintained on a disk array. This is very fast and has little impact on the application's performance. The replica can itself form the backup, or it can be streamed to tape without further interruption to the application's use of the source database. Depending on the hardware and software with which it is created, a replica may be an exact duplicate (mirror, snapclone), or a virtual copy (snapshot) of the data being backed up.

In ZDB, **replication** (the process of creating or maintaining a replica) is the critical factor in minimizing interruption to the application.

Figure 1: Zero downtime backup and instant recovery concepts

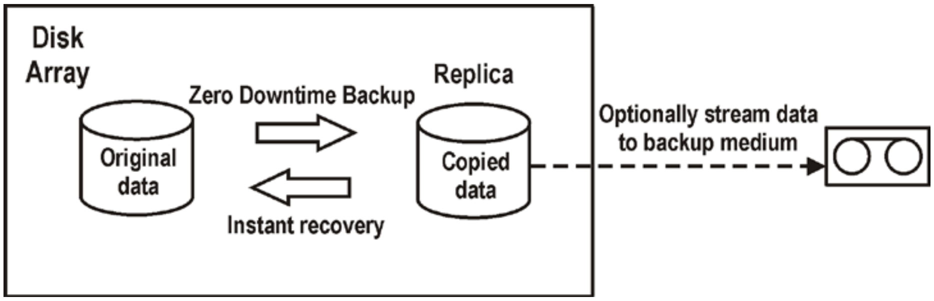

Instant recovery

Instant recovery requires a replica to exist on the same disk array to which the data is to be restored. Application and backup systems are disabled and the contents of the replica are restored directly to their original locations. During instant recovery, a data copy operation is performed in which data located on the source volume is replaced by data located on the target volumes. This operation is performed internally within the disk array, involving no other backup medium or device. This makes instant recovery very fast.

After the restore is completed, the sections of the database or filesystem concerned are returned to their states at the time the replica was created and the application system can be re-enabled. Depending on the application or database concerned, this may be all that is required. In some cases, additional action is required for full recovery, such as applying archived transaction log files that have been backed up separately.

With instant recovery, lost or corrupt data is replaced with known good data, which was previously created as a replica. Such data is handled on the complete storage volume level.

The remainder of the process depends on the application being recovered:

- Where a filesystem has been replicated, this step is all that is required to return the data to its state at the moment the replica was created.
- For a database application, you may need to perform additional operations to fully recover the database after performing instant recovery, such as restoring and applying transaction log files. In this way, you may be able to recover the database to a later point in time than that at which the replica was created, if log files for that time exist (commonly known as roll forward). This usually involves the use of another backup medium or device. For more information, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

In instant recovery sessions that use only Data Protector disk array agents, you cannot define which backup objects specified in the backup specification should be restored; only a complete backup object set can be selected for instant recovery and, hence, only the set of replicas can be restored. Additionally, on UNIX systems with configured LVM, not only the volumes constituting the replica are restored, but entire volume groups in which these volumes reside are returned to the state they were in when the replica was created.

Instant recovery process

The following is an example of instant recovery:

Figure 2: Instant recovery example

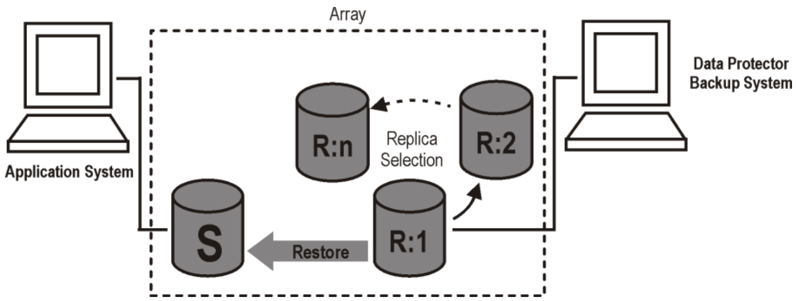

1. Decide which replica you want to restore and select the ZDB session that created it.
2. Select the instant recovery options, which are primarily provided for selecting the instant recovery method and the data safety level.
3. Optionally, perform a preview of the instant recovery session to provide an extra level of security.

   NOTE: Instant recovery preview is not available in instant recovery sessions that use the Data Protector Microsoft Volume Shadow Copy Service integration.

4. Start the instant recovery.

Data Protector then:

1. Starts processes on the application system and the backup system.
2. Extracts the session information from the IDB and the array-specific information associated with the session from the ZDB database.
3. Performs the necessary checks to verify that all required conditions for a successful instant recovery are met (including any instant recovery options specified).
4. Prepares the application system by deactivating any volume groups (on UNIX systems with configured LVM) and dismounts any filesystems associated with the replica.
5. Restores the original data.
6. Re-enables any volume groups that it disabled and re-mounts any filesystems that it dismounted.

After instant recovery, the contents of the source volumes are returned to the state they were in when the replica was created.

HP 3PAR native agent for ZDB and IR

To support ZDB and IR for 3PAR, Data Protector provides a native agent, which uses the SMI-S interface for controlling the disk array. The SMI-S agent creates snapshots (virtual copies)—specified in the backup specification—from the source volumes and then data is copied from snapshots to a (tape) device.

To integrate HP 3PAR StoreServ Storage with Data Protector, you need to install the HP P6000 / HP 3PAR SMI-S Agent component on the application and backup systems. In addition, to perform ZDB-to-disk+tape or ZDB-to-tape sessions, you need to install the General Media Agent component, regardless of the operating system.

Licensing

Data Protector supports two licensing schemes:

- **Traditional licensing** (based on features and backup targets): In this scheme, you need to purchase the ZDB license separately. For more information on this licensing scheme, see the Traditional licensing section.
- **Capacity-based licensing**: This scheme includes the ZDB license. That is, you do not have to purchase the ZDB license separately. For more information on this licensing scheme, see the Capacity-based licensing section.

Traditional licensing

Data Protector supports three different backup targets: snapshot, disk, and tape. Depending on which targets are used, you can license one or all targets as they can be combined. The product structure is modular and offers a lot of flexibility. You can select the license based on the Data Protector functionality that best meets your requirements.

The Data Protector product structure and licensing consists of three main categories:

- Starter Packs: A management server (Cell Manager) is supported on HP-UX, Windows, and Linux.
- Backup targets such as tape drive licenses, referred to as Drive Extensions. For one drive, advanced backup-to-disk and ZDB are licensed by capacity.
- Data Protector Functional Extensions: The functional extensions licenses are required once per instance (system, library, and terabyte) for online backup of databases and applications, the Manager-of-Managers functionality, libraries with more than 60 media slots, encryption, Instant Recovery, NDMP, and Granular Recovery Extension (GRE).

NOTE: The UNIX product licenses operate on the UNIX, Windows, and Novell NetWare platforms, supporting the functionality regardless of the platform. However, the Windows product licenses operate on the Windows, Novell NetWare, and Linux platforms only. Passwords are bound to the Cell Manager and are valid for the entire Data Protector cell. Clients do not require any license for file system or disk image backups.

Capacity-based licensing

The capacity-based licensing is based on the volume of primary data protected by Data Protector and includes unlimited use of enterprise protection features. The capacity is measured in "Front End Terabytes" or “Front End TB”.

The total amount of Front End Terabytes is defined as the aggregate amount of data from all systems being backed up. Per system it is measured as the largest full (that is, the total amount of source data protected).

The licensing in this product structure is perpetual. That is, it covers all existing or new servers, storage, applications, and so on.

The following features are included in the capacity-based license:

- Cell Managers and Manager-of-Managers
- Tape Drives and Libraries
- Online Backup and Granular Recovery Extensions
- Zero Downtime Backup and Instant Recovery
- Advanced Backup to Disk and NDMP

The following licenses are not included in this product structure and must be ordered separately:

- Software encryption

Solution description

This solution involves the HP Data Protector 8.1 and HP 3PAR Storage Systems integration to perform backup and recovery of the Oracle database.

The setup used for the solution consists of:

- Application server on which the application is running. The application data is stored on the HP 3PAR Storage System and is accessible through SAN.
- HP 3PAR Storage System provides the storage for the application.
- Backup server is used to move data from replica to backup device (for example, tape).

Figure 3: Solution setup

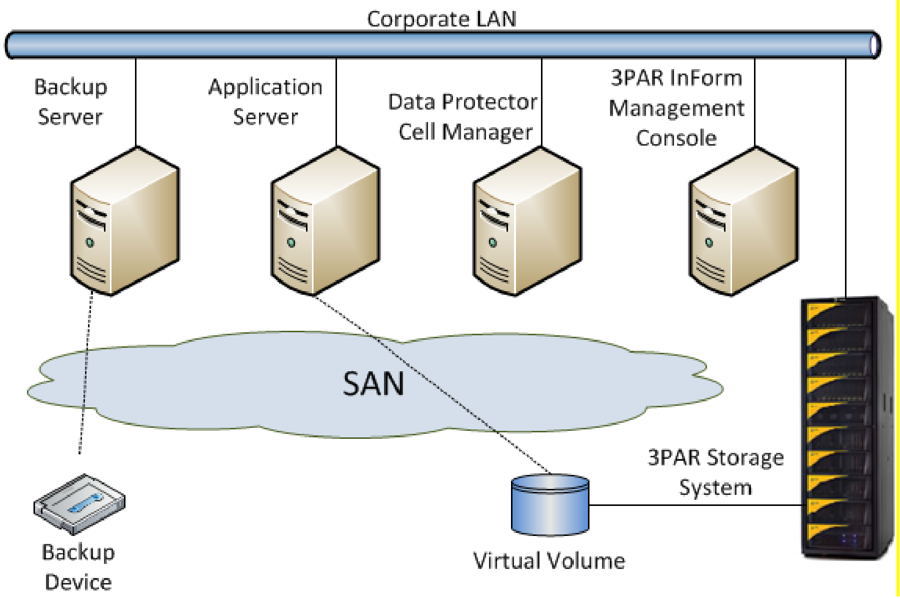

- Both application and backup servers have Data Protector client software installed. Data Protector Cell Manager is installed on one of the servers. Application and backup systems are the clients of Data Protector Cell Manager. Server on which Data Protector Cell Manager is installed does not require SAN access.
- 3PAR InForm Management Console is installed on one of the servers and manages the HP 3PAR Storage System.

Once the backup session is started, Data Protector will resolve the volumes involved in the backup. During this phase, Data Protector finds out on which storage volumes the selected objects are located. Finally, it requests the creation of a snapshot for each of those volumes by using the HP

3PAR SMI-S provider. This copy is a replica or a virtual copy snapshot in 3PAR terms. The 3PAR Storage System maintains integrity of the data on the snapshots using copy-on-write technology. In the following figure, creating a virtual copy is shown by the red line.

Figure 4: Creation of virtual copy snapshot

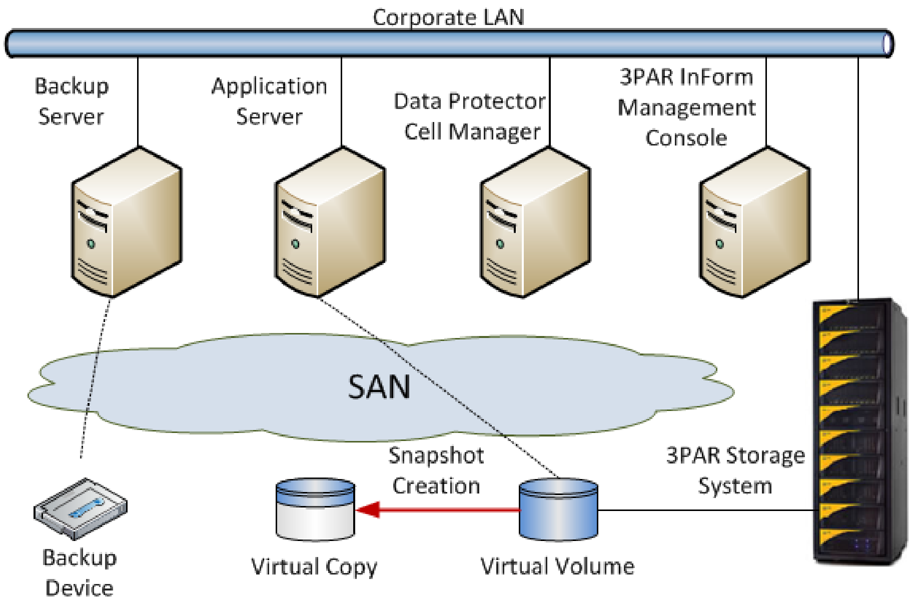

Solution workflow

The following workflow illustrates the key tasks required to perform zero downtime backup and instant recovery for the Oracle database using the HP 3PAR StoreServ Storage integration with Data Protector.

Figure 5: Solution workflow

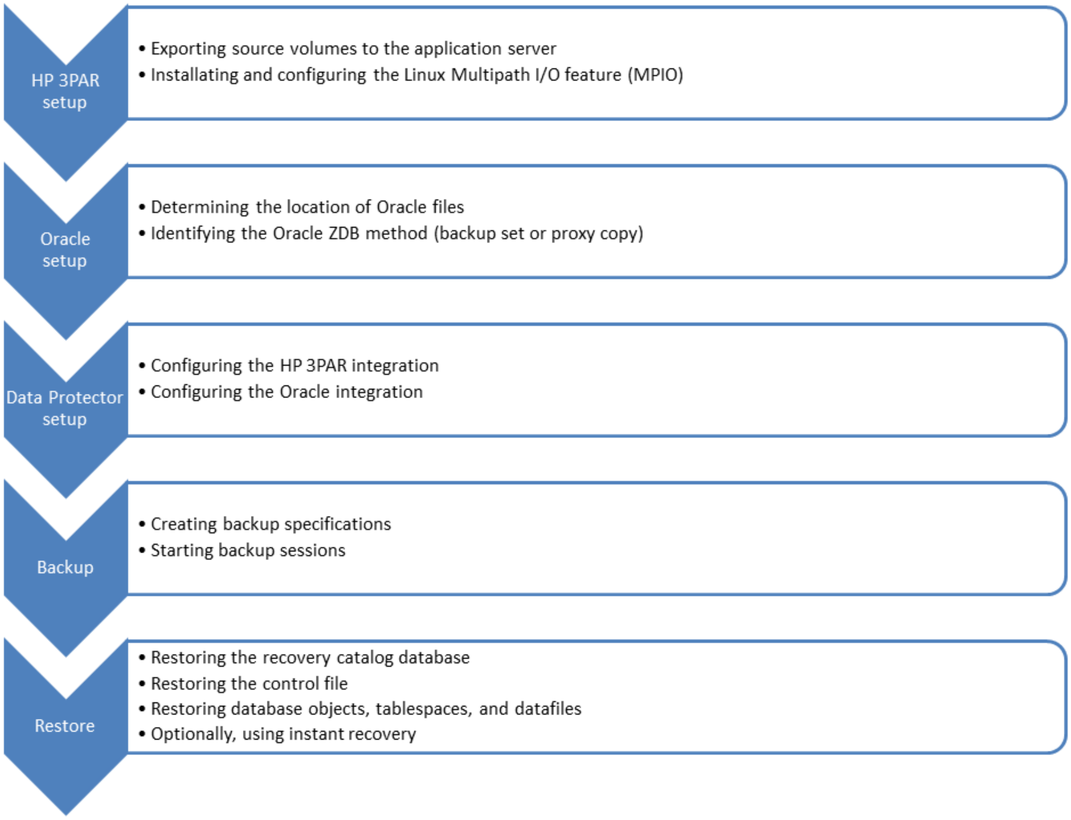

HP 3PAR setup

- Exporting source volumes to the application server
- Installating and configuring the Linux Multipath I/O feature (MPIO)

Oracle setup

- Determining the location of Oracle files
- Identifying the Oracle ZDB method (backup set or proxy copy)

Data Protector setup

- Configuring the HP 3PAR integration
- Configuring the Oracle integration

Backup

- Creating backup specifications
- Starting backup sessions

Restore

- Restoring the recovery catalog database
- Restoring the control file
- Restoring database objects, tablespaces, and datafiles
- Optionally, using instant recovery

HP 3PAR storage system setup

To prepare an HP 3PAR Storage System for integration with Data Protector, perform the following steps:

- Export source volumes to the application server
- Install and configure Linux MPIO with 3PAR

Exporting source volumes to the application server

Prior to starting with exporting of source Virtual Volumes to the application, ensure that the following has been correctly configured:

- All application and backup servers have been properly zoned with the 3PAR Storage System
- Host objects for all application and backup systems have been created and configured. When configuring host objects, select the operating system to populate Persona or environment (see Table **2**):

Table 1: 3PAR Storage System Host Configuration

Operating System	**Persona**
Red Hat Enterprise Linux	1 – Generic
Solaris 9/10	
SuSE	
Windows 2003	
Citrix Xen Server 5.x/6.x	
Solaris 11	2 – Generic-ALUA
Windows 2008/2008 R2	
Exanet	6 - Generic-legacy
HP-UX	7 – HPUX-legacy
IBM VIO Server	8 – AIX-legacy
AIX	
Egenera	9 – EGENERA
NetApp/ONTAP	10 – ONTAP-legacy
ESX 4.x/5.x	11 – Vmware
OpenVMS	12 – OpenVMS

If the application host is running as a Guest OS in either Microsoft Hyper-V or VMware EXS environment, and the LUNs are accessed from the Guest OS as Pass through (Hyper-V) or Raw Disk Mapped (VMware) disk, the above rules apply to that host too.

If the virtual volume will be used by either an application running in an MS cluster or as a Clustered Shared Volume, it has to be exported to all clients that are part of the cluster, or that are configured as the Hyper-V server.

Installating and configuring the Linux Multipath I/O feature (MPIO)

Prerequisite

Install appropriate multi-path device management software on the application system and the backup system.

Linux systems: HP Device Mapper Multipath Enablement Kit for HP Disk Arrays 4.2.0 or newer version.

To configure MPIO for 3PAR Storage System, perform the following steps:

1. Start the multipath daemon.
2. Run the following command for configuring the daemon to start during the system startup:
   Red Hat Enterprise Linux: `chkconfig multipathd on`
   SUSE Linux Enterprise Server: `chkconfig boot.multipath on`
3. Navigate to the `/etc/multipath.conf` file and add the following line to the `defaults` section of the file:

   ```
   no_path_retry      fail
   ```

This prevents the multipath device management software from queuing for unavailable disk volumes.

Ensure that the `no_path_retry` parameter value is not overridden by analogous entries in the device sections of the same file in which the corresponding disk array storage systems are configured.

Oracle database setup

To get an Oracle zero downtime backup work, the datafiles must be located onto volumes of a disk array. For HP EVA and HP 3PAR, these volumes need to be replicated using the disk array-based mirroring or snapshot facilities.

The following files have to be stored on a different physical volume/volume group than the datafiles. If they were, an Instant Recovery would overwrite them thus preventing a complete roll-forward recovery.

- redo log files
- archive log files
- control files

These files may either reside on local disks of the application hosts or other disks/volumes on the disk array, which are not replicated during Data Protector ZDB processes.

The archive log directory is continuously accessed by the database in terms of copying current redo log files over to the archive directory. Therefore, the archive log directory cannot be unmounted, even for the short time of the split or creation of snapshots. Due to this mount limitation, the archive logs can never be brought into a consistent state to perform a split or create a snapshot/snapclone. The only way to back up the archive logs is to back them up locally on the application host using a non-ZDB backup specification.

The control file is a very important file for the database; therefore, you should not overwrite it during restore or instant recovery. It should have its own location, which is not affected by an IR (disk restore).

Determining the location of Oracle files

The following Oracle database components need to be placed on volumes different than the datafile volume:

- redo log files
- archive log files
- control files

To determine the location of these components, execute SQL queries on the database using `sqlplus`. As an Oracle user, run the following commands on an Oracle database.

NOTE: The database needs to be in the mounted/open mode.

```
$ sqlplus "/ as sysdba"

      SQL*Plus: Release 12.1.0.1.0 Production on Tue Nov 26 13:00:14 2013

      Copyright (c) 1982, 2013, Oracle.  All rights reserved.

      Connected to:
      Oracle Database 12c Enterprise Edition Release 12.1.0.1.0 - 64bit
      Production
      With the Partitioning, OLAP, Advanced Analytics and Real
      Application Testing options

SQL> select INSTANCE_NAME, STATUS from v$instance;

      INSTANCE_NAME    STATUS
      ---------------- ------------ -------------------------------------
      APPN             OPEN
```

Determine the location of the redo log files.

```
SQL> select member from v$logfile;
      MEMBER
      -------------------------------------------------------------------
      /LOG/redo03.log
      /LOG/redo02.log
      /LOG/redo01.log
```

Determine the location of the archive log file. First, check the `Archive destination` entry for the location. If the archive destination is USE_DB_RECOVERY_FILE_DEST, check the setting of that parameter by using the `show parameter` command.

```
SQL> ARCHIVE LOG LIST

      Database log mode              Archive Mode
      Automatic archival             Enabled
      Archive destination            /LOG/
      Oldest online log sequence     65
      Next log sequence to archive   67
      Current log sequence           67
```

Determine the location of the control file(s).

```
SQL> SELECT name FROM v$controlfile;

      NAME
      -------------------------------------------------------------------
      /DATA/control01.ctl
      /DATA/control02.ctl
```

Determine the location of the Oracle parameter file (pfile) or server parameter file (spfile).

```
SQL> show parameter pfile;

      NAME               TYPE        VALUE
      ---------------    ----------- ------------------
      spfile             string      /orahome/oracle/dbs/spfileAPPN
                                     .ora
```

Oracle ZDB methods

Oracle databases can be backed up by user-managed scripts or by Oracle Recovery Manager (RMAN), which is a backup, restore, and recovery tool. RMAN offers two methods to perform ZDB:

- Backup set
- Proxy copy

Backup set ZDB

A backup set, which is a logical unit, contains one or more Oracle backup objects. Backup objects are operating system files like data files, control files, or archived redo logs. With the Oracle backup set ZDB method, the entire data to be backed up is provided to Data Protector through the Oracle API—the data is streamed through the Data Protector Oracle integration Media Management Library (MML).

To perform the Oracle ZDB with backup sets, the Oracle software must run on the backup server. After the replicas are created, an Oracle instance of the production database is started on the backup server and a normal Oracle online backup is performed.

For more information on the Oracle backup set ZDB method, see *HP Data Protector 8.1 Zero Downtime Backup Integration Guide*.

Proxy copy ZDB

A proxy copy is a special type of backup in which RMAN turns over the control of the data transfer to a media manager that supports this feature. The PROXY option of the BACKUP command specifies that a backup should be a proxy copy. For each file that you attempt to back up using the BACKUP PROXY command, RMAN queries the media manager to determine whether it can perform a proxy copy. If the media manager cannot proxy copy the file, RMAN uses conventional backup sets to perform the backup. An exception occurs when you use the PROXY ONLY option, which causes Oracle to issue an error message when it cannot proxy copy. Oracle records each proxy-copied file in the control file. RMAN uses this data to resynchronize the recovery catalog. You can use the `V$PROXY_DATAFILE` view to obtain the proxy copy information.

To perform a proxy copy ZDB, you do not have to install Oracle on the backup server.
For more information on the Oracle proxy copy ZDB method, see *HP Data Protector 8.1 Zero Downtime Backup Integration Guide*.

Data Protector setup

This section provides information on configuring HP 3PAR and Oracle integrations with Data Protector.

Installing the HP 3PAR SMI-S Agent component

You must install the HP P6000 / HP 3PAR SMI-S Agent component on the application and backup servers (see Figure 6). In addition, you must install:

- The General Media Agent component on each server that acts as the backup server.
- The Data Protector Oracle integration on the application and backup servers. For more information, see the Configuring the Oracle and Data Protector integration section.

NOTE: The application server can also double as the backup server. In such a case, both the HP P6000 / HP 3PAR SMI-S Agent and General Media Agent components must be installed on the application server.

Limitation: On Linux, only two host configurations are supported. That is, application and backup servers must be running on different machines.

Figure 6: Adding HP 3PAR Agent and General Media Agent components

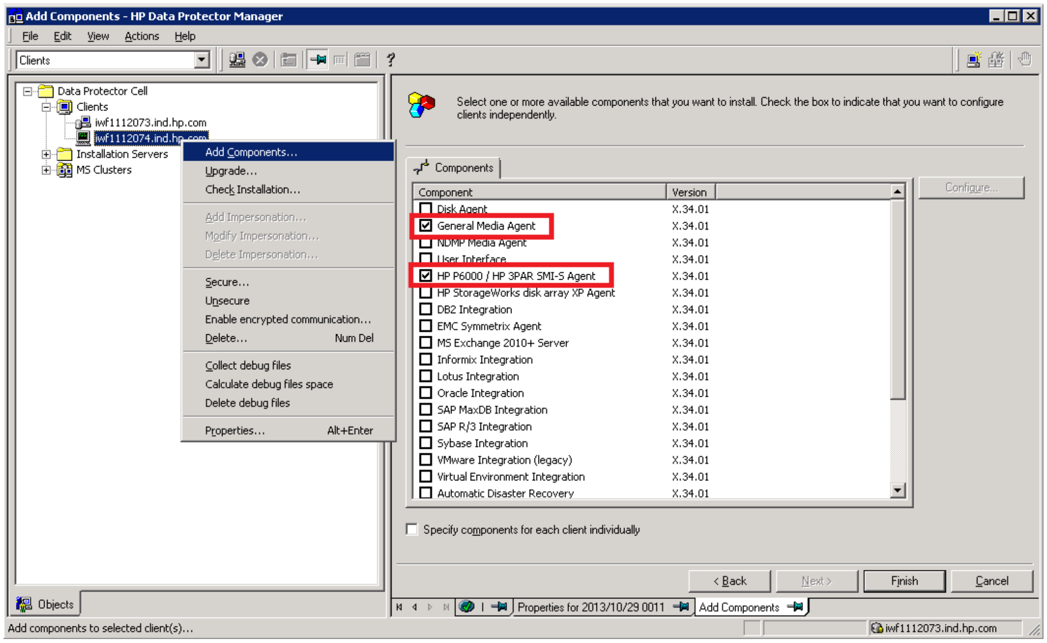

Select one or more available components that you want to install. Check the box to indicate that you want to configure clients independently.

Component	Version
☐ Disk Agent	X.34.01
☑ General Media Agent	X.34.01
☐ NDMP Media Agent	X.34.01
☐ User Interface	X.34.01
☑ HP P6000 / HP 3PAR SMI-S Agent	X.34.01
☐ HP StorageWorks disk array XP Agent	X.34.01
☐ DB2 Integration	X.34.01
☐ EMC Symmetrix Agent	X.34.01
☐ MS Exchange 2010+ Server	X.34.01
☐ Informix Integration	X.34.01
☐ Lotus Integration	X.34.01
☐ Oracle Integration	X.34.01
☐ SAP MaxDB Integration	X.34.01
☐ SAP R/3 Integration	X.34.01
☐ Sybase Integration	X.34.01
☐ VMware Integration (legacy)	X.34.01
☐ Virtual Environment Integration	X.34.01
☐ Automatic Disaster Recovery	X.34.01

☐ Specify components for each client individually

Add components to selected client(s)...

Configuring the HP 3PAR StoreServ Storage integration

Before you start with the configuration, ensure that all prerequisites are fulfilled.

Prerequisites

- Install Data Protector licenses and components:
    - Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
    - Install the following Data Protector component on the application system and the backup system: HP P6000 / HP 3PAR SMI-S Agent

For licensing information and installation and upgrade instructions, see the HP Data Protector Installation and Licensing Guide.

- Make sure the same operating system version is installed on the application system and the backup system.
- If the application system and the backup system reside in a Data Protector cell with secured clients, ensure that access between both systems is allowed in both directions.
- Source volumes must have snapshot space (copy space) in a storage system's Common Provisioning Group (CPG) associated with.
- HP 3PAR CIM service should be up and running on the InForm, as it may not be running after the installation of InForm.
    - Use the following HP 3PAR CLI commands to manage the CIM server: `setcim, startcim, stopcim, and showcim`
- Install a supported version of the InForm OS.
    - InForm OS Version 3.1.2 and 3.1.2 MU2 are supported.
    - For ZDB/IR with Oracle ASM, InForm OS Version 3.1.2 MU2 is a pre-requisite.
- Ensure that the usual ports for the SMI-S communication are open:
    - Default HTTP port: 5988
    - Default HTTPS port: 5989

To prepare the Data Protector HP 3PAR StoreServ Storage integration for use with a storage system of the HP 3PAR StoreServ Storage family, you must perform the mandatory configuration step.

Connection configuration data

To be able to connect to a CIMOM provider and perform ZDB-to-disk, ZDB-to-disk+tape, and IR sessions, the Data Protector HP P6000 / HP 3PAR SMI-S Agent needs the following information:

- Fully qualified domain name or IP address of the system where the CIMOM service is running

  In case the system has multiple IP addresses configured, the address by which the system can be accessed by the Data Protector ZDB agent should be used.

- Whether the connection uses Secure Sockets Layer (SSL)
- Port number of the port on which the CIMOM service is accepting requests
- Username and password

  These credentials must belong to a 3PAR StoreServ system user account with the *Edit* or *Super* privilege level in the following 3PAR StoreServ system virtual domains, depending on the effective disk array configuration:

    - Domain of the application system and the source volumes—When the source volumes and the application system belong to a specific domain
    - All domains of a domain set—When the application system and the source volumes belong to this domain set
    - All existing domains—When the application system and the source volumes do not belong to any domain

For more information on using the HP 3PAR StoreServ Storage authorization system, see the HP 3PAR StoreServ Storage documentation.

You need to provide the preceding information in advance for each CIMOM provider to which the Data Protector HP P6000 / HP 3PAR SMI-S Agent should connect. It is stored in the HP 3PAR StoreServ Storage part of the SMISDB.

Configuration procedure

To add the required user credentials for an application system where the CIMOM service is running, use the Data Protector `omnidbzdb` command. Proceed as follows:

1. Identify the source volumes that will be involved in the ZDB-to-disk or ZDB-to-disk+tape sessions.
2. Identify the 3PAR StoreServ system virtual domains or domain set to which the application system and the source volumes belong.
3. Choose a disk array user account that has a proper privilege level on the corresponding domains. Identify and write down its username and password that you will need in the next step.
4. Using the `omnidbzdb --diskarray 3PAR --ompasswd --add` command, add the username and password that you acquired in the previous step to the ZDB database, providing the name of the application system you identified in Step 1 of this procedure.

   For command syntax and usage examples, see the `omnidbzdb` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidbzdb` man page.

5. Using the `omnidbzdb --diskarray 3PAR --ompasswd --check` command, verify that the HP P6000 / HP 3PAR SMI-S Agent can connect to the disk array using the configured user authentication data.
6. Using the `omnidbzdb --diskarray 3PAR --ompasswd --list <IP address>` command list the connection configuration data for connections to the CIMOM providers available to the system with the specified IP address.

TIP: For each application system, you can add user credentials of multiple disk array user accounts. When several are configured for the same system, the HP 3PAR Agent checks user accounts in alphabetical order and uses the first account with sufficient privileges: the *Edit* or *Super* privilege level on the application system and the source volumes.

For information on performing other tasks related to management of user credentials in the ZDB database, see the `omnidbzdb` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidbzdb` man page.

HP 3PAR user requirements and supported combinations

For the HP 3PAR native agent, you must configure at least one user per HP 3PAR array. You need to consider the following requirements while configuring the user(s):

- Users with the “browse” privilege are not supported.
- If the source volumes and the application host are not in a domain, the user must have “edit” rights on all domains. In such a case, only the single user should be configured for the HP 3PAR disk array.
- If the source volumes and the application host belong to a domain, the user needs at least “edit” rights for this specific domain.
- Hosts can also be part of a 3PAR domain set. This allows a host to be accessed from multiple domains. In such a case, the user needs to have at least “edit” rights on all domains contained in the domain set.

- If multiple users are configured for one HP 3PAR disk array, the agent attempts to locate the user who can access (see) the volumes. For each disk array, users are verified in the alphabetical order, and the first user who can "see" the disk is used.

The following table lists the supported combinations (highlighted in green). Yellow rows are not supported, as the credentials for the HP 3PAR/P10000 CIMOM provider connection configuration must not be a user account with the Super privilege level; instead, it must be the Edit privilege level.

Table 2: Supported combinations for HP 3PAR users

User	Privilege	Volume	Hosts	Backup	IR
Domain A	Edit	Domain A	Domain A	Success	Success
Domain A	Edit	Domain B	Domain B	FAILED	FAILED
Domain A	Edit	Domain A	Domain B	FAILED	FAILED
Domain A	Edit	No Domain	No Domain	FAILED	FAILED
Domain A	Browse	Domain A	Domain A	Success	FAILED
No Domain	super	Domain A	Domain A	Success	Success
No Domain	Super	Domain B	Domain A	FAILED	FAILED
No Domain	super	No Domain	No Domain	Success	Success
No Domain	Edit	Domain A	Domain A	Success	Success
No Domain	Edit	Domain B	Domain A	FAILED	FAILED
No Domain	Edit	No Domain	No Domain	Success	Success
Domain A Domain B	Edit none	Domain A	Domain Set (A,B)	Success	FAILED
Domain A Domain B	None edit	Domain A	Domain Set (A,B)	FAILED	FAILED
Domain A Domain B	Edit none	Domain B	Domain Set (A,B)	FAILED	FAILED
Domain A Domain B	None edit	Domain B	Domain Set (A,B)	Success	FAILED
Domain A Domain B	Edit edit	Domain A	Domain Set (A,B)	Success	Success
Domain A Domain B	Edit edit	Domain B	Domain Set (A,B)	Success	Success
No Domain	super	Domain Set (A,B)	Domain Set (A,B)	Success	Success

For more information on different user rights, see the HP 3PAR disk array documentation.

Data Protector Oracle Server ZDB integration

You can employ a variety of backup strategies to best meet your system priorities. If database availability is the highest priority, your backup strategy should include online backups that are performed frequently to minimize the recovery time.

The advantages of using the Data Protector Oracle ZDB integration are as follows:

- ZDB reduces the performance degradation of the application system.
- The tablespaces are in the backup mode (online backup) or the database is shut down (offline backup) only during the short period required to create a replica (split the mirror disks or create snapshots).
- The load to the application system is significantly reduced. Following the replica creation, tape backup can be started on the copied data using a separate backup system.

The Data Protector Oracle ZDB integration offers online and offline backup of your Oracle Server System (application system). The online backup concept is widely used since it enables high application availability. Offline backup requires shutting down the database while creating a replica and, therefore, does not offer high availability.

ZDB methods and Oracle versions

The installation, upgrade, configuration, and parts of backup flow are different depending on the selected Oracle ZDB method. The procedures for configuration of backup specifications and starting or scheduling backups are the same, regardless of the Oracle ZDB method.

Backup and restore types

Backup

Using Data Protector, you can perform the following types of backup:

- Online ZDB to disk, ZDB to tape, and ZDB to disk+tape.
  During the creation of a replica, the database on the application system is in hot backup mode. If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the streaming of the data to tape media is subsequently performed on the backup system.

- Offline ZDB to disk, ZDB to tape, and ZDB to disk+tape.
  During the creation of a replica, the database is shut down on the application system. Therefore, the database is not available during the short time that it takes to create the replica. If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the streaming of the data to tape media is subsequently performed on the backup system.

With both online and offline ZDB to tape or ZDB to disk+tape, a standard Data Protector (non-ZDB) backup of the recovery catalog and the control file is started automatically, after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTES:

- Backup of the recovery catalog and control file is not performed with ZDB to disk. The Oracle Recovery Manager utility (RMAN) is not aware of ZDB-to-disk sessions.
- Backup of archived logs cannot be done with the Data Protector Oracle ZDB integration. Backup of archive logs and control file has to be done following the standard Data Protector Oracle integration backup procedure.

Restore

Using Data Protector and the disk array integrations, you can perform the following types of restore:

- Restoring from backup media to the application system on LAN (standard Data Protector restore) and using RMAN on the application system, you can:
  - recover a whole database
  - recover a part of a database
  - recover a whole database as it was at a specific point in time
- Restoring data using instant recovery. Note that the database recovery is done using the incremental backups performed using the standard Data Protector Oracle integration backup.

Integration concepts

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

Components

The software components involved in backup and restore processes are:

- The Oracle Recovery Manager (RMAN)
- The Data Protector Oracle integration software

Integration functionality overview

The Data Protector Oracle Integration agent (ob2rman.pl) works with RMAN to manage all aspects of the following operations on the Oracle target database:

- Database startup and shutdown
- Backups (backup and copy)
- Recovery (restore, recovery, and duplication)

For more information on the Data Protector Oracle integration concepts (such as how integration works, non-ZDB flow, and restore flow), see *HP Data Protector 8.1 Zero Downtime Backup Integration Guide*.

Configuring the Oracle and Data Protector integration

Prerequisites

- Oracle Server and Data Protector client systems must be correctly installed and configured. See the:
    - Latest support matrices at http://support.openview.hp.com/selfsolve/manuals for an up-to-date list of supported versions, platforms, devices, and other information.
    - *HP Data Protector Installation and Licensing Guide* for instructions on installing Data Protector on various architectures and on installing the Data Protector 3PAR integration with Oracle.
    - *Oracle Recovery Manager User's Guide and References* for Oracle concepts and backup/recovery strategies.
    - *Oracle Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle backup terminology and concepts.
    - *Oracle Enterprise Manager User's Guide* for information on backup and recovery with the Oracle Enterprise Manager, as well as information about SQL*Plus.
- The Data Protector 3PAR integration must be correctly installed and configured. For installation, see the *HP Data Protector Installation and Licensing Guide*. For configuration, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.
- ***Oracle Server configurations with ASM:*** The disk array must support creation of replicas with cross-volume data consistency.
- The Oracle Server software must be installed on the application system and the Oracle target database must be open or mounted there.
- The Oracle recovery catalog database must be properly configured and open.
- Oracle net services must be properly configured and running (on the application system) for the Oracle target database and the recovery catalog. The net services are needed for the Data Protector Oracle agent to be connected to the Oracle database on the application system through Oracle.

For more information about different connection options, see the *Oracle Recovery Manager User's Guide and References*.

NOTE: The Data Protector Oracle integration uses RMAN for backup and restore. RMAN connection to a target database requires a dedicated server process. To ensure that RMAN does not connect to a dispatcher when the target database is configured for a shared server, the net service name used by RMAN must include `(SERVER_DEDICATED)` in the `CONNECT_DATA` attribute of the connection string.

- ***Oracle Real Application Clusters (RAC):*** Each node must have a dedicated disk for storing archive logs. Such disks must be NFS mounted on all other RAC nodes. However, if the archive logs are not on an NFS mounted disk, you must modify the archive log backup specification.
- ***RAC:*** The control file must be created on a shared disk and be accessible from all RAC nodes, and the `OB2_DPMCTL_SHRLOC` environment variable must point to this location.

Before you begin

- Test whether the Oracle Server system and the Cell Manager communicate properly. Configure and run a Data Protector filesystem backup and restore on the Oracle Server system.
- Identify the Oracle database user that will be used by Data Protector for backup. This user must have the `SYSDBA` privilege granted. For example, it could be the Oracle user `sys`, which is created during database creation. See the Oracle documentation for more information on user privileges in Oracle.
- On Windows systems, if the Oracle target database and the Oracle recovery catalog are installed on two different systems, configure a *domain* user account that is a member of the Administrators group on both systems.
  On Windows Server 2003 systems with the Oracle target database installed, you need to restart the `Data Protector Inet` service under a Windows domain user account that has the appropriate Oracle database permissions for running backups and restores.
  For information on how to change the `Data Protector Inet` service account, see the *HP Data Protector Help* index: "Inet, changing account".
  However, for other supported Windows operating systems, you can use user impersonation instead. For details on setting accounts for the Inet service user impersonation, see the *HP Data Protector Help* index: "Inet user impersonation".
- In case of the backup set method, if the Oracle database is installed on symbolic links, create these symbolic links on the backup system, too.
- From the application system, using SQL*Plus, connect to the target database and recovery catalog by specifying the user, password, and net connect identifier. Connect to the target database as the database administrator and to the recovery catalog database as the recovery catalog owner.

  Example

  If the user name for the target database is `system`, password `manager`, net service name `PROD`, and the user name and password for the recovery catalog is `rman` and the net service name `RMANCAT`, then the commands will be as follows:

```
sqlplus /nolog
SQL> connect system/manager@PROD as sysdba;
Connected.
SQL> connect rman/rman@RMANCAT;
Connected.
```

- For online backup only, enable the Oracle automatic log archiving:

  1. Shut down the Oracle target database instance on the application system.
  2. Back up the entire database using a filesystem backup.
  3. Select the location for archive logs:

     - If SPFILE is used, run:
       ```
       alter system set log_archive_dest=path_to_archive_logs
       SCOPE=SPFILE;
       ```

     - If the init.ora file is used, run:
       ```
       log_archive_start=true
       log_archive_dest=path_to_archive_logs
       ```

       The default path of the file is:
       Windows systems: `ORACLE_HOME\database\initDB_NAME.ora`
       UNIX systems: `ORACLE_HOME/dbs/initDB_NAME.ora`

       where, `DB_NAME` is the name of the Oracle database instance.

  4. Mount the target database. To enable the archive log mode, start SQL*Plus and type:

     ```
     startup mount;
     alter database archivelog;
     alter database open;
     ```

     Example
     If the user name for the target database is `system`, password `manager`, instance name `PROD`, and the user name and password for the recovery catalog is `rman`, then the commands will look like:

     ```
     sqlplus /nolog
     SQL> connect system/manager@PROD as sysdba;
     Connected.
     SQL> startup mount;
     SQL> alter database archivelog;
     Statement processed.
     SQL> archive log start;
     Statement processed.
     SQL> alter database open;
     ```

  5. Back up the entire database.

Backup set method

For the Oracle backup set method, you should:

- Ensure that the Oracle software on the backup system and application system have the same directory structure. That means that ORACLE_HOME for both Oracle installations has to be identical.
- Ensure that the following files are the same on the application system and the backup system. In addition, check that the permissions are identical as on the application system:
  - `names.ora`
    Default path: `ORACLE_HOME/network/admin/names.ora`

  - `initDB_NAME.ora`
    Default path: `ORACLE_HOME/dbs/initDB_NAME.ora`

- `orapwDB_NAME`
  Default path: `ORACLE_HOME/dbs/orapwDB_NAME`
- `admin/DB_NAME`
  Default path: `ORACLE_BASE/admin/DB_NAME`

- Ensure that the Oracle net services on the application system and the backup system have the same directory structure. This can be accomplished by either NFS sharing of the files, manually copying the files from the application system to the backup system, or by using the UNIX `rdist` or `tar` commands to distribute the files from the application system.
- Test whether the Oracle user can log in to the Oracle target database as the Oracle database administrator and to the Oracle recovery catalog database as the Oracle recovery catalog owner from the backup system:

  1. Export `ORACLE_HOME`, `DB_NAME`, and on UNIX systems also `SHLIB_PATH` variables.
  2. Using SQL*Plus, connect to the Oracle recovery catalog database by specifying the user (recovery catalog owner), password, and net connect identifier.
  3. Connect to the Oracle target database locally using the Oracle Net software as the Oracle database administrator with the `SYSDBA` role.

     Example
     If the DB_NAME of the target database is `PROD`, the `DB_NAME` of the Oracle recovery catalog database is `RMANCAT`, and `ORACLE_HOME` is `/oracle/PROD`, then the commands will be:

     ```
     root# su - ora
     ora# id
     uid=101(ora) gid=101(dba)

     ora# export DB_NAME=PROD
     ora# oracle/PROD/bin/sqlplus
     SQL> connect rman/rman@RMANCAT
     Connected.

     SQL> connect system/manager as sysdba
     SQL> connect system/manager@PROD as sysdba;
     Connected.
     ```

- Test whether the user `root` and the Oracle administrator (for example, the user `oracle`) can connect to the target database and the recovery catalog database using the RMAN command on the backup system:

  1. Log on as the Oracle database administrator to the backup system (for example, the user `oracle`).
  2. Execute the RMAN command and connect to the target database and the recovery catalog database.

     Example
     If the DB_NAME of the target database is `PROD`, the `DB_NAME` of the Oracle recovery catalog database is `RMANCAT`, and `ORACLE_HOME` is `/oracle/PROD`, then the commands will be:

     ```
     root# su – ora
     ora# id
     uid=101(ora) gid=101(dba)
     ```

```
ora# export DB_NAME=PROD
ora# rman target system/manager catalog rman/rman
Recovery Manager: Release x.x.x.x.x - Production
RMAN-06005: connected to target database: PROD
RMAN-06008: connected to recovery catalog database
RMAN> exit
Recovery Manager completed.
```

Configuring Oracle user accounts

Decide under which user accounts you want backups to run. Data Protector requires the following user accounts:

- Oracle operating system user account (see the Configuring Oracle operating system user accounts section)
- Oracle database user accounts

Configuring Oracle operating system user accounts

For each Oracle database, Data Protector requires an operating system user account that has Oracle rights to back up the database. This user account usually belongs to the DBA user group (OSDBA user). The user account under which the Oracle database is running has these rights. For example, to find such a user on UNIX systems, run:

```
ps -ef|grep ora_pmon_DB_NAME
```

or

```
ps -ef|grep ora_lgwr_DB_NAME
```

The following table explains how to configure users on different operating systems:

Client system	Description
UNIX system	Ensure that the Oracle user oracle from the Oracle Inventory group (`oinstall`) has been added to the Data Protector admin user group. For details on adding users, see the *HP Data Protector Help* index: "adding users". Add the `OSDBA` user account and `root` user account from both the application and backup systems to the Data Protector `admin` or `operator` user group. The `OSDBA` user on the backup system must have the same numerical user ID and group ID as the `OSDBA` user on the application system (for example, `uid=101(ora) gid=101(dba)`). **TIP:** To find the user ID, connect to a system under this user account and run: `#id`
Windows system	On Windows systems, Data Protector connects to the Oracle database using the Data Protector `Inet` service on the related system. By default, the service runs under the Local System account, which is automatically added to the Data Protector `admin` user group. However, if you have restarted the Data Protector `Inet` service on the application and backup systems under `OSDBA` user accounts, you need to add new users to the Data Protector `admin` or `operator` user group.

For information on adding users to Data Protector user groups, see the *HP Data Protector Help* index: “adding users”.

NOTE: The OSDBA user account for the backup system needs to be added to a Data Protector user group only if you plan to use the Oracle backup set ZDB method.

Clusters

In cluster environments, ensure to add the following users to the Data Protector `admin` or `operator` user group:

- OSDBA user for all physical nodes
- OSDBA user for the virtual server (applicable for MC/ServiceGuard clusters)

Figure 7: Example user configuration in a cluster environment

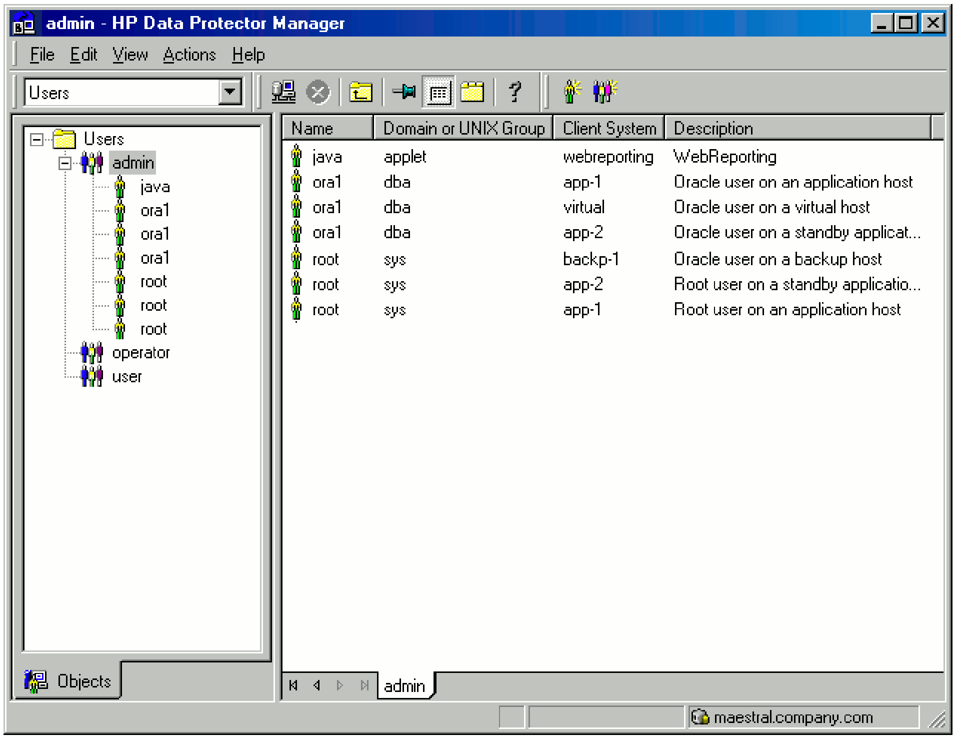

Name	Domain or UNIX Group	Client System	Description
java	applet	webreporting	WebReporting
ora1	dba	app-1	Oracle user on an application host
ora1	dba	virtual	Oracle user on a virtual host
ora1	dba	app-2	Oracle user on a standby applicat...
root	sys	backp-1	Oracle user on a backup host
root	sys	app-2	Root user on a standby applicatio...
root	sys	app-1	Root user on an application host

Configuring Oracle database users accounts

Identify or create the following Oracle database user accounts.

User	Description
Primary database user	Required to log in to the primary database.
Recovery catalog user	The owner of the recovery catalog (for example, `rman`). Required to log in to the catalog database. Needed if you use the recovery catalog. If you are using Oracle 11g R2 or later, ensure that the owner of the Oracle recovery catalog: • is granted the CREATE ANY DIRECTORY and the DROP ANY DIRECTORY system privileges, which are required to use the Data Pump Export (expdp) and the Data Pump Import (impdp) utilities. • has `SELECT` permissions on `sys.v$instance` view. Start SQL*Plus and type:

	`grant select on v_$instance to`
	`recovery_catatalog_user;`

Standby database user	Required to log in to the standby database. Applicable only in Oracle Data Guard environments. Needed to back up the standby database.

You need to provide these user accounts when you configure the Oracle database.

Configuring Oracle databases using Data Protector

Configuration of an Oracle database consists of providing Data Protector with the following data:

- Oracle Server home directory
- Login information to the target database
- Optionally, login information to the recovery catalog database
- Optionally, login information to the standby database
- Optionally, ASM-related information
- Backup method to be used and the related options

During the configuration, the `util_oracle8.pl` command, which is started on the application system, saves the specified parameters in the Data Protector Oracle database specific configuration file on the Cell Manager.

Ensure that the database is open during the configuration procedure and that you are able to connect to the database. To configure an Oracle database, you can use the Data Protector GUI or the Data Protector CLI.

Using the Data Protector GUI

Configure an Oracle database when you create the first ZDB backup specification for the database. Start with the procedure described in the Creating backup specifications section, and at Step 10, proceed as follows:

1. In the Configure Oracle dialog box, click the **General** tab, and specify the pathname of the Oracle Server home directory.

Figure 8: Configuring Oracle - General (UNIX)

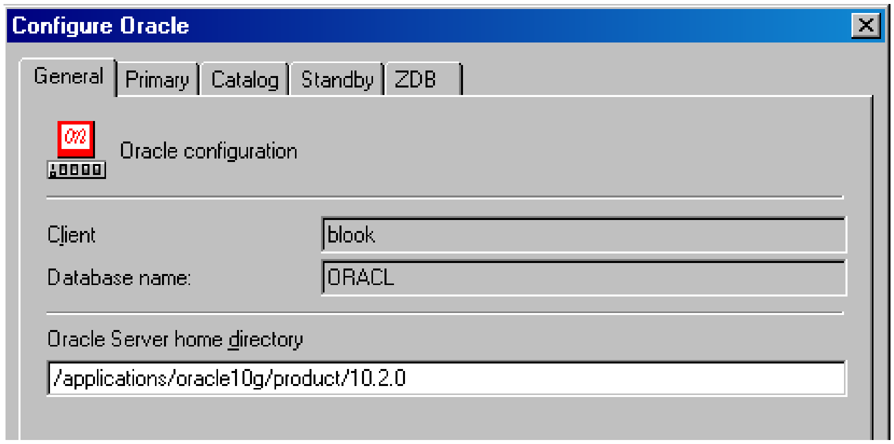

Client: blook

Database name: ORACL

Oracle Server home directory: /applications/oracle10g/product/10.2.0

2. In the **Primary** page, specify the login information to the primary database.

NOTE: The user must have the SYSDBA privilege granted.

In **Services**, type the net service name for the primary database instance. The backup will be performed on the system where this database instance resides.

RAC: List all net services names for the primary database separated by a comma.

Figure 9: Configuring Oracle - Primary

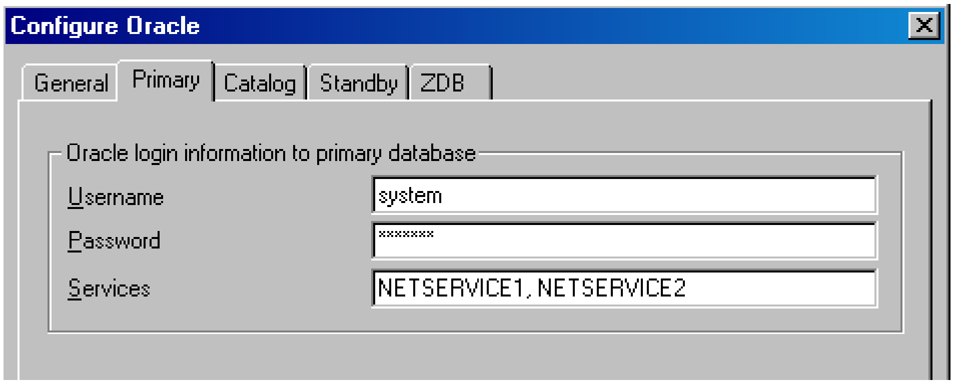

Oracle login information to primary database

Username: system

Password: *******

Services: NETSERVICE1, NETSERVICE2

3. In the **Catalog** page, select **Use target database control file instead of recovery catalog** to use the primary database control file.

   To use the recovery database catalog as an RMAN repository for backup history, select **Use recovery catalog** and specify the login information to the recovery catalog. For ZDB, you must use the recovery catalog, and the user specified must be the owner of the recovery catalog.

   In **Services**, type the net service name for the recovery catalog.

Figure 10: Configuring Oracle - Catalog

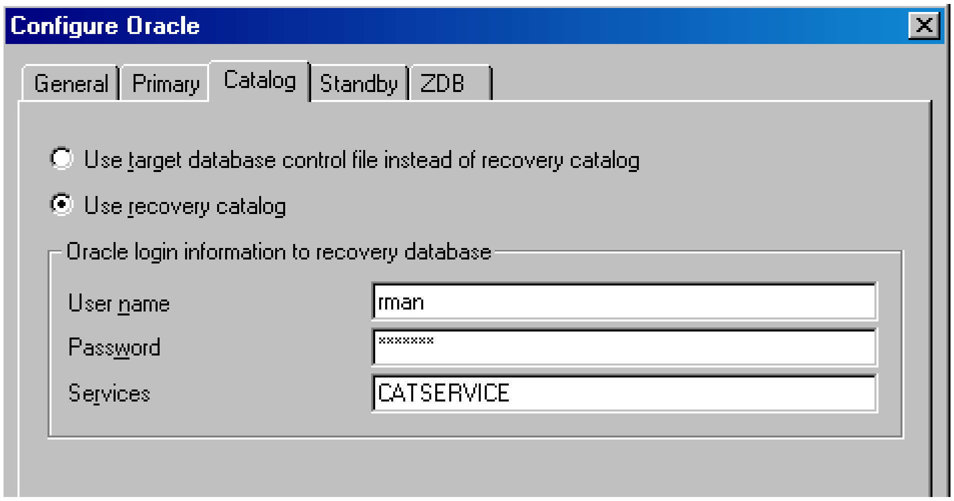

○ Use target database control file instead of recovery catalog

● Use recovery catalog

Oracle login information to recovery database

User name: rman

Password: *******

Services: CATSERVICE

4. If you have Oracle Data Guard configuration for non-ZDB sessions and if you intend to back up a standby database, configure the standby database too.

   In the **Standby** page, select **Configure standby database** and specify the login information to the standby database. In **Services**, type the net service name for the standby database instance.

RAC: List all net services names for the standby database separated by a comma.

Figure 11: Configuring Oracle - Standby

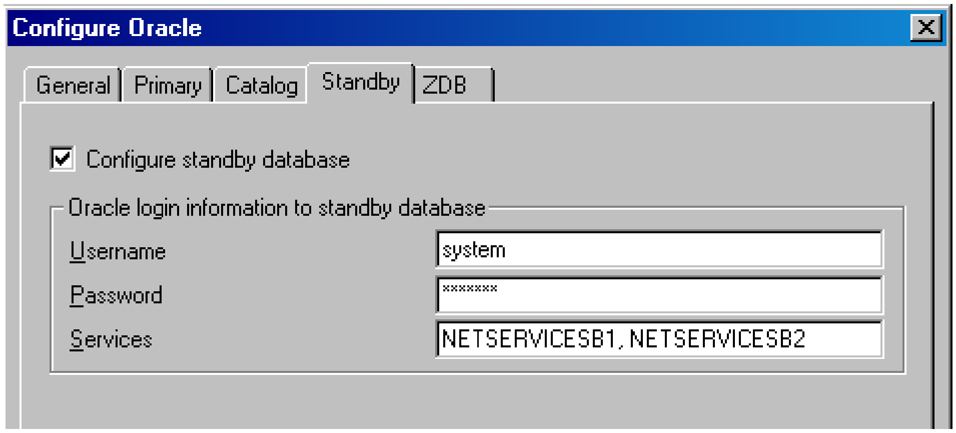

☑ Configure standby database

Oracle login information to standby database

Username: system

Password: *******

Services: NETSERVICESB1, NETSERVICESB2

5. In the **ZDB** page, select **Backup method** and then select **PROXY** or **BACKUP SET** in the drop-down list.

   In **Backup control file copy location**, you can specify the location on the source volumes where a backup copy of the current control file will be made during the ZDB-to-disk backup.

   If you do not specify the location, `ob2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

   If your backup method is *backup set* and if your database instance uses PFILE (and not SPFILE), select the **Parameter file (PFILE)** option and specify the pathname of `PFILE` residing on the application system.

Figure 12: Configuring Oracle - ZDB

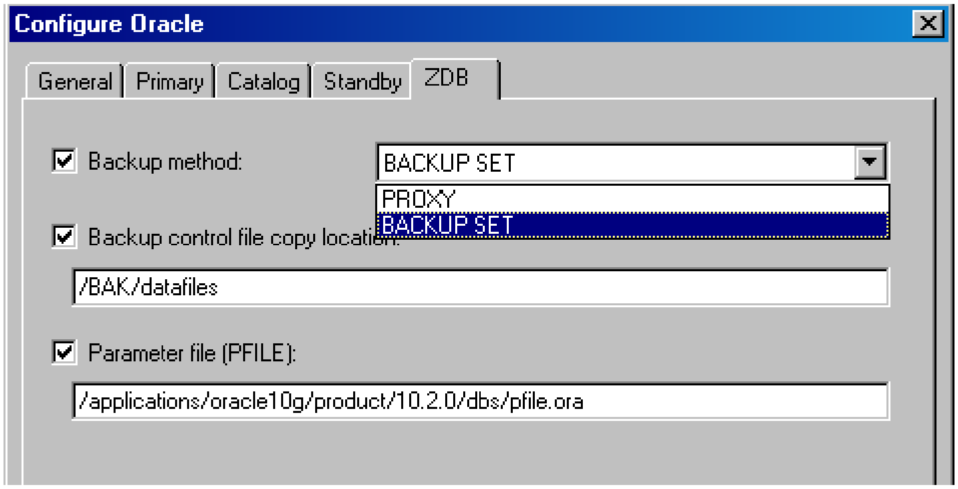

☑ Backup method: BACKUP SET

☑ Backup control file copy location: /BAK/datafiles

☑ Parameter file (PFILE): /applications/oracle10g/product/10.2.0/dbs/pfile.ora

6. Click **OK**.

   The Oracle database is configured. Exit the GUI or proceed with creating the backup specification.

Using the Data Protector CLI

1. On UNIX systems, log on to the Oracle Server system with an OSDBA user account.
2. On the Oracle Server system, execute:

   Windows systems:
   ```
   perl -I..\lib\perl util_oracle8.pl -config -dbname DB_NAME –orahome
   ORACLE_HOME PRIMARY_DB_LOGIN [CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN]
   [ZDB_OPTIONS] [ASM_OPTIONS] [-client CLIENT_NAME]
   ```

   UNIX systems:
   ```
   util_oracle8.pl -config -dbname DB_NAME -orahome ORACLE_HOME
   PRIMARY_DB_LOGIN [CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN]
   [ZDB_OPTIONS] [ASM_OPTIONS] [-client CLIENT_NAME]
   ```

   where:

   `PRIMARY_DB_LOGIN` is:
   -prmuser PRIMARY_USERNAME
   -prmpasswd PRIMARY_PASSWORD
   -prmservice PRIMARY_NET_SERVICE_NAME_1[,PRIMARY_NET_SERVICE_NAME_2 ...]

   `CATALOG_DB_LOGIN` is:
   -rcuser CATALOG_USERNAME
   -rcpasswd CATALOG_PASSWORD
   -rcservice CATALOG_NET_SERVICE_NAME

   `STANDBY_DB_LOGIN` is:
   -stbuser STANDBY_USERNAME
   -stbpasswd STANDBY_PASSWORD
   -stbservice STANDBY_NET_SERVICE_NAME_1[,STANDBY_NET_SERVICE_NAME_2 ...]

   `ZDB_OPTIONS` are:
   -zdb_method {PROXY | BACKUP_SET}
   [-ctlcp_location BACKUP_CONTROL_FILE_COPY_LOCATION]
   [-pfile PARAMETER_FILE]
   [-bkphost BACKUP_SYSTEM]

   `ASM_OPTIONS` are:
   [-asmhome ASM_HOME]
   [-asmuser ASM_USER -asmpasswd ASM_PASSWORD -asmservice
   ASM_NET_SERVICE_NAME_1[,ASM_NET_SERVICE_NAME_2 ...]]

If you have Oracle Data Guard configuration for *non-ZDB* sessions and if you intend to back up a standby database, you must provide the *`STANDBY_DB_LOGIN`* information.

To configure an Oracle database for ZDB, you must provide the *`ZDB_OPTIONS`* information. If your ZDB method is *backup set*, you must also provide the *`BACKUP_SYSTEM`* information.

The *`ASM_OPTIONS`* options are needed for instant recovery in Oracle Server configurations that use Automatic Storage Management (ASM).

Parameter description

CLIENT_NAME

Name of the Oracle Server system with the database to be configured. It must be specified in a cluster environment or if the ZDB configuration is run on the backup system.
RAC: The virtual server of the Oracle resource group.
Oracle Data Guard: Name of either a primary system or secondary (standby) system.

DB_NAME

Name of the database to be configured.

ORACLE_HOME

Pathname of the Oracle Server home directory.

PRIMARY_USERNAME PRIMARY_PASSWORD

Username and password for login to the target or primary database. Note that the user must have the SYSDBA privilege granted.

PRIMARY_NET_SERVICE_NAME_1 [,PRIMARY_NET_SERVICE_NAME_2, ...]

Net services names for the primary database.
RAC: Each net service name must resolve into a specific database instance.

CATALOG_USERNAME CATALOG_PASSWORD

Username and password for login to the recovery catalog. This is optional and is used only if you use the recovery catalog database as an RMAN repository for backup history.

CATALOG_NET_SERVICE_NAME

Net service name for the recovery catalog.

STANDBY_USERNAME STANDBY_PASSWORD

This is used in Oracle Data Guard environment for backing up a standby database. Username and password for login to the standby database.

STANDBY_NET_SERVICE_NAME_1 [,STANDBY_NET_SERVICE_NAME_2, ...]

Net services names for the standby database.

BACKUP_CONTROL_FILE_COPY_LOCATION

A location on a source volume where a copy of the current control file is made before a ZDB to disk. This is optional and if not specified, ob2rman.pl will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

PARAMETER_FILE

Full pathname of the PFILE residing on the application system. This is optional and used if backup method is backup set and the database instance uses PFILE (and not SPFILE).

BACKUP_SYSTEM

Name of the backup system. It must be specified for a ZDB backup set configuration.

ASM_HOME

Home directory of the ASM instance in an Oracle ASM configuration. Specify this option if the value differs from the home directory of the Oracle database instance.

ASM_USERNAME ASM_PASSWORD

User name and password (authentication credentials) used by the Data Protector Oracle integration agent to connect to the ASM database.

ASM_NET_SERVICE_NAME_1[,ASM_NET_SERVICE_NAME_2 ...]

Name of the net service to be used to access the ASM database. For Oracle environments involving multiple net services, multiple names can be specified.

The message `*RETVAL*0` indicates successful configuration, even if followed by additional messages.

NOTE: If you need to export some variables before starting SQL*Plus, these variables must be defined in the `Environment` section of the Data Protector Oracle global configuration file or using the Data Protector GUI.

Example

The following example represents configuration on a UNIX system of an Oracle database and its recovery catalog with the backup set method used and the parameter file location specified.

The following names are used in the example:

- database name: `oracle`
- Oracle Server home directory: `/app10g/oracle10g/product/10.1.0`
- primary user name: `system`
- primary password: `manager`
- primary net service name 1: `netservice1`
- primary net service name 2: `netservice2`
- recovery catalog user name: `rman`
- recovery catalog password: `manager`
- recovery catalog net service name: `catservice`
- backup system name: `bcksys`

Syntax

```
/opt/omni/lbin/util_oracle8.pl -config -dbname oracle -orahome
/app10g/oracle10g/product/10.1.0 -prmuser system -prmpasswd manager
-prmservice netservice1,netservice2 -rcuser rman -rcpasswd manager
-rcservice catservice -zdb_method BACKUP_SET -pfile
/app10g/oracle10g/product/10.1.0/dbs/pfile.ora -bkphost bcksys
```

Example

The following example shows how to configure an Oracle database in a cluster environment. The database files are managed by ASM. The configuration enables instant recovery in ASM environments:

- Database name: `SUN`
- Oracle Server home directory: `/orahome/ora/app/oracle/product/11.2.0/dbhome_1`
- Primary user name: `sys`
- Primary password: `oracle`
- Primary net service name 1: `SUN1`
- Primary net service name 2: `SUN2`
- Recovery catalog user name: `rman`
- Recovery catalog password: `manager`
- Recovery catalog net service name: `RECO`

- ZDB method: `Backup set`
- Oracle Server system (cluster virtual system): `cluster.company.com`
- Backup system: `backup.company.com`
- ASM home directory: `/oracle/crshome/crshome/crs/app/11.2.0/grid`
- ASM user: `sys`
- ASM user password: `oracle`
- ASM net service name 1: ASMSRV1
- ASM net service name 2: ASMSRV2

To configure the database, execute:

```
opt/omni/lbin/util_oracle8.pl -config -dbname SUN -orahome
/orahome/ora/app/oracle/product/11.2.0/dbhome_1 -prmuser sys -prmpasswd
oracle -prmservice SUN1,SUN2 -rcuser rman -rcpasswd manager -rcservice
RECO -zdb_method BACKUP_SET -bkphost backup.company.com -client
cluster.company.com -asmhome /crshome/crs/app/11.2.0/grid –asmuser sys
–asmpasswd oracle –asmservice ASMSRV1, ASMSRV2
```

Checking the database configuration

You can check the configuration of an Oracle database after you have created at least one backup specification for the database. If you use the Data Protector CLI, a backup specification is not needed.

Using the Data Protector GUI

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

IMPORTANT: Data Protector does not check if the specified user has appropriate Oracle backup permissions.

Using the Data Protector CLI

1. On UNIX systems, log on to the application system with an OSDBA user account.
2. Execute:

   Windows systems:
   ```
   perl -I..\lib\perl util_oracle8.pl -chkconf_smb -dbname DB_NAME
   ```

   UNIX systems:
   ```
   util_oracle8.pl -chkconf_smb -dbname DB_NAME
   ```

Handling errors

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

To get the error description, on the Cell Manager, execute:

- ***Windows systems:*** `Data_Protector_home\bin\omnigetmsg 12 error_number`
- ***UNIX systems:*** `/opt/omni/lbin/omnigetmsg 12 error_number`

IMPORTANT: On UNIX systems, it is possible that although you receive `*RETVAL*0`, backup still fails because Data Protector does not check if the specified user has appropriate Oracle backup permissions.

Checking Oracle configuration for instant recovery

Check if the Oracle configuration is suitable for instant recovery. On the application system, execute:

Windows systems:

```
perl util_oracle8.pl -chkconf_ir -dbname DB_NAME
```

UNIX systems:

```
util_oracle8.pl -chkconf_ir -dbname DB_NAME
```

If the control files, SPFILE, and online redo logs are on the same volume group (if LVM is used) or source volume as datafiles, a warning is displayed stating that instant recovery is not possible.

You can do one of the following:

- Reconfigure the Oracle database instance.
- Set the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR omnirc` options and ignore the warning. However, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery.

For more information on moving the control files and redo logs to source volumes that are not replicated and on setting the `omnirc` options, see the *HP Data Protector Integration Guide*.

Setting environment variables

Use environment variables to modify backup environment to suit your needs. Environment variables are specific to an Oracle database. It means that they can be set differently for different Oracle databases.

Once specified, they are saved to related Data Protector Oracle database configuration files.
For details on how environment variables affect your environment, see the following table.

Environment variable	Default value	Description
`OB2_RMAN_COMMAND_TIMEOUT`	`300 s`	This variable is applicable when Data Protector tries to connect to a target or catalog database. It specifies how long (in seconds) Data Protector waits for RMAN to respond that the connection succeeded. If RMAN does not respond within the specified time, Data Protector aborts the current session.
`OB2_SQLP_SCRIPT_TIMEOUT`	`300 s`	This variable is applicable when Data Protector issues an SQL*Plus query. It specifies how long Data Protector waits for SQL*Plus to respond that the query completed successfully. If SQL*Plus does not respond within the specified time, Data Protector aborts the current session.
`OB2_DPMCTL_SHRLOC`	`N/A`	Defines the location at which the control file is created and from where it is backed up in Data Protector managed control file backup. Data Protector copies the control file to the directory `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp`

		(Windows systems) by default. This variable overrides the default directory with a customer-specified directory. In an Oracle Real Application Clusters (RAC) environments with Oracle version 11.2.0.2 or later, to enable Data Protector managed control file backups and the corresponding restore sessions, ensure this directory resides on a shared disk that all RAC nodes can access.

To set environment variables, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can set a variable when you create a backup specification or modify an existing one:

1. In the Source page of the backup specification, right-click the Oracle database at the top and click **Set Environment Variables**.
2. In the Advanced dialog box, specify the variable name, its value, and click **Add**.

Figure 13: Setting environment variables

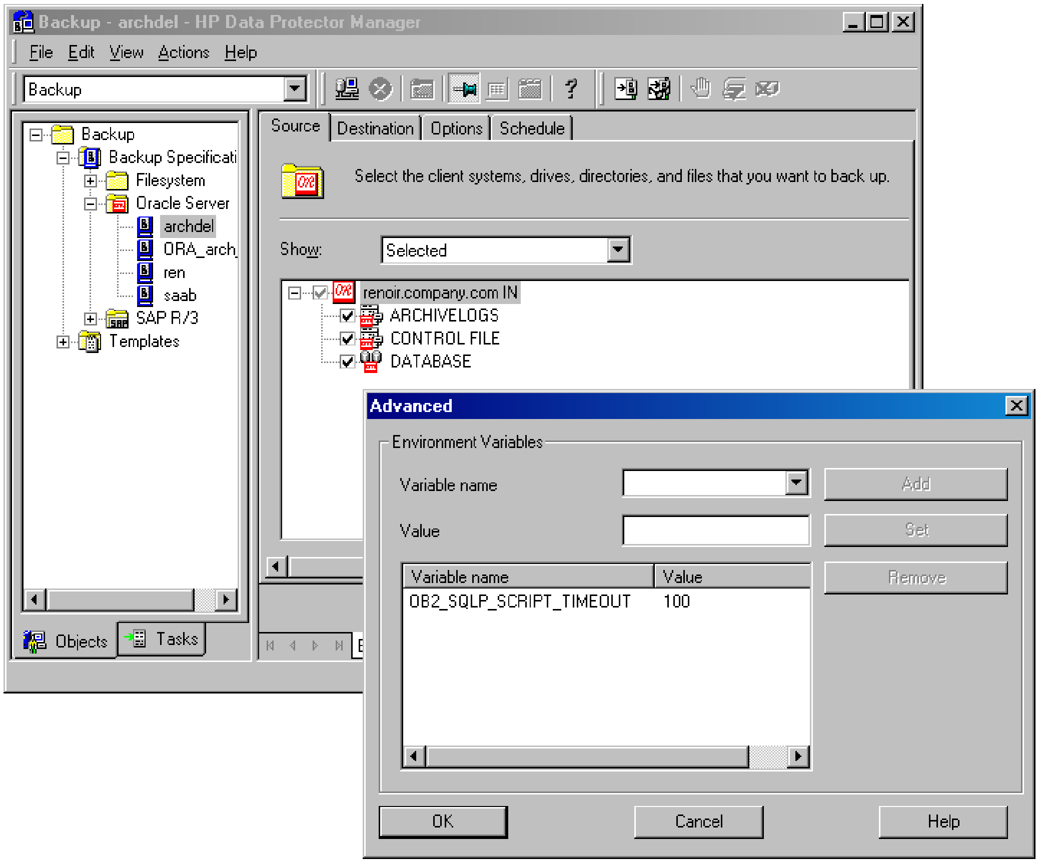

Select the client systems, drives, directories, and files that you want to back up.

Show: Selected

Variable name	Value
OB2_SQLP_SCRIPT_TIMEOUT	100

3. Click **OK**.

Using the Data Protector CLI

Execute:

```
util_cmd –putopt Oracle8 DatabaseName Variable Value –sublist Environment
```

For details, see the util_cmd man page or the *HP Data Protector Command Line Interface Reference*.

Example

To set the environment variable OB2_RMAN_COMMAND_TIMEOUT to 100 seconds for the Oracle database INST2, execute:

```
util_cmd –putopt Oracle8 INST2 OB2_RMAN_COMMAND_TIMEOUT 100 –sublist
Environment
```

Backup

To configure an Oracle ZDB, perform the following steps:

1. Configure the devices you plan to use for a backup. For instructions, see the *HP Data Protector Help* index: "configuring devices".
2. Configure media pools and media for a backup. For instructions, see the *HP Data Protector Help* index: "creating media pools".
3. Ensure that you are able to connect to the database.
4. Configure a non-ZDB backup specification and run the backup of Oracle data on the application system to verify that you have properly configured the Oracle environment. For information on how to create a non-ZDB backup specification, see the *HP Data Protector Integration Guide*.
5. Create a Data Protector Oracle ZDB backup specification as explained in the following section.

Creating backup specifications

Online ZDB

To perform an online ZDB of an Oracle database, the database has to run in the ARCHIVELOG mode.

Offline ZDB

To perform an offline ZDB, create a ZDB backup specification only.

Cluster-aware systems

Before you perform an offline ZDB in a cluster environment, take the Oracle Database resource offline and bring it back online after the replica is created. This can be done using the Oracle fscmd command line interface commands in the Pre-exec and Post-exec commands for the client system in a particular backup specification, or by using the Cluster Administrator.

You cannot perform a ZDB of the archived redo log files. Therefore, you need to create two backup specifications:

- ZDB backup specification for backing up database files
- Standard Data Protector Oracle integration backup specification for backing up the application system archived log files

Procedure

To create an Oracle ZDB backup specification:

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

3. In the Create New Backup dialog box, select the following:
    a. From the Oracle Server list, select a template based on the required ZDB method:
        - For backup set: `SMB_BackupSet_Database`
        - For proxy copy: `SMB_Proxy_Database`
    b. From the Backup type drop-down list, select **Snapshot or split mirror backup**.
    c. From the Sub type drop-down list, select **HP 3PAR**.

Figure 14: Create New Backup dialog box

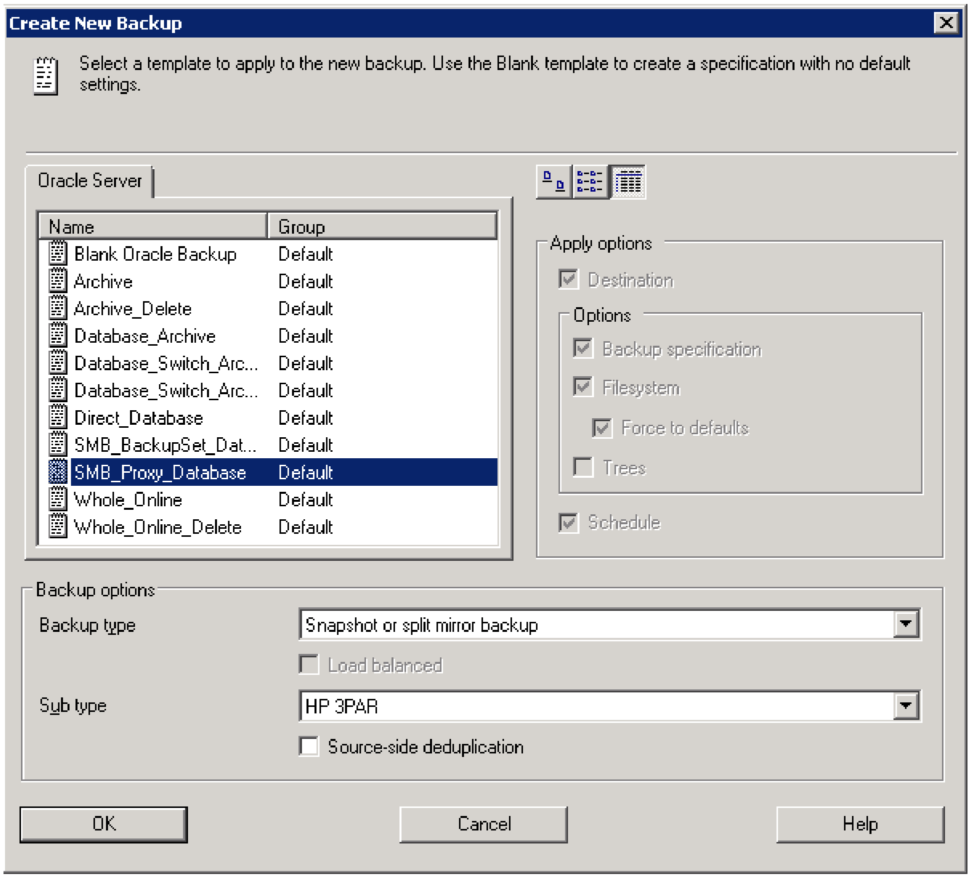

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Oracle Server

Name	Group
Blank Oracle Backup	Default
Archive	Default
Archive_Delete	Default
Database_Archive	Default
Database_Switch_Arc...	Default
Database_Switch_Arc...	Default
Direct_Database	Default
SMB_BackupSet_Dat...	Default
SMB_Proxy_Database	Default
Whole_Online	Default
Whole_Online_Delete	Default

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

□ Trees

☑ Schedule

Backup options

Backup type: Snapshot or split mirror backup

□ Load balanced

Sub type: HP 3PAR

□ Source-side deduplication

4. Click **OK**
5. In Application system, select the Data Protector Oracle integration client. In a non-RAC cluster environment, select the virtual server.

    RAC: Select the virtual server of the Oracle resource group.
    In Backup system, select the backup system.

    Select other HP 3PAR-specific backup options. To enable instant recovery, select the **Track the replica for instant recovery** option.

Figure 15: HP 3PAR backup options

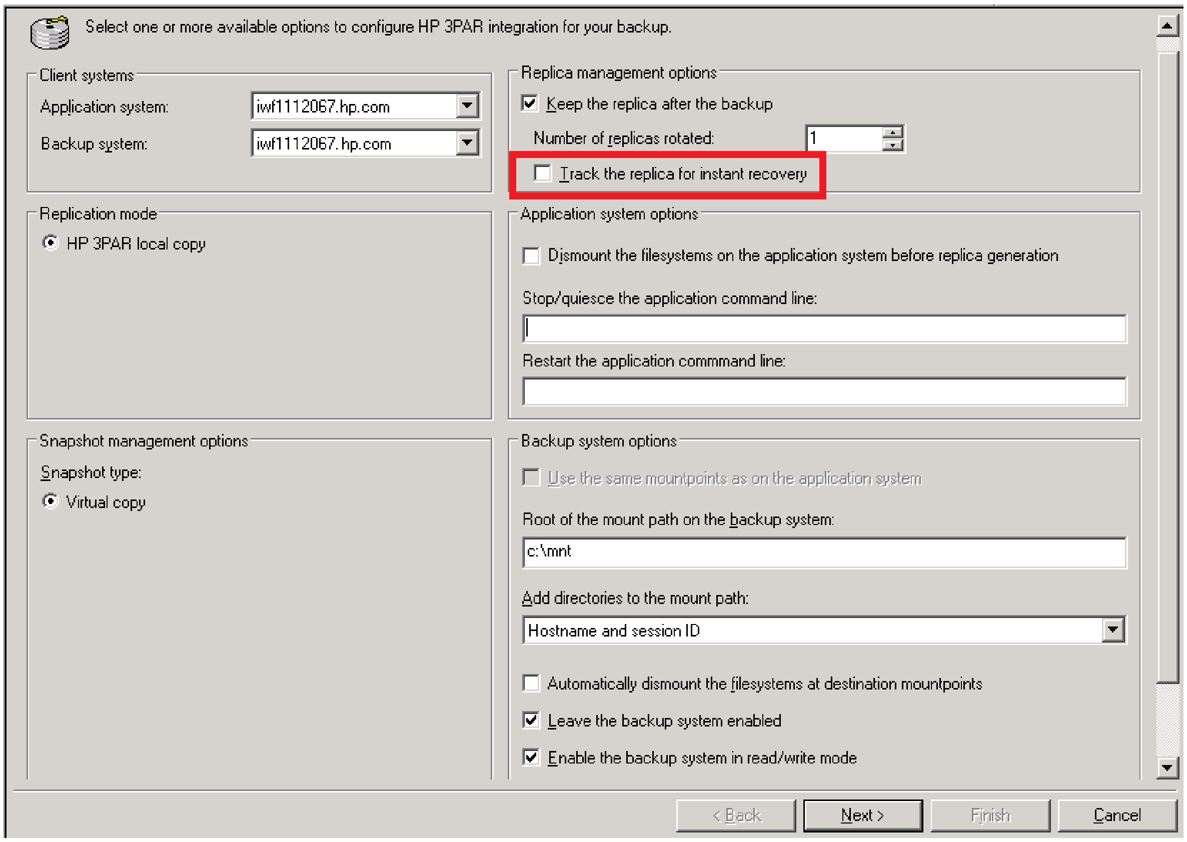

Select one or more available options to configure HP 3PAR integration for your backup.

Client systems

Application system: iwf1112067.hp.com

Backup system: iwf1112067.hp.com

Replica management options

☑ Keep the replica after the backup

Number of replicas rotated: 1

□ Track the replica for instant recovery

Replication mode

● HP 3PAR local copy

Application system options

□ Dismount the filesystems on the application system before replica generation

Stop/quiesce the application command line:

Restart the application commmand line:

Snapshot management options

Snapshot type:

● Virtual copy

Backup system options

□ Use the same mountpoints as on the application system

Root of the mount path on the backup system: c:\mnt

Add directories to the mount path: Hostname and session ID

□ Automatically dismount the filesystems at destination mountpoints

☑ Leave the backup system enabled

☑ Enable the backup system in read/write mode

6. Click **Next**.
7. In Application database, type the name of the database to be backed up. You can obtain the database name using the following SQL*Plus command:

   ```
   SQL>select name from v$database;
   ```

8. In Username and Group/Domain name, specify the OSDBA user account under which you want the backup to start (for example, the user name `ora`, group DBA). You can configure this user as described in the Configuring Oracle user accounts section.

   NOTE: For Windows Server 2008 systems, it is not mandatory to specify these options. If they are not specified, the backup runs under the Local System account.

   Ensure that this user is added to the Data Protector `admin` or `operator` user group and has the Oracle database backup rights. This user becomes the backup owner.

Figure 16: Specifying an Oracle Server system

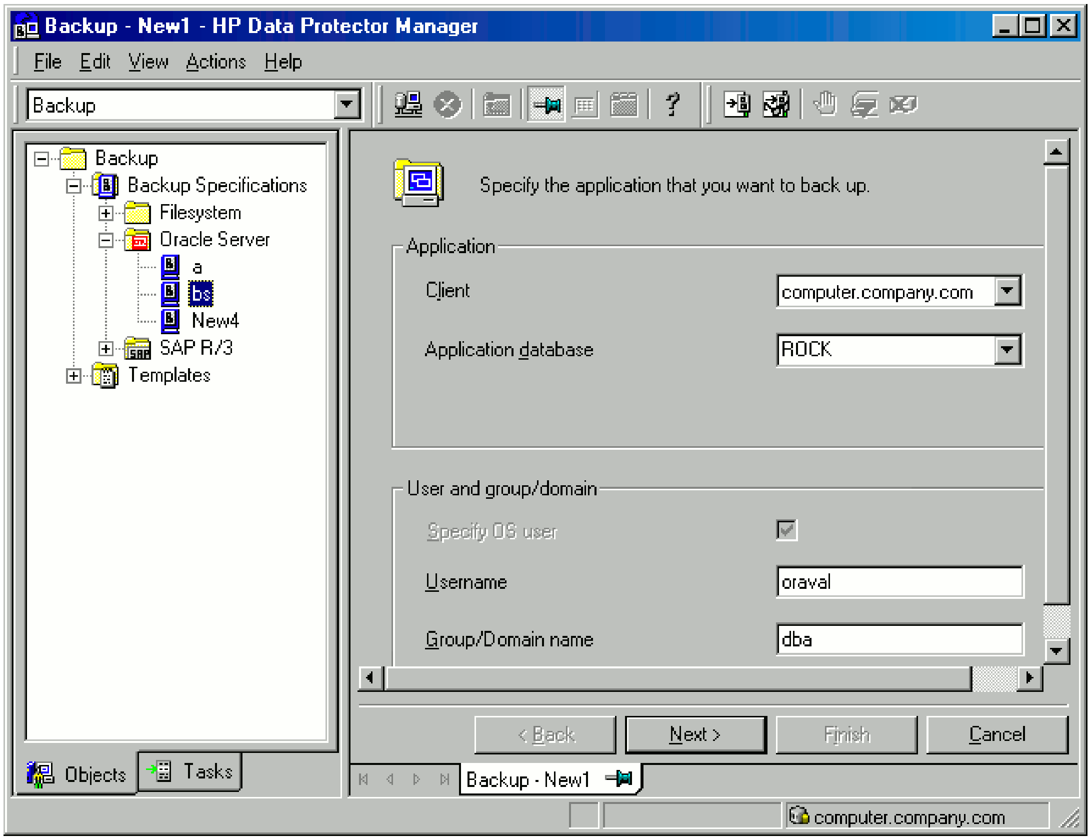

Specify the application that you want to back up.

Application

Client: computer.company.com

Application database: ROCK

User and group/domain

Specify OS user ☑

Username: oraval

Group/Domain name: dba

9. Click **Next**. Data Protector performs a configuration check. The check is started under the specified OSDBA user account. If it completes successfully, the OSDBA user and group are saved in both the Oracle database-specific configuration file and the Oracle system global configuration file, overriding previous values, if any.
10. If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box appears. Configure the Oracle database for use with Data Protector as described in the Configuring Oracle databases section.
11. Select the Oracle database objects to be backed up.

    NOTE: Since temporary tablespaces do not contain permanent database objects, RMAN and Data Protector do not back them up. For more information, see the Oracle documentation.

Figure 17: Selecting backup objects

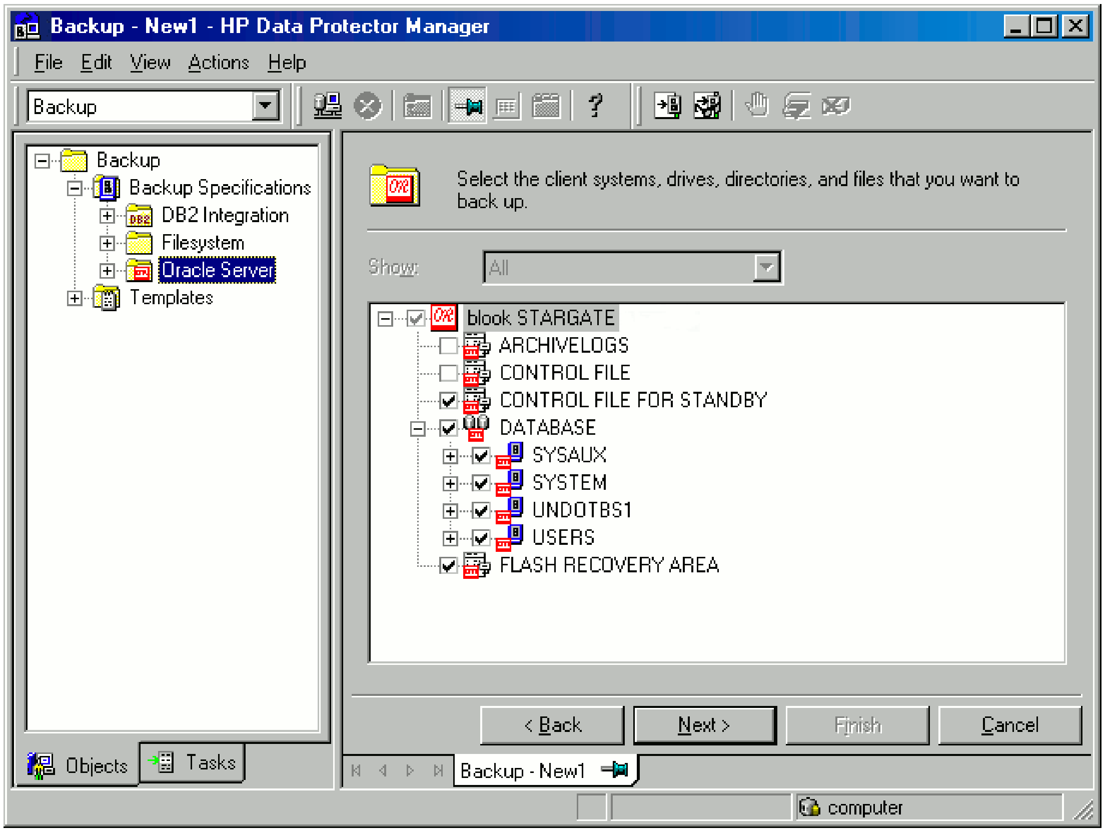

Select the client systems, drives, directories, and files that you want to back up.

- ☑ blook STARGATE
  - □ ARCHIVELOGS
  - □ CONTROL FILE
  - ☑ CONTROL FILE FOR STANDBY
  - ☑ DATABASE
    - ☑ SYSAUX
    - ☑ SYSTEM
    - ☑ UNDOTBS1
    - ☑ USERS
  - ☑ FLASH RECOVERY AREA

12. Click **Next**.

    If the backup method configured for this instance does not correspond to the method in the backup specification, Data Protector will display a warning and abort the configuration.

13. Select the device(s) you want to use for the backup. Click **Properties** to set the device concurrency, media pool, and pre-allocation policy. For more information on these options, click **Help**.

    You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror. For detailed information on the object mirror functionality, see the *HP Data Protector Help* index: “object mirroring”.

14. Click **Next** and then set the backup options.
    The following table provides details on the Oracle backup options.

Table 3: Backup options

Option	Description
Disable recovery catalog auto backup	By default, Data Protector backs up the recovery catalog after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the recovery catalog.
Disable Data Protector managed control file backup	By default, Data Protector backs up the Data Protector managed control file after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the Data Protector managed control file.

Back up standby database	This option is ignored for ZDB.
RMAN Script	You can edit the Oracle RMAN script section of the Data Protector Oracle backup specification. The script is created by Data Protector during the creation RMAN Script of a backup specification and reflects the backup specification's selections and settings. You can edit the script only after the backup specification has been saved. For information on editing the RMAN script, see the *HP Data Protector Integration Guide*.
Pre-exec, Post-exec	Specify a command or RMAN script that will be started by `ob2rman.pl` on the Oracle Server system before the backup (pre-exec) or after it (post-exec). RMAN scripts must have the `.rman` extension. Do not use double quotes. For example, you can provide scripts to shut down and start an Oracle instance. For examples of shutting down and starting an Oracle instance on a UNIX system, see Appendix A – Examples of pre-exec and post-exec scripts. Provide the pathname of the command or the RMAN script.
Backup offline	Select this option to perform an offline ZDB session. This option stops the database before creating a replica, and restarts it after the replica is created. For more information on this option, see the *HP Data Protector Integration Guide*.

For information on other the Backup Specification Options and Common Application Options, press **F1**.

15. After setting the backup options, click **OK**.
16. Click **Next** and then save the backup specification. It is recommended that you save all Oracle backup specifications in the Oracle group.

    IMPORTANT: The word `DEFAULT` is a reserved word and, therefore, must not be used for backup specification names or labels of any kind. Do not use any punctuation in the names of backup specifications, as the Oracle channel format is created from the backup specification name.

Figure 18: Saving the backup specification

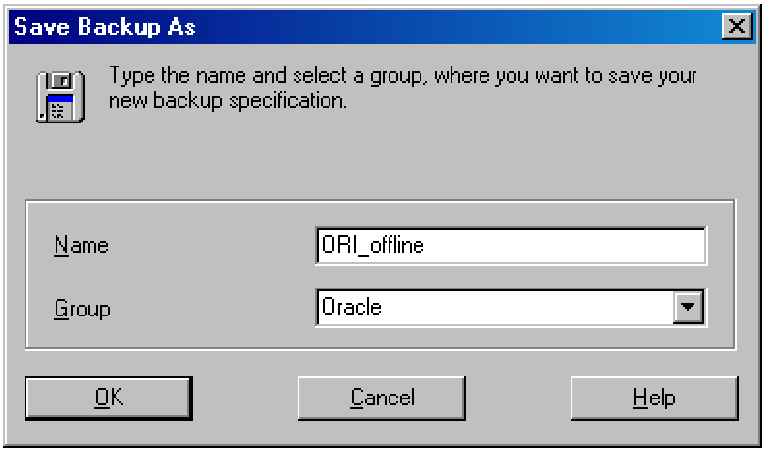

Save Backup As

Type the name and select a group, where you want to save your new backup specification.

Name: ORI_offline

Group: Oracle

17. Click **OK**.

To start the backup, see the following section

Starting backup sessions

To run a ZDB-to-disk, ZDB-to-tape, or ZDB-to-disk+tape backup session of an Oracle database, use any of the following methods:

- Schedule a backup of an existing Oracle ZDB backup specification using the Data Protector Scheduler. See the Scheduling backup sessions section.
- Start an interactive backup of an existing Oracle ZDB backup specification using the Data Protector GUI or the Data Protector CLI. See the Running an interactive backup section.

Considerations

Before running an Oracle ZDB session, consider the following:

- It is not possible to start a ZDB, restore, or instant recovery sessions using the same source volume on the application system at the same time. A ZDB, restore, or instant recovery session must be started only after the preceding session that is using the same source volume on the application system has finished the ZDB session or restore; otherwise, the session will fail.
- For the backup set method, if the Oracle database is installed on symbolic links, then these symbolic links have to be also created on the backup system.

Scheduling backup sessions

Scheduling a backup session means setting the time, date, and type of a backup that starts unattended once the scheduling options are defined and saved in the backup specification.

For more information on scheduling, see the HP Data Protector Help index: “scheduled backups”.

To schedule an Oracle ZDB backup specification, proceed as follows:

1. In the Data Protector Manager, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup Specifications** and then **Oracle Server**.
3. Double-click the backup specification you want to schedule and click the **Schedule** tab.
4. In the Schedule page, select a date in the calendar and click **Add** to open the Schedule Backup dialog box.
5. Specify **Recurring**, **Time options**, **Recurring options**, and **Session options**.

   NOTE: Only the Full backup type is supported.

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. You can run a ZDB-to-disk or a ZDB-to-disk+tape session only if the **Track the replica for instant recovery** option is selected in the backup specification.

Figure 19: Selecting ZDB-to-disk or ZDB-to-disk+tape session

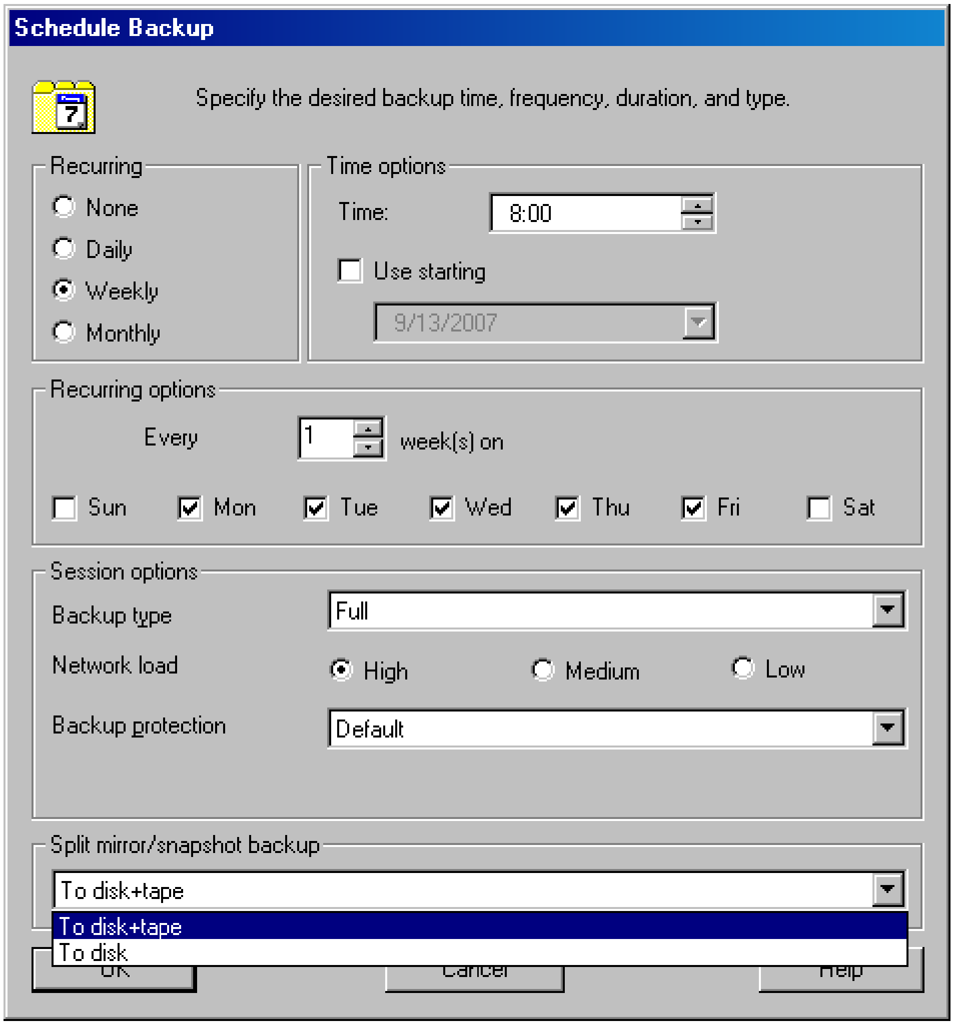

Schedule Backup

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

Recurring
○ None
○ Daily
● Weekly
○ Monthly

Time options
Time: 8:00
☐ Use starting
9/13/2007

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

Split mirror/snapshot backup
To disk+tape
- To disk+tape
- To disk

6. Click **OK** and then **Apply** to save the changes.

Running an interactive backup

You can start an interactive backup any time after a backup specification is created and saved. You can use the Data Protector GUI or CLI.

Starting a backup using the GUI

To start an interactive ZDB session of an Oracle database using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Oracle Server**. Right-click the backup specification you want to use and click **Start Backup**.
3. In the Start Backup dialog box, select the **Network load** option. For information on network load, click **Help**.

   NOTE: Only the Full backup type is supported.

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. You can run a ZDB-to-disk or a ZDB-to-disk+tape session only if the **Track the replica for instant recovery** option is selected in the backup specification.

Figure 20: Selecting ZDB to disk or ZDB to disk+tape session - interactive backup

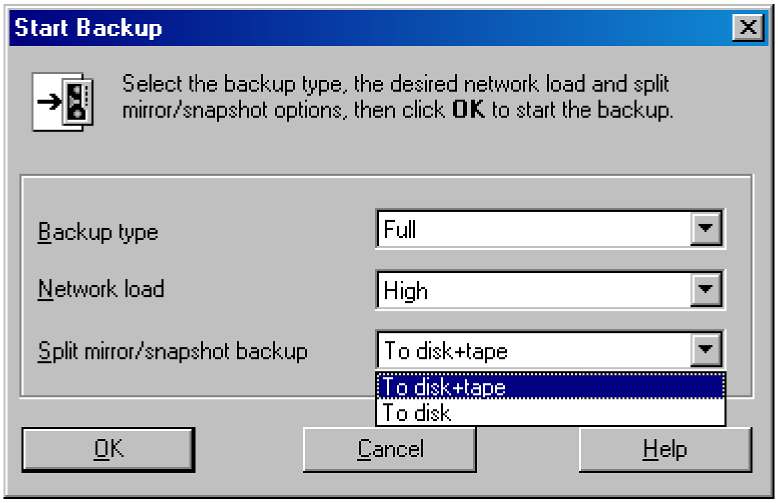

Start Backup

Select the backup type, the desired network load and split mirror/snapshot options, then click **OK** to start the backup.

Backup type: Full

Network load: High

Split mirror/snapshot backup: To disk+tape

4. Click **OK**.

Starting a backup using the CLI

To start an Oracle ZDB-to-tape or ZDB-to-disk+tape session using the Data Protector CLI, execute:

```
omnib -oracle8_list Name
```

To start an Oracle ZDB-to-disk session using the Data Protector CLI, execute:

```
omnib -oracle8_list Name -disk_only
```

where, `Name` is the name of the backup specification. For more information on the `omnib` command, see its man page or the *HP Data Protector Command Line Interface Reference*.

NOTE: It is not possible to run a ZDB-to-disk or a ZDB-to-disk+tape session if the **Track the replica for instant recovery backup** option is not selected in the backup specification.

Switching between Oracle backup methods

You can switch between the Oracle backup methods by reconfiguring the Data Protector Oracle integration for each database. It is not possible to select the method during the backup specification creation.

IMPORTANT: When switching between the Oracle backup set and proxy-copy methods, you must carefully follow the instructions given bellow to ensure a successful switch between both methods and to ensure that during a restore or recovery RMAN does not select backup objects backed up using different methods in one restore session. If such a mixed set is used, the restore procedure will fail.

To switch between the backup methods:

1. Successfully back up the entire database using the currently selected method.
2. To avoid selecting backup specifications with a backup method different than the current backup method, remove or move all ZDB backup specifications belonging to the selected database instance. The backup specifications are located on the Cell Manager in:

   Windows systems: Data_Protector_home\Config\Server\BarLists\Oracle8
   UNIX systems: /etc/opt/omni/server/barlists/oracle8

3. Re-configure the database with the new method selected while creating a new Oracle ZDB specification.

4. Optionally, if you switch from backup set to proxy-copy, you may:

    a. On the Cell Manager, remove the file:

       Windows systems:
       Data_Protector_home\Config\Server\Integ\Config\Oracle8\client_name%initDB_NAME_bckp.ora

       UNIX systems:
       /etc/opt/omni/server/integ/config/Oracle8/client_name%initDB_NAME_bckp.ora

    b. Remove the Oracle software from the backup system.

5. Perform ZDB of the entire database.

IMPORTANT: If you need to perform a restore from a time between the start and the end of the first backup of the entire database using the new backup method, RMAN may try to use backup files from the old method through a channel allocated for the files from the old method and the restore will fail.

Restore

You can restore the following database objects using the Data Protector GUI or RMAN:

- Control files
- Datafiles
- Tablespaces
- Databases
- Recovery Catalog Databases

Using the Data Protector GUI, you can also duplicate a production database (see the Duplicating an Oracle database section).

The following are the available methods in Data Protector for restoring database objects:

- Standard restore from backup media to the application system on LAN (see the Restoring from backup media to the application system on LAN section).
- Instant recovery (see the Instant recovery and database recovery section).

Prerequisites

- An instance of Oracle must be created on the system to which you want to restore or duplicate the database.
- The database must be in the `Mount` state if the whole database is being restored, or in the `NoMount` state if the control file is being restored or a database duplication is performed.
- You must be able to connect to the database.
- On Windows systems, when performing a restore from backup using the Oracle backup set ZDB method, set the `omnirc` option `ZDB_SMISA_AUTOMOUNTING` on the application system to 2, in order to enable automatic volume mounting on the local system.

Restoring from backup media to the application system on LAN

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

- Data Protector GUI (see the Restoring Oracle using the Data Protector GUI section)
- RMAN (see the Restoring Oracle using RMAN section)

Restoring Oracle using the Data Protector GUI

For restore, RMAN scripts are generated with necessary commands, depending on selections made in the GUI. To use additional commands, use them manually from RMAN.

Restoring database items in a disaster recovery

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

1. Restore the recovery catalog database (if it was lost).
2. Restore the control file.
3. Restore the entire database or data items.

Changing the database state

Before you restore any database item or you perform a duplication of a database, ensure that the database is in the correct state:

Table 4: Required database states

Item to restore	Database state
Control file, duplicating a database	`NoMount` (started)
All other items	`Mount`

To change the database to the correct state, execute:

```
sqlplus /nolog
SQL>connect user/password@service as sysdba;
SQL>shutdown immediate;
```

To change the database to the `NoMount` state, execute:

```
SQL>startup nomount;
```

To change the database to the `Mount` state, execute:

```
SQL>startup mount;
```

Restoring the recovery catalog database

The Oracle recovery catalog database is exported using the Oracle export utility to a binary file and backed up by Data Protector. This file has to be restored back to the disk and then imported into the Oracle database using the Oracle import utility. Data Protector provides a facility to do this automatically using the Oracle integration.

To restore the recovery catalog database:

1. Ensure that the recovery catalog database is in the **Open** state.
2. Remove the recovery catalog from the database (if it exists), using the RMAN command `DROP CATALOG`.
3. In the Data Protector GUI, switch to the **Restore** context.

4. Under Restore Objects, expand **Oracle Server**, expand the system on which the database, for which you want to restore the recovery catalog, resides, and then click the database.
5. In the Restore action drop-down list, select **Perform RMAN Repository Restore**.
6. In the Results Area, select **RECOVERY CATALOG**.

   If you want to change the recovery catalog login information, right-click **RECOVERY CATALOG** and click **Properties**. In Recovery Catalog Settings, specify the login information for recovery catalog.

Figure 21: Recovery catalog settings dialog

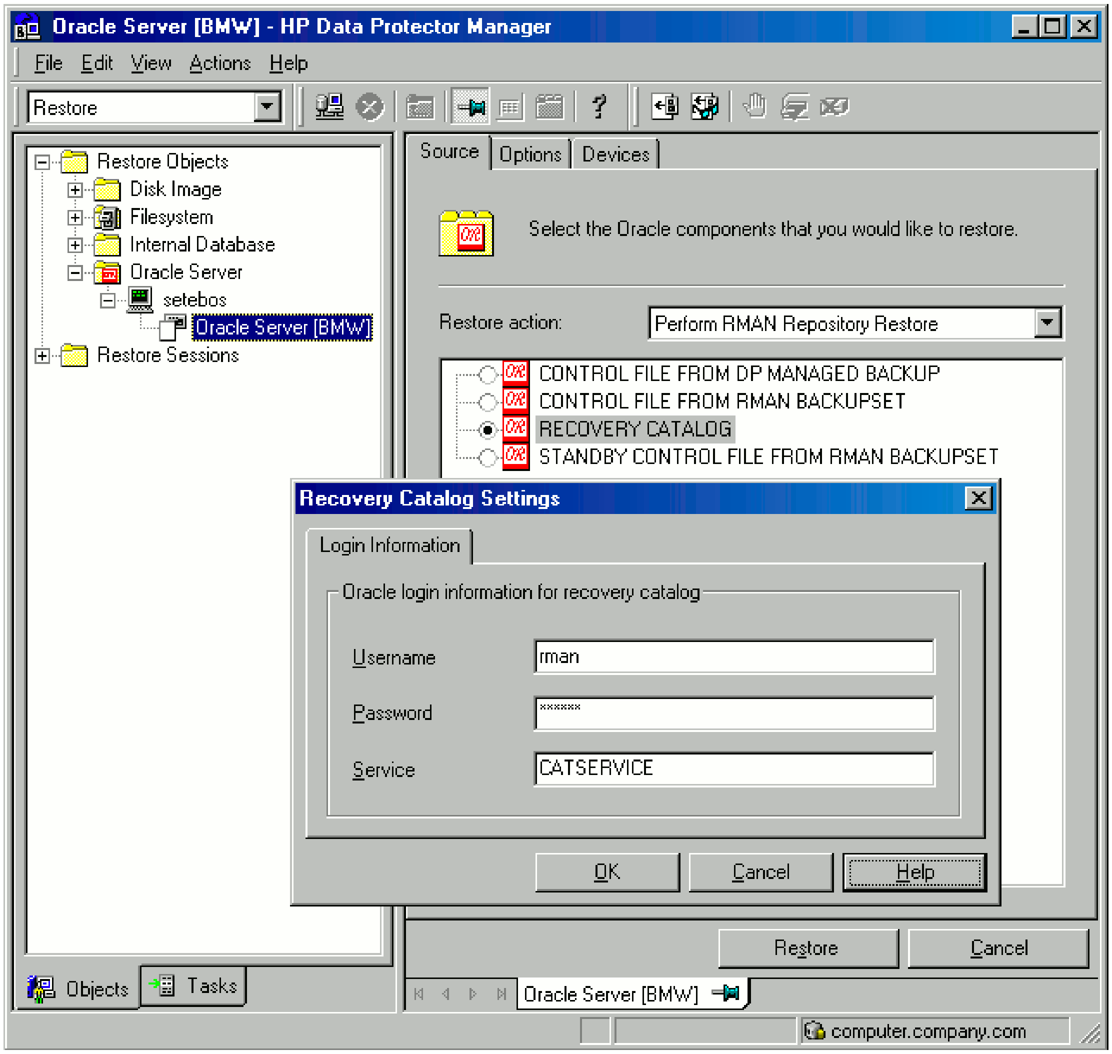

Oracle Server [BMW] - HP Data Protector Manager

Select the Oracle components that you would like to restore.

Restore action: Perform RMAN Repository Restore

○ CONTROL FILE FROM DP MANAGED BACKUP
○ CONTROL FILE FROM RMAN BACKUPSET
● RECOVERY CATALOG
○ STANDBY CONTROL FILE FROM RMAN BACKUPSET

Recovery Catalog Settings

Login Information

Oracle login information for recovery catalog

Username: rman
Password: ******
Service: CATSERVICE

7. In the Options page, perform the following:
   a. In User name and User group, specify the user name and password for the recovery catalog database.
   b. From the Session ID drop-down list, select the Session ID.

   For more information, see the Restore, recovery, and duplicate options section.

8. Click **Restore**.

Proceed to restore the control file.

Restoring the control file

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

Depending on the type of the control file backup, the following types of restore are possible when restoring the control file:

- Restoring from Data Protector managed control file backup (CONTROLFILE FROM DP MANAGED BACKUP)

  The control file was backed up automatically by ob2rman.pl at the end of a backup session, unless the option Disable Data Protector managed control file backup was selected.

  The recovery catalog is not required for this restore option. The control files (ctrlDB_NAME.dbf) are restored to:

  Windows systems: Data_Protector_home\tmp
  UNIX systems: /var/opt/omni/tmp

  NOTE: In Oracle Real Application Clusters (RAC) environments with Oracle versions 11.2.0.2 and later, the control files are created at, backed up from, and restored to the location defined by the OB2_DPMCTL_SHRLOC variable. This directory must reside on a shared disk and be accessible from all RAC nodes in order for restore sessions to succeed.

  After the restore, execute the following script:

  ```
  run {
  allocate channel 'dev0' type disk;
  restore controlfile from 'TMP_FILENAME';
  release channel 'dev0';
  }
  ```

  where, TMP_FILENAME is the location to which the file was restored.

- Restoring from RMAN backup set (CONTROLFILE FROM RMAN BACKUPSET)

  The recovery catalog is required.

A backup session can contain more than one type of the control file backup. To restore the control file:

1. Open the sqlplus window and change the database to the nomount state. See the Changing the database state section.
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under Restore Objects, expand **Oracle Server**, expand the system on which the database, for which you want to restore the control file, resides, and then click the database.
4. In the Restore Action drop-down list, select **Perform RMAN Repository Restore**. In the Results area, select the control file for restore.
5. In the Options page, from the Client drop-down list, select the system on which the Data Protector Oracle integration agent (ob2rman.pl) will be started. To restore the control file to a different database than it is selected, click **Settings** and specify the login information for the target database.

Set the other restore options. For information, see the Restore, recovery, and duplicate options section.

6. Click **Restore**.

Proceed to restore the Oracle database objects.

Restoring Oracle database objects

Before you restore Oracle database objects, ensure that you have an up-to-date version of the recovery catalog database and the control file. They contain the database structure information. If you do not have up-to-date versions of these files, restore them as described in the preceding sections.

To restore Oracle database objects:

1. Change the database to the mount state. See the Changing the database state section.
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under Restore Objects, expand **Oracle Server**, expand the system on which the database, for which you restore the database objects, resides, and then click the database.
4. In the Restore action drop-down list, select the type of restore you wish to perform. For information on the options, see the Restore, recovery, and duplicate options section.

   IMPORTANT: If you do not select Perform Restore and Recovery or Perform Recovery Only, you will have to recover the database objects manually using RMAN. For information, see the Restoring Oracle using RMAN section.

Figure 22: Source page

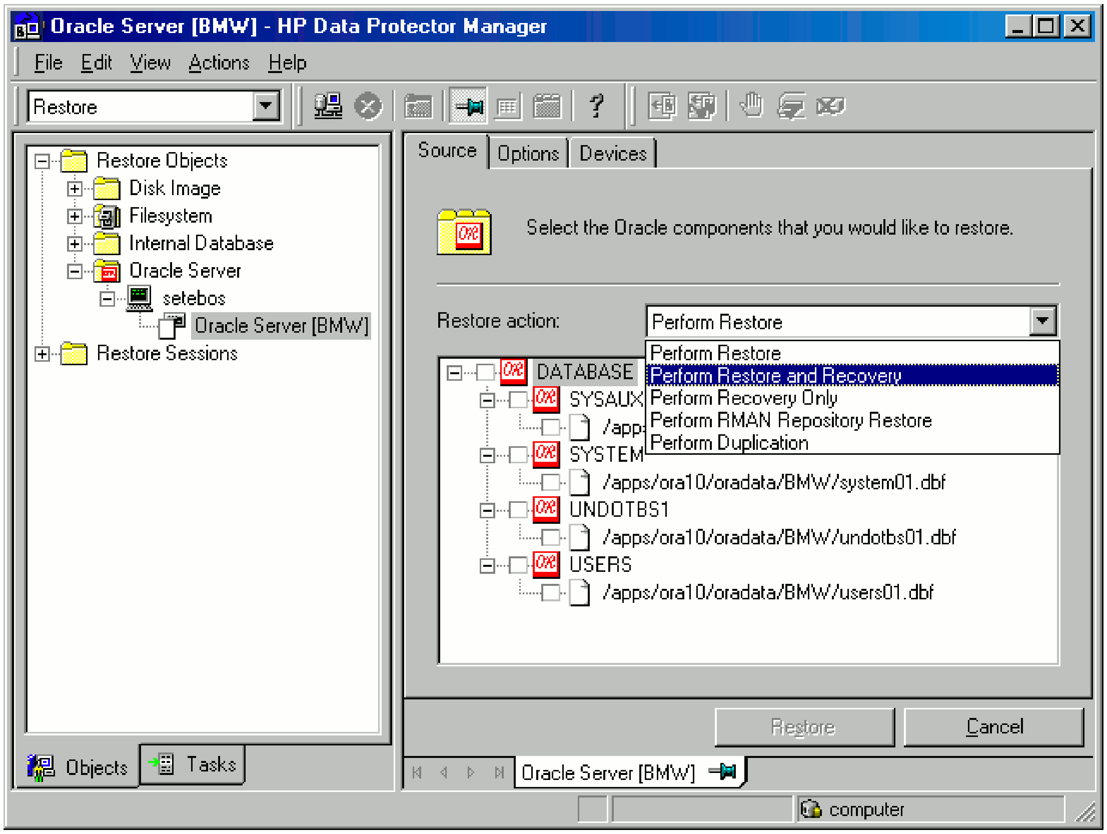

Select the Oracle components that you would like to restore.

Restore action: Perform Restore

5. In the Results Area, select objects for restore. If you are restoring datafiles, you can restore the files to a new location. Right-click the database object, click **Restore As**, and in the Restore As dialog box, specify the new datafile location.

   NOTE: When restoring to a new location, current datafiles will be switched to the restored datafile copies only if you have selected **Perform Restore and Recovery** from the **Restore action** drop-down list.

6. In the Options page, from the Client drop-down list, select the system on which the Data Protector Oracle integration agent will be started. To restore the database objects to a different database than it is selected, click **Settings** and specify the login information for the target database.

Figure 23: Options page

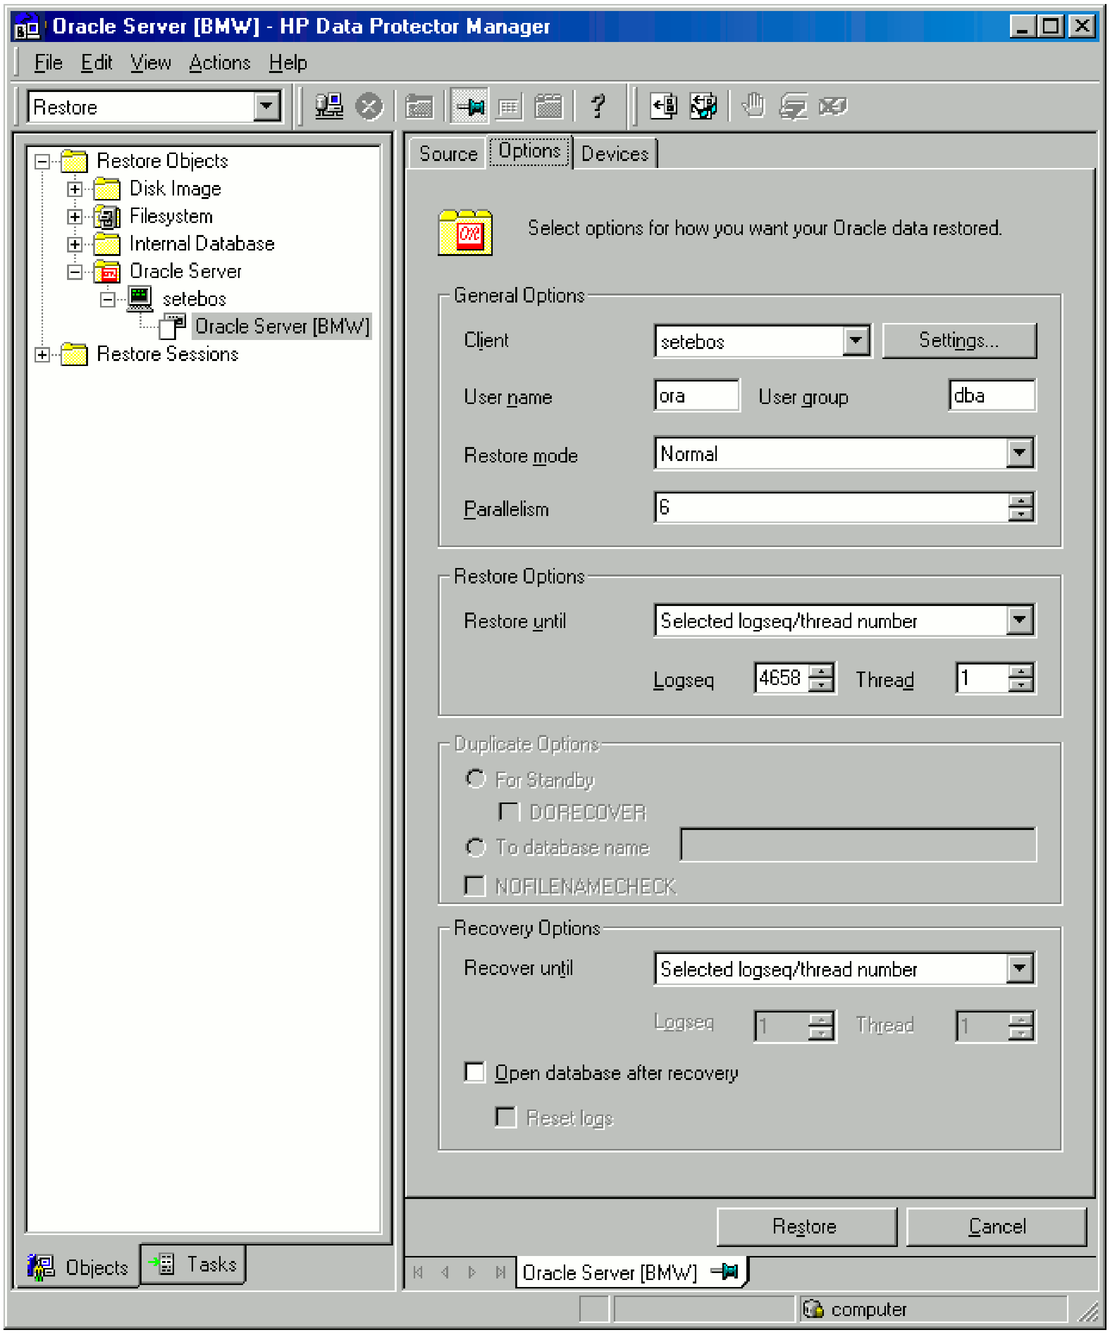

Oracle Server [BMW] - HP Data Protector Manager

Restore Objects
- Disk Image
- Filesystem
- Internal Database
- Oracle Server
  - setebos
    - Oracle Server [BMW]

Restore Sessions

Select options for how you want your Oracle data restored.

General Options

Client: setebos

User name: ora

User group: dba

Restore mode: Normal

Parallelism: 6

Restore Options

Restore until: Selected logseq/thread number

Logseq: 4658

Thread: 1

Duplicate Options

○ For Standby

□ DORECOVER

○ To database name

□ NOFILENAMECHECK

Recovery Options

Recover until: Selected logseq/thread number

Logseq: 1

Thread: 1

□ Open database after recovery

□ Reset logs

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

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

Figure 24: Devices page

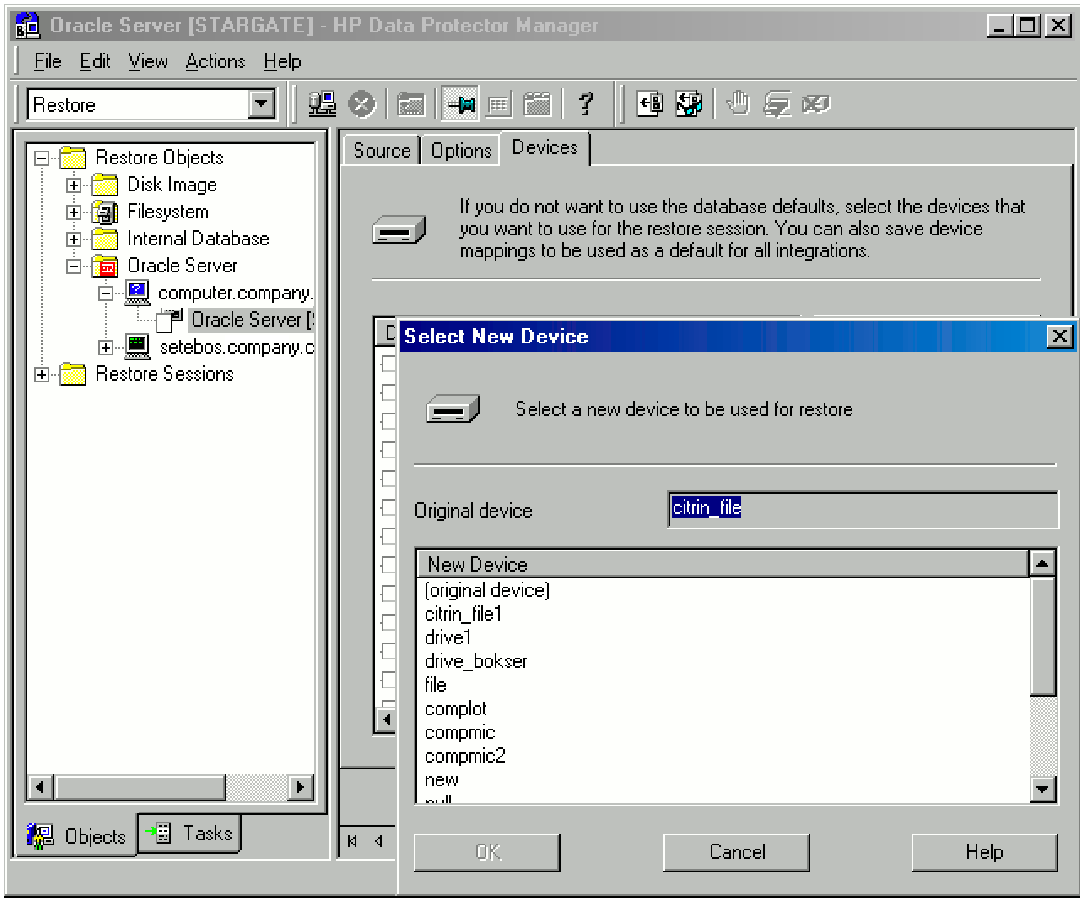

If you do not want to use the database defaults, select the devices that you want to use for the restore session. You can also save device mappings to be used as a default for all integrations.

Select a new device to be used for restore

Original device: citrin_file

New Device:
(original device)
citrin_file1
drive1
drive_bokser
file
complot
compmic
compmic2
new
null

7. Click **Restore**.

After the restore:

1. Change the database to the correct state. If you selected **Perform Restore and Recovery** or **Perform Recovery Only** in the Source page, then the database is automatically changed to the **Open** state by Data Protector.
2. If you performed an Oracle database restore and recovery until point in time, and the session has finished successfully, reset the database to register the new incarnation of database in the recovery catalog.

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

   ```
   rman target Target_Database_Login catalog Recovery_Catalog_Login
   RMAN> RESET DATABASE;
   RMAN> exit
   ```

3. If you did not choose to use Data Protector to recover the database objects and if you have all archived redo logs on disk, perform the following task after the database is restored:

Open a command line window and enter the following commands:

```
sqlplus /nolog
SQL>recover database;
SQL>connect user/password@service as sysdba;
SQL>alter database open;
```

Restoring tablespaces and datafiles

To restore tablespaces and datafiles:

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

   ```
   sqlplus /nolog
   SQL>connect user/password@service as sysdba;
   SQL>alter database datafile 'datafile name' offline;
   ```

   If you are restoring a tablespace enter:

   ```
   SQL>alter tablespace tablespace_name offline;
   ```

2. After completing the restore, get the datafiles and tablespaces back online with the following procedure:

   a. Open a command line window and enter the following commands:

      ```
      sqlplus /nolog
      SQL>connect user/password@service as sysdba
      ```

   b. If you are restoring a datafile, enter:

      ```
      SQL>alter database datafile 'datafile_name' online;
      ```

   c. If you are restoring a tablespace, enter:

      ```
      SQL>alter tablespace tablespace_name online;
      ```

Duplicating an Oracle database

Perform a production database duplication to create:

- A standby database which has the same DBID as the production (primary) database. With this, you can:

  - Create a new standby database.
  - Re-create a standby database after:

    - Loss of entire standby database
    - Primary database control file was restored or recreated
    - Database point-in-time recovery was performed on the primary database
    - Switchover or failover of database roles occurred

- An independent copy, with a unique DBID, which can be used for data mining or testing purposes.

Prerequisites

- The entire primary database with the archived logs must be backed up.
- Archive logs, which have not been backed up to tape since the last full backup and are required for duplication must be available on the duplicate system with the same path names as on the target system (system with the production database to be duplicated).
- Net service name for the auxiliary instance must be configured.
- When duplicating a database on the same system on which the target database resides, set all `*_PATH, *_DEST, DB_FILE_NAME_CONVERT, and LOG_FILE_NAME_CONVERT` initialization parameters appropriately. Thus, the target database files will not be overwritten by the duplicate database files.

To duplicate a production database:

1. On the system where the selected database will be duplicated, change the Oracle auxiliary database instance to the nomount state.
2. In the Context List of the Data Protector GUI, click **Restore**.
3. Under Restore Objects, expand **Oracle Server**, expand the system on which the production database resides, and then click the production database which you want to duplicate.

   If there are several such systems, select the system on which you want the Data Protector Oracle integration agent (`ob2rman.pl`) to be started.

4. In the Restore Action drop-down list, select **Perform Duplication**.
5. In the Options page, from the Client drop-down list, select the system on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started.
6. Click **Settings** to specify the login information (a user name, password, and net services name) for the auxiliary database.

   If you do not provide the login information, the duplication session will fail.

7. In User name and User group, specify the user name and group for the `OSDBA` account, which will be used by the Data Protector Oracle integration agent.
8. In Parallelism, specify the number of RMAN auxiliary channels to be allocated for database duplication. Set duplicate options.
9. If you are creating a new database copy (not for standby), specify the **Recover until** option to recover the duplicated database until a specified point in time.

Figure 25: Oracle duplicate options

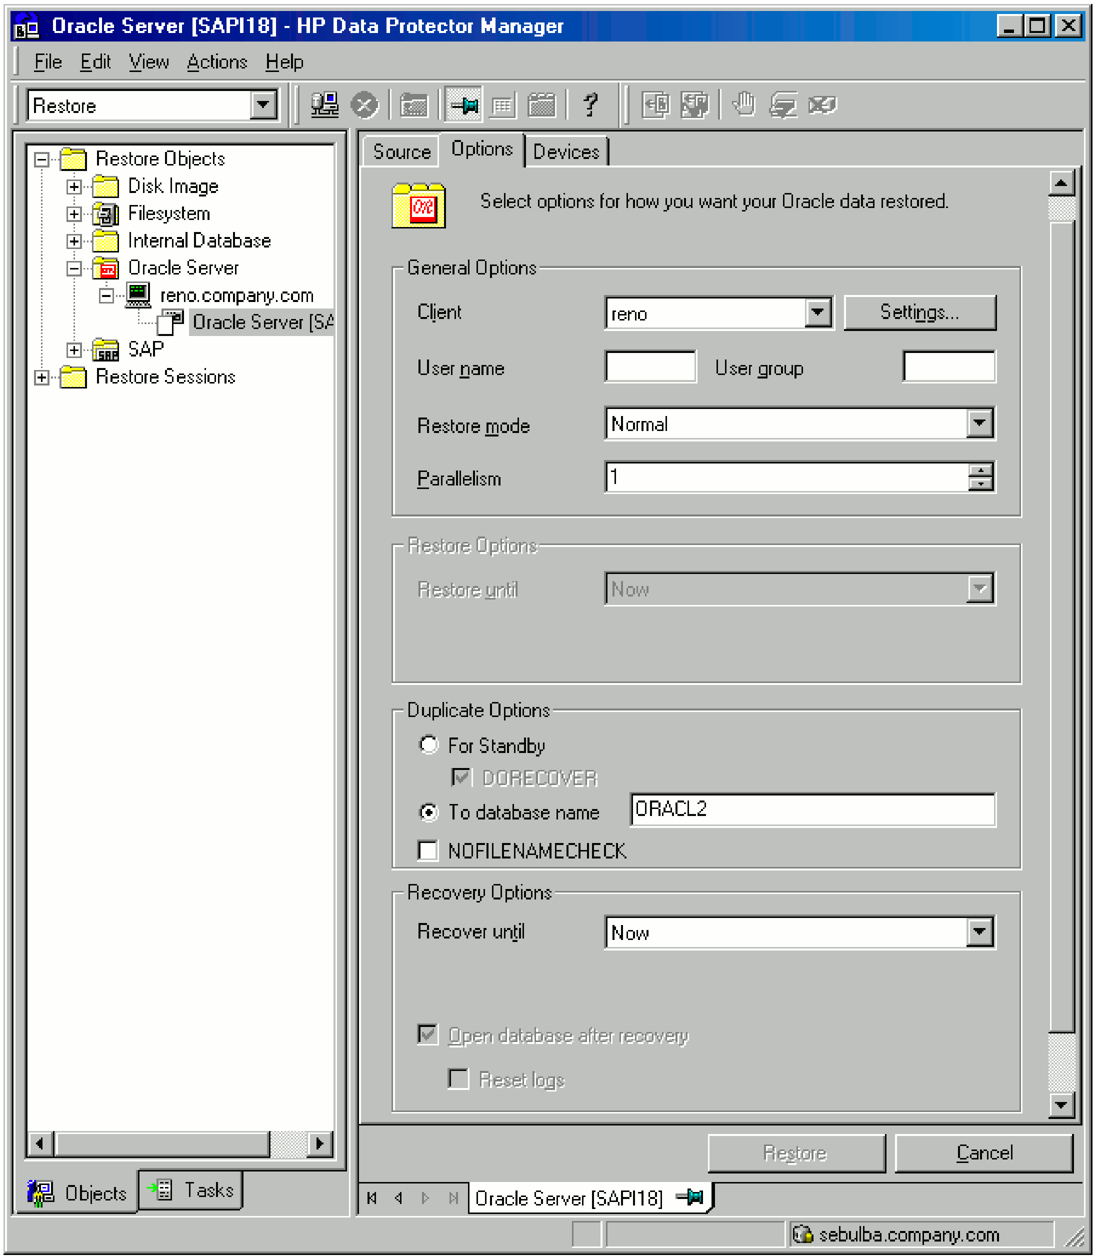

Select options for how you want your Oracle data restored.

General Options

Client: reno

User name:

User group:

Restore mode: Normal

Parallelism: 1

Restore Options

Restore until: Now

Duplicate Options

○ For Standby

☑ DORECOVER

● To database name: ORACL2

□ NOFILENAMECHECK

Recovery Options

Recover until: Now

☑ Open database after recovery

□ Reset logs

10. Click **Restore**.

    When the standby database is created, it is left mounted. Start the managed recovery process (log apply services) manually. For information on using the RMAN commands to duplicate a database, see the Oracle documentation.

Instant recovery and database recovery

The Data Protector instant recovery functionality is used only to restore the target volumes on which the database files are located. The database recovery part is performed after instant recovery by the RMAN utility. During database recovery, incremental backups and archive log backups performed after ZDB to disk or ZDB to disk+tape are restored from tape. Only those archive logs that do not reside on the target volumes are restored.

NOTE: Instant recovery is not supported for LUNS exported as volume set although ZDB-to-tape and disk+tape are supported for the same.

IMPORTANT: If the Oracle control file, online redo logs, and SPFILE are on the same source volumes as datafiles and if you enable instant recovery by setting the `omnirc` options, the control file, SPFILE, and online redo logs are overwritten during the instant recovery.

Prerequisites

- The control file that reflects the internal database structure at the time of backup must be available on the application system. If necessary, restore the appropriate control file from a tape backup.
- The recovery catalog must be open.

RAC preparation steps

In case of RAC, set the following option in the `omnirc` file:

```
ZDB_IR_VGCHANGE=vgchange -a s
```

The instant recovery procedure is the same as without RAC. However, if you need to perform instant recovery to some other node than the one that was backed up, perform the following procedure before the standard instant recovery procedure:

1. Make sure that the MC/SG virtual package is running on the target node.
2. Set the `OB2BARHOSTNAME` environment variable as the virtual hostname before running the configuration from the command line:

   ```
   export OB2BARHOSTNAME=virtual_hostname
   ```

Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Shut down the Oracle database instance using `sqlplus`. In case of RAC, shut down all instances.

   For example:
   ```
   sqlplus
   sql> shutdown immediate
   sql> exit
   ```

2. In the Context List, click **Instant Recovery**.
3. Expand **Oracle Server** and select the ZDB-to-disk or ZDB-to-disk+tape session from which you want to perform the restore.
4. In the Source tab, select the objects to recover. Only whole databases can be selected. Set other options as required for HP 3PAR.
5. At this point, you can decide whether to perform a database recovery immediately after an instant recovery or not:

   - To perform only an instant recovery, click **Restore**.

     NOTE: You can perform a database recovery at a later time either from the Data Protector Manager Restore Context or manually using the RMAN CLI.

   - To perform a database recovery immediately after an instant recovery, click the **Options** tab, select **Recovery**, and then select the database recovery options. For a recovery until a selected time, logseq/thread number, or SCN number, it is recommended to reset the log files.

6. Click **Restore** or **Preview**.

   NOTE: The Preview feature only checks if the replica can be restored. It does not check if the database recovery will be successful.

Data Protector recovers the database after completing the following tasks:

- Performing instant recovery by switching the database to the mount state
- Restoring the necessary incremental backups and archived redo logs from tape
- Applying the redo logs

If you reset the logs, reset the database too. Otherwise, during the next backup, Oracle will try to use the logs that were already reset and the backup will fail. Login to the target and recovery catalog database and execute:

```
rman target Target_Database_Login catalog Recovery_Catalog_Login
RMAN> RESET DATABASE;
RMAN> exit
```

Oracle database recovery after the instant recovery

To recover the Oracle database after the instant recovery has been performed, perform the following steps:

1. Change the Oracle database to the mount state by connecting to the target database using `sqlplus` and then running the following command:

   ```
   startup mount
   ```

2. To recover the database, you can use one of the following options:

   - Perform a recovery from the Data Protector Manager Restore Context using the following procedure:

     a) Expand **Oracle Server** and select the database to recover. In the Source tab, under Restore action, select **Perform recovery only**.

Figure 26: Selecting the database for recovery

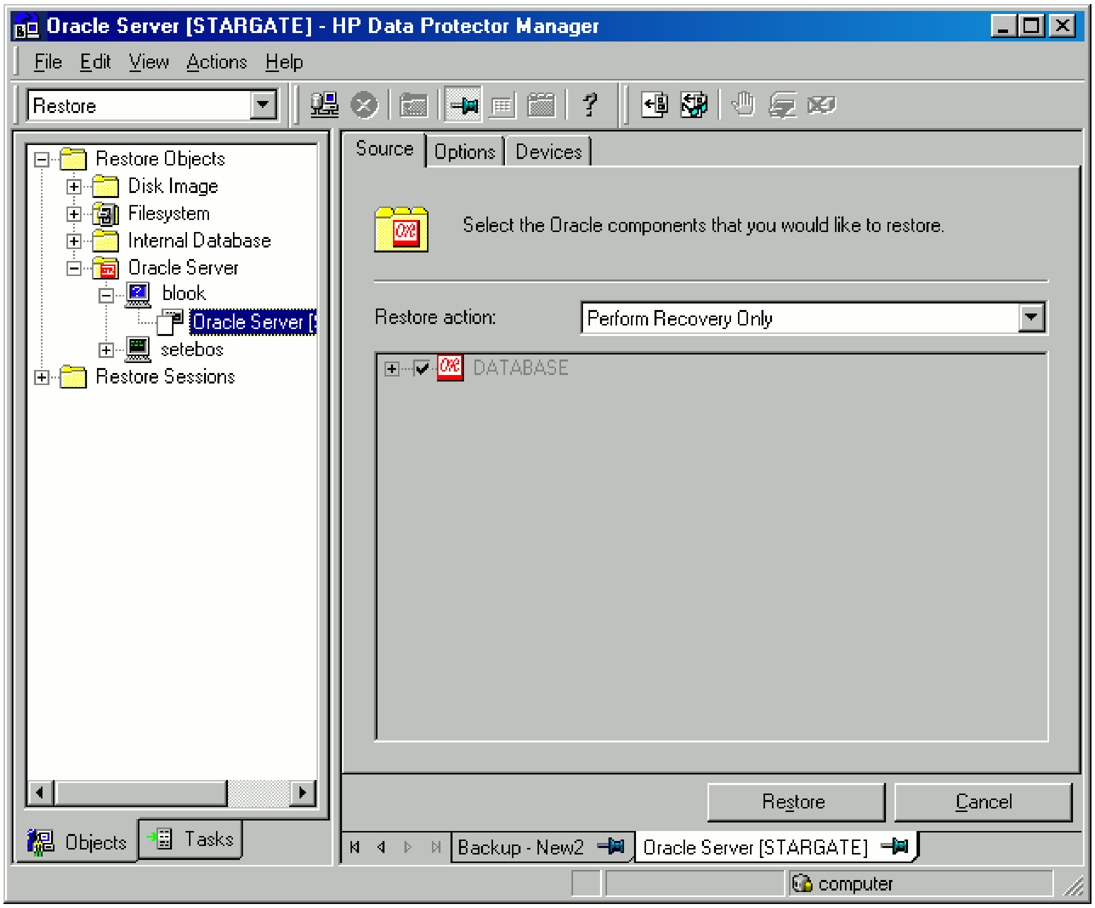

Select the Oracle components that you would like to restore.

Restore action: Perform Recovery Only

b) In the Options tab, select the recovery options.

c) Click **Restore**.

- Perform a manual database recovery using RMAN. Run the following RMAN script to recover the database:

```
run {
allocate channel dev1 type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
recover database;
sql 'alter database open';
release channel dev1;
}
```

For additional examples on how to recover the database after an instant recovery, see Appendix A – Examples of Restoring Oracle using RMAN.

Aborting sessions

You can abort currently running sessions by clicking the **Abort** button. During a session, if RMAN or SQL*Plus does not respond, Data Protector automatically aborts the session. By default, Data Protector waits for the response for 5 minutes. Using `omnirc` options or environment variables, `OB2_RMAN_COMMAND_TIMEOUT` and `OB2_SQLP_SCRIPT_TIMEOUT`, you can modify this time interval.

For details on setting environment variables, see the Setting environment variables section.
For details on setting the corresponding `omnirc` options, see the *HP Data Protector Help* index: “omnirc option”. Note that environment variables override `omnirc` options.

Appendix A – Examples of restoring Oracle using RMAN

Data Protector acts as a media management software for the Oracle system; therefore, RMAN can be used for a restore. This appendix contains only examples of how you can perform a restore. The examples provided do not apply to all situations where a restore is needed.

See the Oracle Recovery Manager User's Guide and References for detailed information on how to perform:

- Restore and recovery of the database, tablespace, control file, and datafile
- Duplication of a database

This section includes the following examples of restore:

- Example of full database restore and recovery
- Example of point-in-time restore

For additional examples, see the *HP Data Protector Integration Guide*.

Preparing the Oracle database for restore

You can perform the restore of an Oracle database when the database is in the mount mode. However, when you are performing the restore of tablespaces or datafiles, only a part of the Oracle database can be put offline.

Prerequisites

The following requirements must be met before you start a restore of an Oracle database:

- Ensure that the recovery catalog database is open. If the recovery catalog database cannot be brought online, you may need to restore the recovery catalog database. For information on restoring the recovery catalog database, see the Restoring the recovery catalog database section.
- Check which ZDB method (proxy-copy or backup set) was used for the backup session that you plan to restore.
- Control files must be available. If the control files are not available, you must restore them. For more information, see the *Oracle Recovery Manager User's Guide and References*.

  If you have to perform a restore of the recovery catalog database, you must perform this restore first. Only then can you perform a restore of other parts of the Oracle database. When you are sure that the recovery catalog database files are in place, start the recovery catalog database.

- Ensure that the following environment variables are set:
  - ORACLE_BASE
  - ORACLE_HOME
  - ORACLE_TERM
  - DB NAME
  - PATH
  - NLS_LANG
  - NLS_DATE_FORMAT

  Example

  ```
  ORACLE_BASE=/opt/oracle
  ORACLE_HOME=/opt/oracle/product/10.1.0
  ORACLE_TERM=HP
  ```

```
DB_NAME=PROD
PATH=$PATH:/opt/oracle/product/10.1.0/bin
NLS_LANG=american
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

- Verify that the `/etc/oratab` file has the following line:

  Windows systems: `PROD:Oracle_home\product\10.1.0:N`
  UNIX systems: `PROD:/opt/oracle/product/10.1.0:N`

  The last letter determines whether the database will automatically start upon boot-up (Y) or not (N).

Connection strings used in the examples

In the examples below, the following connection strings are used:

- Target connection string for target database:

  ```
  sys/manager@PROD
  ```

  where, `sys` is the username, `manager` is the password and `PROD` is a net service name.

- Recovery catalog connection string for recovery catalog database:

  ```
  rman/rman@CATAL
  ```

  where, `rman` is the username and password and `CATAL` is a net service name.

SBT_LIBRARY parameter

On Windows and UNIX systems, set the `SBT_LIBRARY RMAN` script parameter to point to the correct platform-specific Data Protector MML. The parameter must be specified for each RMAN channel separately. For details on the Data Protector MML location, see the *HP Data Protector Integration Guide*.

In the following examples, the `SBT_LIBRARY` parameter is set to `/opt/omni/lib/libob2oracle8.so`, which is the correct path for 32-bit Solaris systems.

Example of full database restore and recovery

To perform a full database restore and recovery, you also need to restore and apply all the archive logs. To perform a full database restore and recovery:

1. Log in to the Oracle RMAN:

   Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
   UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the full database restore and recovery:

For a non-ZDB or ZDB backup set session, use the following:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

For a ZDB proxy-copy session, use the following:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

You can also save the script into a file and perform a full database restore using the saved files. The procedure in such cases is as follows:

1. Create a file restore_database in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.

2. Start the full database restore:

   Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_datafile`

   UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_datafile`

Example of point-in-time restore

To perform a point-in-time restore, you also need to restore and apply the archive logs to the specified point-in-time. To perform a point-in-time database restore and recovery:

1. Log in to the Oracle RMAN:

   Windows systems: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`

   UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

2. Start the point-in-time restore:

   For a non-ZDB or ZDB backup set session, use the following:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
set until time 'Mar 14 2004 11:40:00';
restore database;
recover database;
sql 'alter database open';
release channel 'dev1';
}
```

For a ZDB proxy-copy session, allocate one channel for restoring proxy-copy sessions and one channel for database recovery. Release the proxy-copy channel before the recovery using the following:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
allocate channel 'dev2' type 'sbt_tape' parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME)';
set until time 'Mar 14 2006 11:40:00';
restore database;
release channel 'dev1';
recover database;
sql 'alter database open';
release channel 'dev2';
}
```

3. After you have performed a point-in-time restore, reset the database in the Recovery Catalog.

You can also save the script into a file and perform a point-in-time restore using the saved files:

1. Create a file `restore_PIT` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` (Windows systems) directory.
2. Start the point-in-time restore:

   Windows systems: ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_PIT

   UNIX systems: ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_PIT

Appendix B – Examples of pre-exec and post-exec scripts

Pre-exec example

The following is an example of a script that shuts down an Oracle instance:

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/sqlplus ]; then
$ORACLE_HOME/bin/sqlplus << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
shutdown
EOF
echo "Oracle database \"$DB_NAME\" shut down."
exit 0
else
echo "Cannot find Oracle SQLPLUS ($ORACLE_HOME/bin/sqlplus)."
exit 1
fi
```

Post-exec example

The following is an example of a script that starts an Oracle instance:

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/sqlplus ]; then
$ORACLE_HOME/bin/sqlplus << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
startup
EOF
echo "Oracle database \"$DB_NAME\" started."
exit 0
else
echo "Cannot find Oracle SQLPLUS ($ORACLE_HOME/bin/sqlplus)."
exit 1
fi
```

Appendix B – Limitations

3PAR ZDB agent

- 3PAR volumes connected using iSCSI are not supported.
- Snapshots of snapshots is not supported.
- 3PAR remote copy is not supported.

Instant recovery

- Source volumes presented as the 3PAR volume set are not supported.
- Source volumes presented using the HostSet and Port-Present presentation types are not supported.
- On Linux, only two host configurations are supported, that is, application and backup servers must not be the same host.
- The target volumes from which the copy back needs to happen should not be presented to any host.
- No support for “Retain source for forensics“.