



Hewlett Packard
Enterprise

Data Protector

Software Version: 10.02

Zero Downtime Backup Integration Guide

Document Release Date: December 2017

Software Release Date: December 2017

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation updates

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

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

To check for recent software updates, go to <https://softwaresupport.hpe.com/patches>.

To verify that you are using the most recent edition of a document, go to <https://softwaresupport.hpe.com/manuals>.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

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

Support

Visit the HPE Software Support Online web site at <https://softwaresupport.hpe.com>.

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

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

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests

- Download software patches
- Access product documentation
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

To find more information about access levels, go to <https://softwaresupport.hpe.com/web/softwaresupport/access-levels>.

Contents

Chapter 1: Data Protector Oracle Server ZDB integration	15
Introduction	15
Backup and restore types	16
Integration concepts	18
Oracle backup set ZDB concepts	22
Backup process	24
Oracle proxy-copy ZDB concepts	26
Backup process	27
Configuring the integration	30
Prerequisites	30
Limitations	31
Before you begin	33
Backup set method	34
Configuring the HPE P6000 EVA Disk Array Family in ASM environments	36
Configuring the HPE P9000 XP Disk Array Family in ASM environments	36
Cluster-aware systems	37
Linking Oracle Server with the Data Protector MML	37
Oracle 12c CDB and PDB mode support	37
Configuring Oracle user accounts	38
Configuring Oracle operating system user accounts	38
Clusters	39
Configuring Oracle database user accounts	40
Configuring Oracle databases	41
Oracle OS authentication	41
Using the Data Protector GUI	42
Using the Data Protector CLI	45
Checking the configuration	49
Using the Data Protector GUI	50
Using the Data Protector CLI	50
Handling errors	50
Checking configuration for instant recovery	50
Setting environment variables	51
Using the Data Protector GUI	52
Using the Data Protector CLI	53
Switching between Oracle backup methods	53
Backup	54
Creating backup specifications	54
Examples of pre-exec and post-exec scripts on UNIX systems	68
Editing the Oracle RMAN script	69
Starting backup sessions	72

- Considerations 72
 - Scheduling backup sessions 72
 - Running an interactive backup 73
 - Starting a backup using the GUI 73
 - Starting a backup using the CLI 74
- Restore 74
 - Prerequisites 76
 - Restoring from backup media to the application system on LAN 76
 - Restoring Oracle using the Data Protector GUI 76
 - Restoring database items in a disaster recovery 76
 - Changing the database state 77
 - Restoring the recovery catalog database 77
 - Restoring the control file 79
 - Restoring Oracle database objects 80
 - Restoring tablespaces and datafiles 84
 - Duplicating an Oracle database 84
 - Restore, recovery, and duplicate options 86
 - Restore action options 86
 - General options 87
 - Duplicate options 88
 - Restore and recovery options 89
 - Restoring Oracle using RMAN 90
 - Preparing the Oracle database for restore 91
 - Connection strings used in the examples 92
 - SBT_LIBRARY parameter 92
 - Example of full database restore and recovery 93
 - Example of point-in-time restore 94
 - Example of tablespace restore and recovery 95
 - Example of datafile restore and recovery 97
 - Example of archive log restore 100
 - Example of database restore using a different device (with the automatic device selection functionality disabled) 101
 - Restoring using another device 101
 - Restoring Oracle to different cell/client 102
 - Instant recovery and database recovery 102
 - Instant recovery using the Data Protector GUI 103
 - Oracle database recovery after the instant recovery 105
 - Oracle in Veritas Cluster instant recovery 107
- Aborting sessions 108
- Troubleshooting 108
 - Before you begin 108
 - Checks and verifications 108
 - Problems 114

- Chapter 2: Data Protector SAP R/3 ZDB integration 125

- Introduction 125
- Integration concepts 126
 - ZDB flow 127
- Data Protector SAP R/3 configuration file 128
 - Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI 130
- Configuring the integration 132
 - Prerequisites 133
 - Before you begin 134
 - Cluster-aware clients 134
 - Configuring user accounts 134
 - Configuring SQL*Net V2 or Net8 TNS listener 135
 - Checking the connection 136
 - Authentication password file 137
 - Enabling archived logging 137
 - Sharing directories on the application system 138
 - UNIX application system 139
 - Windows application system 139
 - Choosing authentication mode 140
 - Configuring SAP R/3 databases 140
 - Before you begin 140
 - Using the Data Protector GUI 140
 - Using the Data Protector CLI 144
 - Handling errors 145
 - Checking the configuration 145
 - Using the Data Protector GUI 146
 - Using the Data Protector CLI 146
 - Configuring the SAP R/3 parameter file 147
- Backup 147
 - Considerations 148
 - Creating backup specifications 149
 - Modifying backup specifications 157
 - Scheduling backup sessions 157
 - Previewing backup sessions 157
 - Using the Data Protector GUI 157
 - Using the Data Protector CLI 157
 - What happens during the preview? 158
 - Starting backup sessions 158
 - Backup methods 158
 - Using the Data Protector GUI 158
 - Using the Data Protector CLI 158
 - Using the SAP BRTOOLS 159
 - Configuring SAP compliant ZDB sessions 160
 - Using the Data Protector GUI 160
 - Using the Data Protector CLI 161

Manual balancing	161
Restore	162
Considerations	164
Standard restore	164
Instant recovery	167
Considerations	167
Instant recovery using the Data Protector GUI	167
Database recovery options	169
Instant recovery using the Data Protector CLI	170
Instant recovery from replicas containing the control file	171
Restoring using another device	171
Using the Data Protector GUI	172
Using the Data Protector CLI or SAP commands	172
Localized SAP R/3 objects	172
Monitoring sessions	172
Troubleshooting	173
Before you begin	173
General troubleshooting	173
Prerequisites on the SAP side of the integration	173
Configuration problems	174
Backup problems	176
Restore problems	177
Prerequisites on the SAP side of the integration	179
Configuration problems	180
Backup problems	181
Restore problems	182
Verifying the prerequisites (Oracle side)	185
Verifying the prerequisites (SAP side)	186
Verifying the configuration	186
Verifying the backup configuration	187
Verifying restore	188
Configuration and backup problems	190
Restore problems	193
Chapter 3: Data Protector Microsoft SQL Server ZDB integration	194
Introduction	194
Integration concepts	195
Configuring the integration	196
Prerequisites	196
Before you begin	196
Data Protector SQL Server configuration file	197
Configuring an SQL Server cluster	197
Configuring SQL Server instances	198
Using the Data Protector GUI	198

- Using the Data Protector CLI 200
- Changing and checking configuration 200
 - Using the Data Protector GUI 200
 - Using the Data Protector CLI 201
- Backup 201
 - Creating ZDB specifications 202
 - SQL Server-specific backup options 212
 - Scheduling backups 214
 - Starting backup sessions 214
 - Using the Data Protector GUI 214
 - Using the Data Protector CLI 214
- Restore 215
 - Before you begin 215
 - Standard restore 215
 - Restore options 218
 - Restoring to a different SQL Server instance or/and different SQL Server 220
 - Instant recovery 221
- Monitoring sessions 223
- Troubleshooting 223
 - Before you begin 224
 - Checks and verifications 224
 - Problems 225

Chapter 4: Data Protector Microsoft Exchange Server 2010+ ZDB integration 229

- Introduction 229
- Integration concepts 230
 - Supported environments 230
 - Standalone environments 230
 - DAG environments 231
- Configuring the integration 233
 - Prerequisites 233
 - Limitations 235
 - Before you begin 235
 - Configuring user accounts 235
 - Windows domain user account for backup and restore sessions 235
 - User account for executing Exchange Management cmdlet operations 236
- Backup 236
 - Backup types 236
 - Microsoft Exchange Server backup types 237
 - ZDB backup types 237
 - VSS backup types 238
 - Backup parallelism 238

- Replica rotation in DAG environments 238
- Backup considerations 239
- Object operations considerations 240
- Creating backup specifications 240
- Modifying backup specifications 249
- Scheduling backup sessions 249
- Previewing backup sessions 250
 - Using the Data Protector GUI 250
 - Using the Data Protector CLI 250
 - What happens during the preview? 250
- Starting backup sessions 250
 - Using the Data Protector GUI 250
 - Using the Data Protector CLI 251
- Backup objects 251
- Restore 252
 - Restore methods 252
 - Repair all passive copies with failed status 253
 - Restore to the latest state 253
 - Restore to a point in time 253
 - Restore to a new mailbox database 254
 - Restore files to a temporary location 254
 - Restore destination 254
 - Restoring to a standalone database 254
 - Restoring to an active copy 255
 - Restoring to a passive copy 255
 - Restoring data to a new database 255
 - Restoring data to a temporary location 256
 - Instant recovery in DAG environments 256
 - Restore chain 256
 - Restore chain during instant recovery 257
 - Restore parallelism 257
 - Finding information for restore 257
 - Using the Data Protector GUI 257
 - Using the Data Protector CLI 258
 - Standard restore 258
 - Restoring using the Data Protector GUI 258
 - Restoring using the Data Protector CLI 265
 - Restoring using another device 268
 - Instant recovery 268
 - Performing instant recovery using the Data Protector GUI 268
 - Performing instant recovery using the Data Protector CLI 274
 - Restore options 276
- Monitoring sessions 281
- Troubleshooting 281
 - Before you begin 281

Checks and verifications	282
Problems	282
Chapter 5: Data Protector Microsoft SharePoint Server VSS based solution ..	286
Introduction	286
Backup	286
Limitations	287
Restore	287
Installation and configuration	287
ZDB prerequisites	287
Microsoft Office SharePoint Server 2007	287
Microsoft SharePoint Server 2010	287
Licensing	288
Installing the integration	288
Configuring the integration	290
Configuring user accounts	290
Backup	290
Prerequisites	290
Limitations	290
Recommendations	290
How the command works	291
Microsoft SharePoint Server 2010	291
Microsoft SharePoint Server 2013	291
Considerations	292
The command syntax	292
Option description	293
Starting Windows PowerShell	296
Creating backup specifications (examples)	298
Modifying backup specifications	299
Source page	299
Destination page	299
Options page	300
Starting backup sessions (examples)	300
Scheduling backup sessions	300
Restore	301
Before you begin	302
Restoring data	303
Considerations	303
Prerequisites	304
Restoring using the Data Protector GUI	304
Restoring using the Data Protector CLI	304
Limitations	305
After the restore	305
Troubleshooting	305

- Before you begin 306
- Checks and verifications 306
- After restore, you cannot connect to the Central Administration webpage 306
- Backup fails with the error Failed to resume Service Windows SharePoint Services Help Search 307
- After restore, a quiesce operation fails 307
- After restore, you cannot connect to the FAST Search Server 308
- The SharePoint_VSS_backup.ps1 script stops responding and the farm stays in read only mode 308
- SharePoint Search service application not operational after restore 309

Chapter 6: Data Protector Virtual Environment ZDB integration for VMware ..311

- Introduction 311
 - Recommendations 312
- Integration concepts 312
 - Supported environments 312
 - Backup process 314
 - Backup concepts 315
 - What is backed up? 315
 - Virtual machines 315
 - Virtual machine templates 316
 - vStorage Image backup method 316
 - Snapshot management 316
 - Backup types 317
 - Changed block tracking 318
 - Non-Changed Block Tracking (Non-CBT) backup 321
 - Quiescence 322
 - Prerequisites 323
 - Limitations 323
 - Considerations for Quiescence Operations 323
 - Disk space requirements 324
 - Free space required option 324
 - Backup disk buffer 325
 - Backup parallelism 326
 - Backup considerations 326
- Restore concepts 327
 - Restore of VMware objects backed up with vStorage Image method 327
 - Restore to a datacenter 327
 - Restore to a directory 328
 - Restore of Nova Instances and Shadow VMs backed up with vStorage Image + Openstack method 329
 - Restore chain 329
 - Power On and Live Migrate 330
 - Restore considerations 331
 - Power On considerations 332

StoreOnce Recovery Manager Central Integration	333
Configuring the integration	335
Recommendations	335
Prerequisites	335
Before you begin	336
Importing and configuring VMware clients	336
Changing the configuration of VMware clients	339
Using the Data Protector GUI	339
Using the Data Protector CLI	340
Checking the configuration of VMware clients	341
Using the Data Protector GUI	341
Using the Data Protector CLI	341
Configuring virtual machines	342
Using the Data Protector GUI	342
Using the Data Protector CLI	345
Customizing the Data Protector behavior with omnirc options	346
Adding the RMC Server details in Data Protector using the Command Line Interface	346
Backup	347
Backup limitations	347
vStorage Image + OpenStack backup method limitations	349
Creating backup specifications	349
Creating backup specifications for RMC backups	355
Modifying backup specifications	358
Scheduling backup sessions	359
Previewing backup sessions	359
Using the Data Protector GUI	359
Using the Data Protector CLI	359
What happens during the preview?	359
Starting backup sessions	359
Using the Data Protector GUI	360
Using the Data Protector CLI	360
Preparing for disaster recovery	360
Restore	361
Restore limitations	361
vStorage Image + OpenStack restore limitations	363
Power On and Live Migrate limitations	363
Finding information for restore	364
Using the Data Protector GUI	364
Using the Data Protector CLI	366
Restoring using the Data Protector GUI	366
Restoring using the Data Protector CLI	378
Recovering virtual machines manually	381
Recovering virtual machines after restore to a directory	381
Recovering with the VM configuration file in the VMX format	381
Recovering with the VM configuration file in the XML format	387

- Recovering virtual machines after restore to a datacenter387
- Restoring using another device 388
- Cleaning up a datastore after a failed restore 388
- Disaster recovery 388
- Instant recovery 389
 - HPE 3PAR ZDB Instant recovery 389
- Monitoring sessions 390
- Troubleshooting 390
 - Before you begin 390
 - Checks and verifications 390
 - Problems 391
- Chapter A: Appendix 406
 - Reconfiguring an Oracle instance for instant recovery 406
 - Examples for moving the control files and redo logs to different locations 407
 - ZDB integrations omnirc options 409
- Send documentation feedback 414

Chapter 1: Data Protector Oracle Server ZDB integration

Introduction

You can employ a variety of backup strategies to best meet your system priorities. If database availability is the highest priority, for instance, your backup strategy should include online backups that are performed frequently to minimize recovery time. This strategy limits downtime, but uses system resources more intensively. The Data Protector zero downtime backup (ZDB) functionality offers online backup capabilities with minimal degradation of the application system performance.

Supported disk arrays

The following disk arrays can be used for zero downtime backup (ZDB) of the Oracle Server data:

- HPE P6000 EVA Disk Array Family (P6000 EVA Array)
- HPE P9000 XP Disk Array Family (P9000 XP Array)
- EMC Symmetrix (EMC)
- Non-HPE Storage Arrays (NetApp storage, EMC VNX, and EMC VMAX storage families)

NOTE:

With the Data Protector EMC and Data Protector non-HPE Storage Array integrations, instant recovery is not supported, and ZDB to tape is the only supported ZDB form.

In Oracle Server configurations that use Automatic Storage Management (ASM), zero downtime backup and instant recovery are supported with the Data Protector P6000 EVA Array and Data Protector P9000 XP Array integrations.

Advantages

The advantages of using Data Protector Oracle ZDB integration are:

- ZDB reduces the performance degradation of the application system.
- The tablespaces are in 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, at leisure, 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. These differences are indicated where appropriate.

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

- 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.

NOTE:

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.

IMPORTANT:

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. For more information on Oracle archive log backup with Data Protector see the *HPE Data Protector Integration Guide*.

NOTE:

On EMC, decision support, application testing, and similar tasks are possible only if the Oracle binaries are installed on the backup system as well. In most cases, however, the Data Protector EMC integration requirement is that application binaries are installed on the application system only.

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

- Using the instant recovery functionality and RMAN on the application system, you can:
 - perform a full database restore and database recovery
 - perform recovery from incremental backup (for ZDB to tape or ZDB to disk+tape)
 - perform recovery from a chain of incremental backups (for ZDB to tape or ZDB to disk+tape)
 - restore a datafile to a location other than its original one

Oracle recovery methods, below provides an overview of recovery methods, depending on the type of backup that was performed and type of recovery required.

Oracle recovery methods

Disk array	Backup types	Recover the whole database until		Recover a part of database until now
		Now	A point in time, logseq/thread, or SCN number	
P9000 XP, P6000 EVA, EMC, non-HPE Storage Arrays	ZDB to tape - online	Restore	Restore	Restore
	ZDB to tape - offline	Restore	Restore ¹	Restore
P9000 XP, P6000 EVA	ZDB to disk - online	Instant recovery+ database recovery	Instant recovery+ database recovery	N/A
	ZDB to disk - offline	Instant recovery	Instant recovery+ database recovery ²	N/A
	ZDB to disk+tape - online	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database 	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database 	Restore

¹ The database must be put in archive mode
² The database must be put in archive mode

Disk array	Backup types	Recover the whole database until		Recover a part of database until now
		Now	A point in time, logseq/thread, or SCN number	
		recovery	recovery	
	ZDB to disk+tape - offline	<ul style="list-style-type: none"> Restore or Instant recovery 	<ul style="list-style-type: none"> Restore or Instant recovery+ database recovery¹ 	Restore

Legend

<i>Restore</i>	Use the Data Protector GUI or RMAN scripts to restore the database from backup media to the application system on LAN.
<i>Instant recovery + database recovery</i>	The following three options are possible: <ul style="list-style-type: none"> Perform instant recovery followed by database recovery from the Data Protector Instant Recovery GUI context or Perform instant recovery first and then perform database recovery from the Data Protector Restore GUI context or Perform instant recovery first and then use RMAN scripts to recover the database.
<i>Instant recovery</i>	Perform instant recovery without database recovery.

See the *HPE Data Protector Concepts Guide* for an overview of ZDB concepts and terminology.

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 a 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:

¹The database must be put in archive mode

- 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)

How does the integration work?

`ob2rman.pl` executes RMAN, which directs the Oracle server processes on the target database to perform backup, restore and recovery. RMAN maintains the required information about the target databases in the recovery catalog, the Oracle central repository of information, and in the control file of a particular target database.

The main information which `ob2rman.pl` provides to RMAN is:

- Number of allocated RMAN channels
- RMAN channel environment parameters
- Information on the database objects to be backed up or restored

For backup, `ob2rman.pl` uses the Oracle target database views to get information on which logical (tablespaces) and physical (datafiles) target database objects are available for backup.

For restore, `ob2rman.pl` uses current control file or recovery catalog (if used) to get information on which objects are available for restore.

Using the Data Protector integration with RMAN, you can back up and restore the Oracle control files, datafiles, and archived redo logs.

The interface from the Oracle server processes to Data Protector is provided by the Data Protector Oracle integration Media Management Library (**MML**), which is a set of routines that allows the reading and writing of data to General Media Agents.

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

A backup that includes all datafiles and current control file that belong to an Oracle Server instance is known as a whole database backup.

These features can be used for online or offline backup of the Oracle target database. However, you must ensure that the backup objects (such as tablespaces) are switched into the appropriate state before and after a backup session. For online backup, the database instance must operate in the ARCHIVELOG mode; whereas for offline backup, objects need to be prepared for backup using the Pre-exec and Post-exec options in the backup specification.

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

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

Backup and restore of an Oracle target database can be performed using the Data Protector User Interface or the RMAN utility.

The heart of the Data Protector Oracle integration is MML, which enables an Oracle server process to issue commands to Data Protector for backing up or restoring parts or all of the Oracle target database files. The main purpose is to control direct interaction with media and devices.

Non-ZDB flow

A Data Protector scheduled or interactive backup is triggered by the Data Protector Backup Session Manager, which reads the backup specification and starts the `ob2rman.pl` command on the Oracle Server under the operating system user account specified in the backup specification. Further on, `ob2rman.pl` prepares the environment to start the backup, and issues the RMAN backup command. RMAN instructs the Oracle Server processes to perform the specified command.

The Oracle Server processes initialize the backup through MML, which establishes a connection to the Data Protector Backup Session Manager. The Backup Session Manager starts the General Media Agent, sets up a connection between MML and the General Media Agent, and then monitors the backup process.

The Oracle Server processes read the data from the disks and send it to the backup devices through MML and the General Media Agent.

RMAN writes information regarding the backup either to the recovery catalog (if one is used) or to the control file of the Oracle target database.

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

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

Restore flow

A restore session can be started using:

- Data Protector GUI
- RMAN CLI

You must specify which objects are to be restored.


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

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

The General Media Agent reads the data from the backup devices and sends it to the Oracle Server processes through MML. The Oracle Server Processes write the data to the disks.

The concept of Oracle integration, data and the control flow are shown in [Data Protector Oracle integration concept, on the next page](#), and the related terms are explained in the following table.

Data Protector Oracle integration concept



Database files can also be managed by **Automatic Storage Management (ASM)**.

Legend

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

Oracle backup set ZDB concepts

See the *HPE Data Protector Concepts Guide* for a general description of ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape and instant recovery concepts.

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 MML.

Depending on the location of the Oracle control file, online redo log files, and SPFILE, the following two options are possible:

- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.
By default, instant recovery for such a configuration is enabled.
- Oracle control file, online redo log files, SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle datafiles.


By default, instant recovery for such a configuration is *not* enabled. You can enable instant recovery by setting the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IRomnirc` options to 1. See [ZDB integrations omnirc options, on page 409](#).

IMPORTANT:

If you enable instant recovery by setting the above mentioned options, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery.

The Oracle archived redo log files do not have to reside on source volumes.

Oracle backup set ZDB concept



Oracle backup set ZDB concept, on the previous page presents only the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on a different volume group (if LVM is used) or source volume than Oracle datafiles. Oracle database files can also be managed by ASM, however some limitations apply to Oracle ASM configurations. For details, see [Limitations, on page 31](#).


For more information on an alternative Oracle backup and restore concept, see [ZDB integrations omnirc options, on page 409](#).

Legend

<i>MA</i>	The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.
<i>SM</i>	The session manager controls backup and restore sessions and writes session information to the IDB.
<i>Disk Array Agent</i>	The disk array agents (ZDB agents) are SYMA (for EMC), SSEA (for P9000 XP Array), and SMISA (for P6000 EVA Array and non-HPE Storage Arrays).
<i>Data Protector MML</i>	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector. This is a Data Protector software library that is linked to the Oracle software.

Backup process

Oracle backup set ZDB flow
 application system



NOTE:

ZDB agents are SYMA for EMC, SSEA for P9000 XP Array, and SMISA for P6000 EVA Array and non-HPE Storage Arrays.

See the *HPE Data Protector Concepts Guide* for a general description of ZDB and instant recovery concepts.

See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape session flows and for the explanation of actions triggered by ZDB options.

This section provides only the information relevant to the Data Protector Oracle ZDB integration.

Operations on a replica (mounting, activating volume/disk groups,...) described below are dependent on or triggered by ZDB options. See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

- Data Protector executes the `ob2rman.pl` command on the backup system. This command retrieves a list of files or disk images to be backed up from the Oracle database on the application system and starts the resolving process. The list is used only to determine the source volumes to be replicated. If the location for control file copy is specified during configuration, `ob2rman.pl` makes a copy of the control file to the specified directory on the application system. This directory has to reside on a disk array source volume.
- When performing an *online* ZDB session, `ob2rman.pl` then sets the Oracle target database into backup mode by issuing the `sqlplus` command “ALTER TABLESPACE BEGIN BACKUP”, starts the procedure to create a replica of the source volumes on which the database is installed; and after the replica is created, takes the database out of backup mode by issuing the `sqlplus` command “ALTER TABLESPACE END BACKUP”.

When performing an *offline* ZDB session, `ob2rman.pl` shuts down the Oracle database, starts the procedure to create a replica of the source volumes on which the database is installed; and after the replica is created, starts up the Oracle database.

- `ob2rman.pl` then starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled and, unless the database is installed on raw partitions, the mount points with the Oracle database files are mounted.
- A ZDB agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.

NOTE:

There must be nothing already mounted on the mount point concerned on the backup system, or the resolving and backup will fail.

- If a ZDB-to-disk session is being performed, at this point the remaining ZDB options are processed and details of the session are written to the ZDB database. The session then finishes. The following steps in this description are not performed, therefore RMAN is not given any information about ZDB-to-disk session.
- If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the processing continues as follows:
 - `ob2rman.pl` starts the Oracle backup command RMAN on the backup system, and then sends the Oracle RMAN Backup Command Script to the RMAN cmdfile (input command file).
 - RMAN contacts the Oracle database instance on the backup system, which contacts Data Protector via SBT API and initiates a backup.
 - The Oracle database instance on the backup system reads data from the replica and sends it to the Data Protector General Media Agent for writing to the backup device.
 - At the end of data transfer, the backup system is disabled (filesystems are dismounted on all

platforms and volume/disk groups deactivated on UNIX systems) and links are re-established.

- The recovery catalog and the control file are backed up automatically after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE:

A replica of the archive logs is not created; therefore, the archive logs should be backed up from the application system, following the standard Data Protector Oracle archive logs backup procedure.

Oracle proxy-copy ZDB concepts

See the *HPE Data Protector Concepts Guide* for a general description of ZDB-to-disk, ZDB-to-tape, ZDB-to-disk+tape, and instant recovery concepts.

The Data Protector Oracle integration MML supports the Proxy Copy functionality. This enables Data Protector to perform backup using filesystem backup methods.

Depending on the location of the Oracle control file, online redo log files, and SPFILE, the following two options are possible:

- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.

By default, instant recovery is enabled if this option is selected in the GUI.

- Oracle control file, online redo log files, and SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle datafiles.

By default, instant recovery is *not* enabled, even if this option is selected in the GUI. You can enable instant recovery by setting the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IRomnirc` options to 1. See [ZDB integrations omnirc options, on page 409](#).

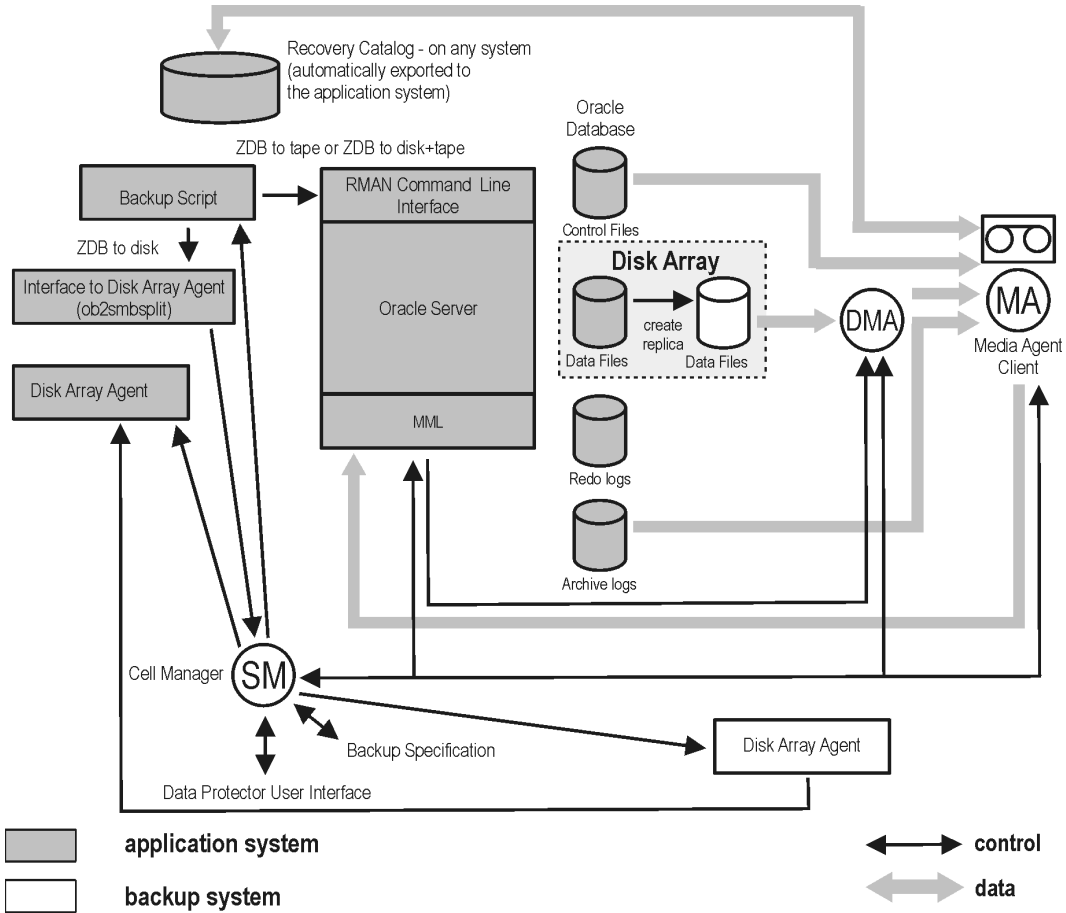
IMPORTANT:

If you enable instant recovery by setting the above mentioned options, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery.

The Oracle archived redo log files do not have to reside on source volumes.

[Oracle proxy-copy ZDB concept, below](#) shows the architecture of the Data Protector Oracle ZDB integration. The figure illustrates the configuration, in which the backup is performed on the backup system. It presents the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on different disk array source volumes than the Oracle data files. For more information on alternative Oracle backup and restore concepts, see [ZDB integrations omnirc options, on page 409](#).


Oracle proxy-copy ZDB concept



<i>MA</i>	The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.
<i>SM</i>	The session manager controls backup and restore sessions and writes the session information to the IDB.
<i>Disk Array Agent</i>	The disk array agents (ZDB agents) are SYMA (for EMC), SSEA (for P9000 XP Array), and SMISA (for P6000 EVA Array and non-HPE Storage Arrays).
<i>MML</i>	The Data Protector Oracle integration Media Management Library, which is a set of routines that enables data transfer between the Oracle Server and Data Protector. This is a Data Protector software library that is linked to the Oracle software.

Backup process

Oracle proxy-copy ZDB flow



See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape sessions flows and for an explanation of actions triggered by ZDB options.

This section provides only the information relevant to the Data Protector Oracle ZDB integration.

Operations on a replica (mounting, activating volume/disk groups...) described below are dependent on or triggered by ZDB options. See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

- In the case of an *offline* ZDB-to-disk+tape or ZDB-to-tape session, `ob2rman.pl` shuts down and opens the database instance in mount state. For both, offline and online ZDB-to-disk+tape or ZDB-to-tape sessions, Data Protector starts RMAN in proxy-copy mode.

In the case of an *offline* ZDB-to-disk session, the database is shut down.

- Data Protector retrieves a list of files or disk images to be included in the replica creation from the Oracle database and starts the resolving process. The list is used only to determine the source volumes to be replicated.

In the case of a *ZDB-to-disk* session, if the location for control file copy is specified during configuration, Data Protector makes a copy of the control file to the specified directory on the application system. This directory has to reside on a disk array source volume.

- In the case of an *online* backup, the Oracle target database is switched into the backup mode.
- `ob2smbsplit` or MML starts the procedure to create a replica of the source volumes on which the database is installed.
- In the case of an *online* backup, the database files are taken out of the backup mode after the replica is created.

In the case of an *offline* backup, the Oracle `alter database open` command is sent after the replica is created.

- Data Protector (for ZDB to disk) or MML (for ZDB to tape or ZDB to disk+tape) starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled (UNIX systems) and, unless the database is installed on raw disks, the mount points containing the Oracle database files are mounted.
- A ZDB agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.
- If a ZDB-to-disk session is being performed, at this point the remaining ZDB options are processed and details of the session are written to the ZDB database. The session then finishes. The following steps in this description are not performed; therefore, RMAN is not given any information about the ZDB-to-disk session.
- MML on the application system sends a request to the Data Protector **data movement agent** (DMA) on the backup system to back up the datafiles to tape.
- The DMA reads data from the backup system and sends it to the General Media Agent to write the actual data to the backup device.

DMA's role is also to disable the General Media Agent requests from accessing the application system. Thus, the database runs on the application system with greatly reduced performance degradation since the backup is performed on the backup system.

- At the end of data transfer, the backup system is disabled (filesystems are dismounted on all platforms and volume/disk groups deactivated on UNIX systems) and links are re-established.
- The recovery catalog and the control file are backed up automatically after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE:

A replica of the archive logs is not created; therefore, the archive logs should be backed up from the application system, following the standard Data Protector Oracle archive logs backup procedure.

Configuring the integration

Prerequisites

- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- You need a license to use the Data Protector ZDB integration with Oracle. Additional licenses are required for instant recovery and for the online extension. For information on licensing, see the *HPE Data Protector Installation Guide*.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector client systems. See the:
 - Latest support matrices at <https://softwaresupport.hpe.com/manuals> for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HPE Data Protector Installation Guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (EMC, P9000 XP Array, P6000 EVA Array, or NetApp Storage) 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.
- A Data Protector disk array integration (EMC, P9000 XP Array, P6000 EVA Array, or NetApp Storage) must be correctly installed and configured. For installation, see the *HPE Data Protector Installation Guide*. For configuration, see the *HPE 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:
 - If a disk array of the HPE P6000 EVA Disk Array Family is used, it must support multissnapping. For configuration details, see [Configuring the HPE P6000 EVA Disk Array Family in ASM environments, on page 36](#).
 - If a disk array of the HPE P9000 XP Disk Array Family is used, it must support the atomic split operation. For configuration details, see [Configuring the HPE P9000 XP Disk Array Family in ASM environments, on page 36](#).

For details about which disk array models and disk array firmware revisions support creation of replicas with cross-volume data consistency, see the latest support matrices at <https://softwaresupport.hpe.com/>.

- The Oracle Server software must be installed on the application system and the Oracle target database must be open or mounted there.
- From Oracle 12c onwards, the Oracle database on Microsoft Windows supports the use of an Oracle home user which is specified at the time of installation. This Oracle home user is used to run the Windows services for Oracle home, and is similar to the Oracle user on Oracle Database on Linux. For backup and restore in Oracle 12c database, if the Oracle integration agent and Media agent are running on the same Windows host, then to avoid shared memory allocation issues Oracle home user should be added to the Windows Backup Operator group.

For more information, see the [Oracle documentation](#).

- 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*.

For details on checking the prerequisites listed above, see [Troubleshooting](#), on page 108.

Note that 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.

- To successfully back up the recovery files residing in the flash recovery area, ensure that you have correctly configured the flash recovery area.
- **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 a NFS mounted disk, you must modify the archive log backup specification. See [Backup of archive logs on RAC cannot be performed](#), on page 119.

- **RAC:** Data Protector Oracle integration must be installed on all the computers (nodes) of Oracle RAC. This is necessary as RMAN connects using Oracle net to the currently active node, and any active node requires SBT library, which is part of the Data Protector Oracle integration.
- **RAC:** With Oracle version 11.2.0.2 and later, 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, from where the control file is backed up.
- On Windows systems, when using the Oracle backup set ZDB method, set the omninc option ZDB_SMISA_AUTOMOUNTING on the backup system to 2, in order to enable automatic volume mounting on the local system.

Limitations

- The MAXPIECESIZE RMAN parameter option is not supported because the restore of multiple backup pieces created during a backup is not possible using the Data Protector Oracle integration.
- In the Oracle Database 10g Release 2, for HP-UX systems, the host name length on which the Oracle database is installed is limited to 8 characters
- **Oracle OS authentication:**

Oracle OS authentication is not supported in the following scenarios:

1. For Oracle ZDB configuration.
 2. For Oracle recovery catalog configuration.
 3. For Oracle data guard database configuration.
- The Oracle recovery catalog database must be used as RMAN repository for backup and restore operations. ZDB using the Oracle control file are not supported. This is set when configuring the database. See [Configuring Oracle databases, on page 41](#).
 - The Oracle database identifier (DBID) must be a unique in a Data Protector cell. If you clone a database you must change the DBID.
 - Preview of zero downtime backup and instant recovery sessions is not available.
 - Using the Oracle proxy-copy ZDB method, individual tablespaces or datafiles cannot be backed up during a ZDB-to-disk or ZDB-to-disk+tape session (instant recovery enabled), only the whole database can be backed up.
 - The Oracle backup set ZDB method is supported on UNIX raw logical volumes only if these were created with LVM or VxVM.
 - When using the Oracle backup set ZDB method, you must reconfigure the Oracle integration if the initialization parameter file has been changed since the last configuration execution. See [Configuring Oracle databases, on page 41](#).
 - The single-host configuration (BC1, TF/1) is not supported for Oracle backup set ZDB sessions.
 - Object copying and object mirroring is not supported for ZDB to disk.
 - Recovery files residing in the **flash recovery area** cannot be backed up using ZDB.
 - **Oracle Server configurations with ASM:** The following limitations apply:
 - Zero downtime backup (ZDB) is only supported with the HPE P6000 EVA Disk Array Family and the HPE P9000 XP Disk Array Family.
 - Only the backup set ZDB method is supported.
 - With the HPE P6000 EVA Disk Array Family, only configurations consisting of a single disk array unit are supported.
 - With the HPE P6000 EVA Disk Array Family, only one instance of ASM should be configured on one storage volume (LUN). If multiple instances are sharing the same LUN, instant recovery will overwrite data from all ASM instances.
 - Instant recovery is only supported with the HPE P6000 EVA Disk Array Family and the HPE P9000 XP Disk Array Family.
 - **Oracle Data Guard:** Standby database is not supported for ZDB.
 - The Data Protector Oracle integration does not support non-ASCII characters in backup specification names.
 - The Oracle integration using Oracle 12c will fail if there are any special characters not supported by Oracle or a space in the password. Out of the special characters supported by Oracle some do not work due to an Oracle bug: a space does not work and neither do the three combinations (a semicolon and a caret, a semicolon and an "at sign", a semicolon and a slash).

Here are the two factors while using special characters in the password:

1. You can only use those special characters which Oracle supports.
2. You should not use any space or any of the three combinations of passwords listed above.

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. For Oracle version 12c the user must have been granted the *SYSBACKUP* privilege. Also, you can use the user with *SYSDBA* privilege for Oracle 12c but first you must set omnirc variable *OB2_ORACLE_USE_SYSDBA* to 1.

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.

For supported Windows operating systems, use user impersonation. For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: “Inet user impersonation”.

- When using 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 look like:

```
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:
Execute:

```
alter system set log_archive_dest=path_to_archive_Logs SCOPE=SPFILE;
```
 - If the *init.ora* file is used:
Execute:

```
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 and 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 backup set method:

- 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. Check also 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 look like:

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

export DB_NAME=PROD
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 look like:

```
su - ora
id
uid=101(ora) gid=101(dba)
export DB_NAME=PROD

rman target system/manager catalog rman/rman
Recovery Manager: Release 10.1.0.2.0 - Production
RMAN-06005: connected to target database: PROD
```

```
RMAN-06008: connected to recovery catalog database
RMAN> exit
Recovery Manager completed.
```

Configuring the HPE P6000 EVA Disk Array Family in ASM environments

Zero downtime backup and instant recovery with the HPE P6000 EVA Disk Array Family are supported for ASM configurations, provided that the P6000 EVA Array supports multisnapping. Additionally, the following prerequisites must be fulfilled:

- To enable zero downtime backup, the ASM-managed files that will be backed up must reside on raw disks, not on raw logical volumes.

Note that the maximum number of source disks that can be involved in multisnapping depends on the firmware revision of the P6000 EVA Array that will be used and the installed Command View (CV) version. If the number of source disks selected for a zero downtime backup session exceeds this limitation, the session is aborted. For limitation details, see the HPE P6000 EVA Disk Array Family documentation.

- The autoextend feature of the Oracle Server ASM must be disabled.
- To enable instant recovery of the Oracle ASM-managed data, set the omnirc variable `SMISA_ALLOW_ASM_IR` to 1. For details, see the HPE Data Protector Help index: "omnirc options".

Configuring the HPE P9000 XP Disk Array Family in ASM environments

In Oracle Server configurations that use Automatic Storage Management (ASM), both ZDB and instant recovery are supported with the HPE P9000 XP Disk Array Family, provided that the following prerequisites are met:

- Data files, control files, and redo log files must reside on separate storage volumes (LUNs). This configuration is needed if you plan to perform instant recovery.
- The disk arrays must support the atomic split operation:
 - Each storage volume (LUN) must be composed of a single LDEV.
 - The storage volumes on which the ASM-managed files reside must belong to consistency groups (CTG) that have unique and non-zero IDs. Storage volumes that belong to the same CTG must be selected together in a backup specification. Otherwise, the backup session will fail. If you plan to back up datafiles, control files and log files in separate sessions, the corresponding storage volumes must reside in separate consistency groups.

Depending on your environment, set the following `omnirc` options on the application system:

- `SSEA_ATOMIC_SPLIT`
- `SSEA_ATOMIC_SPLIT_MULTIPLE_CTGROUPS`
- `SSEA_ATOMIC_SPLIT_MIXED_CONFIG`

At least one option must be set to enable the atomic split operation. For details, see the online Help index: "omnirc options".

- If the ASM instance name differs from +ASM, specify the correct name by setting the Data Protector `omniirc` option `ORA_ASM_LCL_INSTANCE` on the application system.
- Additional requirements for instant recovery:
 - In a cluster environment, on the active cluster node, set the Data Protector `omniirc` option `ORA_ASM_LCL_INSTANCE` to the name of the ASM instance running on that node.
 - If an Oracle ASM instance manages files of more than one database, you must reconfigure Oracle Server to use a separate ASM disk group for each database.
 - The names of the ASM disk groups should not be changed after the backup.
 - The autoextend feature of the Oracle Server ASM must be disabled.

Cluster-aware systems

In cluster environment, if you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name. Set the variable on the Oracle Server system as follows:

Windows systems: `set OB2BARHOSTNAME=virtual_server_name`

UNIX systems: `export OB2BARHOSTNAME=virtual_server_name`

HP-UX with RAC: To enable instant recovery, create an HPE Serviceguard package containing *only* the virtual IP and the virtual hostname parameters and distribute it among the RAC nodes.

Linking Oracle Server with the Data Protector MML

To use the Data Protector Oracle integration, the Oracle Server software needs to be linked with the Data Protector Oracle integration **Media Management Library (MML)** on every system on which an Oracle instance is running.

You do not need to link Oracle Server with the Data Protector MML manually. When you start backups or restores using the Data Protector GUI or CLI, Data Protector automatically links Oracle Server with the correct platform-specific Data Protector MML. However, for testing purposes, you can override this automatic selection. You can manually specify which platform-specific Data Protector MML should be used by setting the Data Protector `SBT_LIBRARY` parameter. On how to set the parameter, see the `util_cmd` man page. The parameter is saved in the Data Protector Oracle instance configuration file.

MML is invoked by the Oracle server when it needs to write to or read from devices using Data Protector.

Oracle 12c CDB and PDB mode support

To backup and restore data, a new feature is introduced in Oracle 12c called Pluggable Database (PDB). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appears to an Oracle Net client as a non-container database (non-CDB). A CDB includes zero, one, or many customer created pluggable databases (PDBs). The PDB is located under container database (CDB) in Oracle 12c.

Each CDB has the following containers:

- Exactly one **root**
The root stores Oracle-supplied metadata and common users. An example of metadata is the source code for Oracle-supplied PL/SQL packages. A common user is a database user known in every container. The root container is named CDB\$ROOT.
- Exactly one **seed PDB**
The seed PDB is a system-supplied template that is used by CDB to create new PDBs. The seed PDB is named PDB\$SEED. You cannot add or modify objects in PDB\$SEED.
- **Zero or more user-created PDBs**
A PDB is a user-created entity that contains the data and code required for a specific set of features. For example, a PDB can support a specific application, such as a human resources or sales application.

Backup Archive and Restore (BAR) GUI lists all the PDBs in the container database, selection of only one or more pluggable databases should be enabled.

Following are the restore scenarios:

- Restore both CDB and PDB
- Restore CDB and no PDB
- Restore one PDB and a test database

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
For details, see [Configuring Oracle operating system user accounts](#) , below.
- Oracle database user accounts
For details, see [Configuring Oracle database user accounts](#), on page 40.

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, execute:

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

or

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

Finding the Oracle user



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

Client system	Description
UNIX system	<p>Ensure that the Oracle user <code>oracle</code> from the Oracle Inventory group (<code>oinstall</code>) has been added to the Data Protector <code>admin</code> user group. For details on adding users, see the <i>HPE Data Protector Help</i> index: "adding users".</p> <p>Add the OSDBA user account and <code>root</code> user account from both the application system and backup system to the Data Protector <code>admin</code> or <code>operator</code> 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, <code>uid=101(ora) gid=101(dba)</code>).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>TIP: To find the user ID, connect to a system under this user account and execute:</p> <pre>#id</pre> </div>
Windows system	<p>On Windows systems, Data Protector connects to the Oracle database using the Data Protector <code>Inet</code> service on the related system. By default, the service runs under the <code>Local System</code> account, which is automatically added to the Data Protector <code>admin</code> user group. However, if you have restarted the Data Protector <code>Inet</code> service on the application system and backup system under OSDBA user accounts, you need to add the new users to the Data Protector <code>admin</code> or <code>operator</code> user group.</p>

For information on adding users to Data Protector user groups, see the *HPE Data Protector Help* index: "adding users".

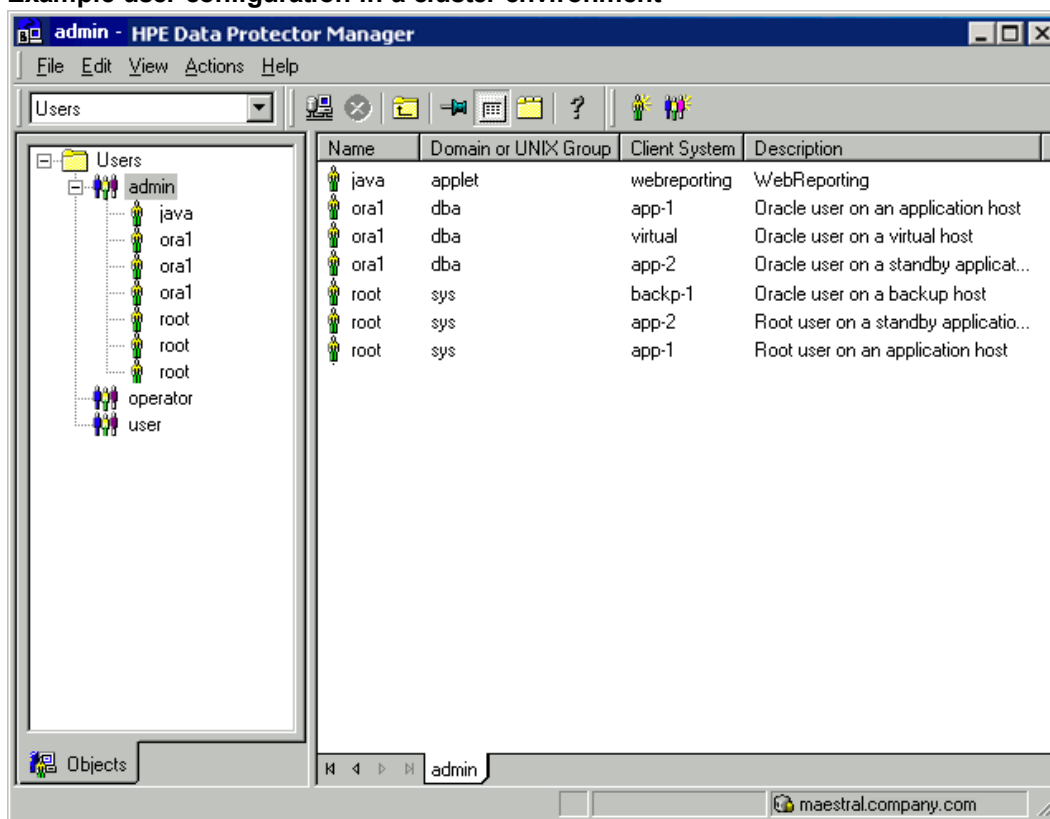
NOTE:
 The OSDBA user account for the backup system needs to be added to a Data Protector user group only if you plan to use the Oracle backup set ZDB method.

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 HPE Serviceguard clusters)
- **UNIX systems:** `root` user for all physical nodes

Example user configuration in a cluster environment



Configuring Oracle database user accounts

Identify or create the following Oracle database user accounts. You need to provide these user accounts when you configure the Oracle database as described in [Configuring Oracle databases, on the next page](#).

Oracle database user accounts

User	Description
Primary database user	Required to log in to the primary database.
Recovery catalog user	<p>The owner of the recovery catalog (for example, rman). Required to log in to the catalog database. Needed if you use the recovery catalog.</p> <p>If you are using Oracle 11g R2 or later, ensure that the owner of the Oracle recovery catalog:</p> <ul style="list-style-type: none"> 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:

User	Description
	<pre>grant select on v_\$instance to recovery_catalog_user;</pre> <p>If you are using Oracle 12c or later, ensure that the owner of the Oracle recovery catalog:</p> <ul style="list-style-type: none">• is granted the EXEMPT ACCESS POLICY permissions
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.

Configuring Oracle databases

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.

NOTE:

With Oracle Server configurations using ASM, to be able to perform instant recovery, the Oracle database must be configured using the Data Protector CLI. This is because the ASM-related parameters cannot be set using the Data Protector GUI. However, if you plan to perform ZDB sessions only, without instant recovery, the database can also be configured using the Data Protector GUI.

Oracle OS authentication

To connect to the database, Data Protector can use Oracle listeners and Oracle OS authentication. For Oracle OS authentication, during database backup user must provide username and usergroup for Oracle OS user which can be used as Oracle OS database administrator. This user must belong to Oracle **dba/sysdba** group. The user information is entered during creation of backup specification see [Creating backup specifications](#) , on page 54

To configure backup specification using OS authentication, perform the following steps:

1. On the Oracle host, create the OS user that will be used for Oracle OS authentication. The user must be in Oracle OS sysdba/dba group. For details how to create Oracle OS database administrator refer to Oracle documentation.
2. Add the user in Data Protector admin user group.
 - a. For host with Oracle agent and Windows platform add impersonation for this user.
3. Create backup specification.
 - a. On the **Specify the application that you want to backup** screen, select **Specify OS user** check box and enter the user name and group for created user.
 - b. On the **Configure Oracle** dialog box in primary tab select **Use OS Authentication** check box.

NOTE:

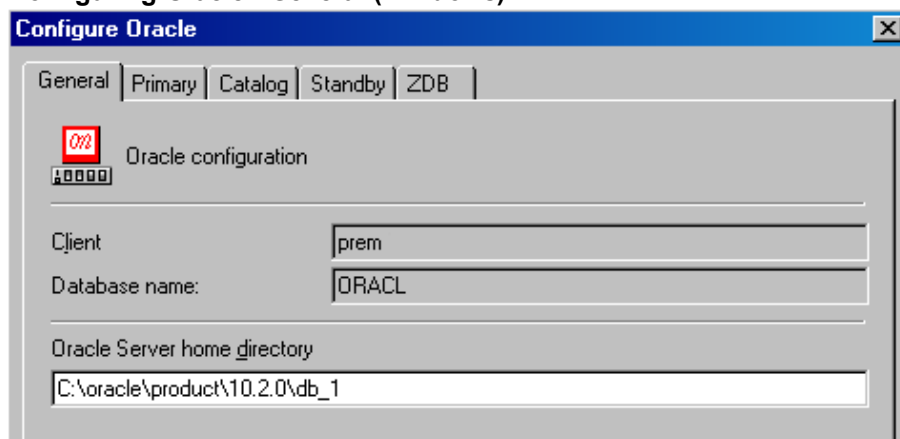
For Oracle RAC environments in the **RAC Database Name** text box enter the global database name.

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 [Creating backup specifications](#) , on page 54 and at [If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed. Configure the Oracle database for use with Data Protector as described in Configuring Oracle databases, on page 41., on page 64](#) proceed as follows:

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

Configuring Oracle - General (Windows)



Configuring Oracle - General (UNIX)



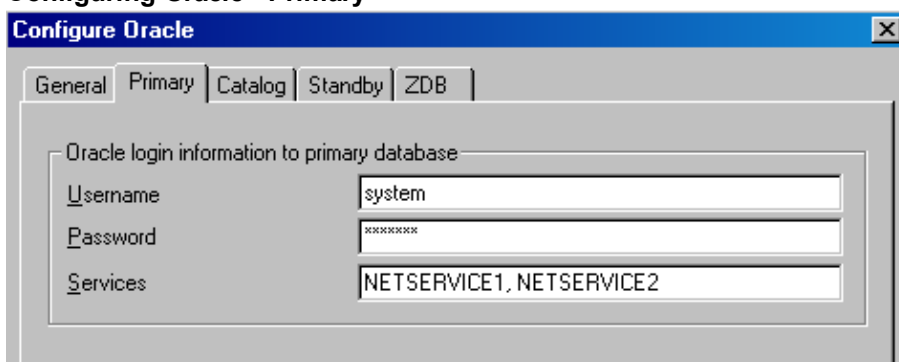
2. In the **Primary** page, specify the login information to the primary database.

Note that the user must have the SYSDBA privilege granted. For Oracle version 12c the user must have been granted the SYSBACKUP privilege. Also you can use the user with SYSDBA privilege for Oracle 12c but first you must set omnirc variable OB2_ORACLE_USE_SYSDBA to 1.


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.

Configuring Oracle - Primary



Oracle OS authentication



In the primary tab select **Use OS Authentication** checkbox.

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

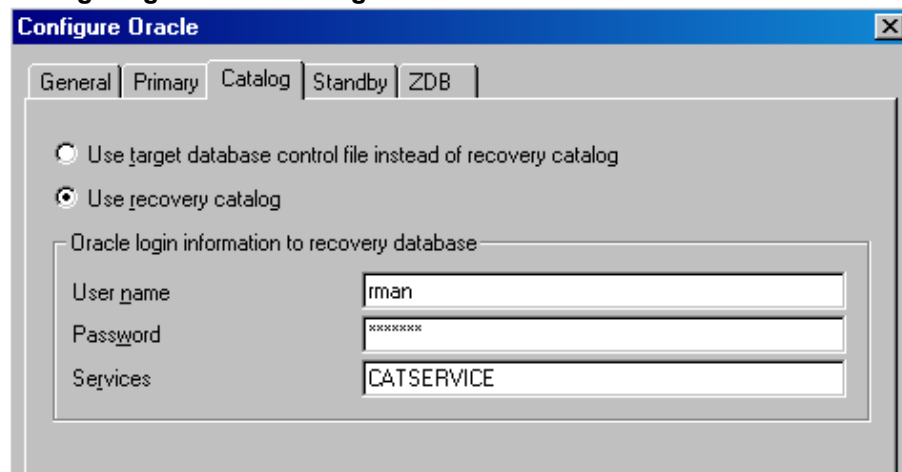
To use the recovery database catalog as an RMAN repository for backup history, select **Use recovery catalog** and specify the login information to the recovery catalog.

Note that for ZDB, you must use the recovery catalog.

The user specified must be the owner of the recovery catalog.

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

Configuring Oracle - Catalog




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

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.

Configuring Oracle - Standby



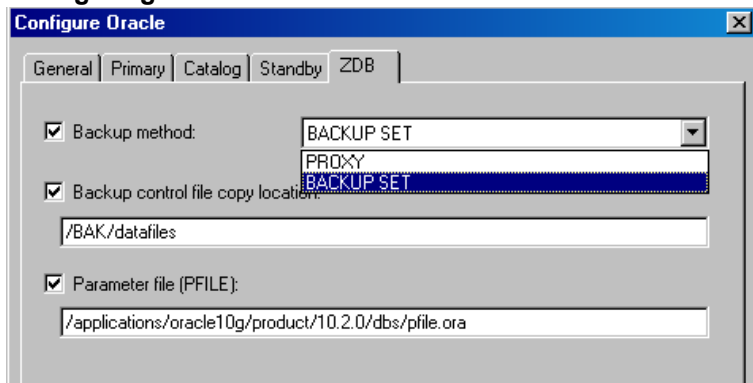
5. In the **ZDB** page, select **Backup method** and then select **PROXY** or **BACKUP SET** in the drop-down list.

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 ZDB to disk.

If you do not specify the location, `ob2rman.p1` 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.

Configuring Oracle - ZDB



6. Click **OK**.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at [Select the Oracle database objects to be backed up.](#), on page 64.

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 | USEOSAUTHENTICATION [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 |  
USEOSAUTHENTICATION [CATALOG_DB_LOGIN] [STANDBY_DB_LOGIN] [ZDB_OPTIONS] [ASM_  
OPTIONS] [-client CLIENT_NAME]
```

where:

PRIMARY_DB_LOGIN

-prouser PRIMARY_USERNAME

-prmpasswd PRIMARY_PASSWORD

USEOSAUTHENTICATION

-useosauth USE_OS_AUT

-racdbname RAC_DB_NAME

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 -asm service 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

<i>CLIENT_NAME</i>	<p>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.</p> <p>RAC: The virtual server of the Oracle resource group.</p> <p>Oracle Data Guard: Name of either a primary system or secondary (standby) system.</p>
<i>DB_NAME</i>	Name of the database to be configured.
<i>ORACLE_HOME</i>	Pathname of the Oracle Server home directory.
<i>PRIMARY_USERNAME</i> <i>PRIMARY_PASSWORD</i>	Username and password for login to the target or primary database. Note that the user must have the SYSDBA privilege granted. For Oracle version 12c the user must have been granted the SYSBACKUP privilege. Also you can use the user with SYSDBA privilege for Oracle 12c but first you must set omnirc variable OB2_ORACLE_USE_SYSDBA to 1.
<i>PRIMARY_NET_SERVICE_NAME_1</i> [, <i>PRIMARY_NET_SERVICE_NAME_2</i> , ...]	<p>Net services names for the primary database.</p> <p>RAC: Each net service name must resolve into a specific database instance.</p>
<i>USE_OS_AUT</i>	<p>This parameter is used to communicate to the configuration script to use Oracle OS authentication instead of username and pass authentication.</p> <p>Value should be 1.</p>
<i>RAC_DB_NAME</i>	This parameter is used in Oracle RAC environments in order to find Oracle SID for current Oracle RAC node. This parameter represent Oracle global database name and this is mandatory on Oracle RAC platforms.
<i>CATALOG_USERNAME</i> <i>CATALOG_PASSWORD</i>	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.
<i>CATALOG_NET_SERVICE_NAME</i>	Net service name for the recovery catalog.
<i>STANDBY_USERNAME</i> <i>STANDBY_PASSWORD</i>	This is used in Oracle Data Guard environment for backing up a standby database. Username and password for login to the standby database.
<i>STANDBY_NET_SERVICE_NAME_1</i> [, <i>STANDBY_</i>	Net services names for the standby database.

<i>NET_SERVICE_NAME_2, ...]</i>	
<i>BACKUP_CONTROL_FILE_COPY_LOCATION</i>	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.p1 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.
<i>PARAMETER_FILE</i>	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).
<i>BACKUP_SYSTEM</i>	Name of the backup system. It must be specified for a ZDB backup set configuration.
<i>ASM_HOME</i>	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.
<i>ASM_USERNAME</i> <i>ASM_PASSWORD</i>	User name and password (authentication credentials) used by the Data Protector Oracle integration agent to connect to the ASM database.
<i>ASM_NET_SERVICE_NAME_1</i> [, <i>ASM_NET_SERVICE_NAME_2</i> ...]	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 *RETVL*0 indicates successful configuration, even if followed by additional messages.

NOTE:
 If you need to export some variables before starting SQL*Plus, listener, or RMAN, 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/sbin/util_oracle8.pl -config -dbname oracle -orahome  
/app10g/oracle10g/product/10.1.0 -prouser system -prpasswd 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/sbin/util_oracle8.pl -config -dbname SUN -orahome  
/orahome/ora/app/oracle/product/11.2.0/dbhome_1 -prouser sys -prpasswd 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 -asmervice ASMSRV1,  
ASMSRV2
```

Checking the 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 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 either:

- Reconfigure the Oracle database instance. See [Reconfiguring an Oracle instance for instant recovery, on page 406](#) on how to move the control files and redo logs to source volumes that are not replicated.
- Or:
- Set the ZDB_ORA_INCLUDE_CF_OLF, ZDB_ORA_INCLUDE_SPF, and ZDB_ORA_NO_CHECKCONF_IRomnirc options and ignore the warning. However, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery. See [ZDB integrations omnirc options, on page 409](#) on how to set the omnirc options.

Setting environment variables

Use environment variables to modify backup environment to suit your needs. Environment variables are Oracle database specific. 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 of how environment variables affect your environment, see [Environment variables , below](#).

NOTE:

Environment variables are not supported on HP OpenVMS systems.

Environment variables

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 its temporary files directory. This variable overrides the default

Environment variable	Default value	Description
		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.

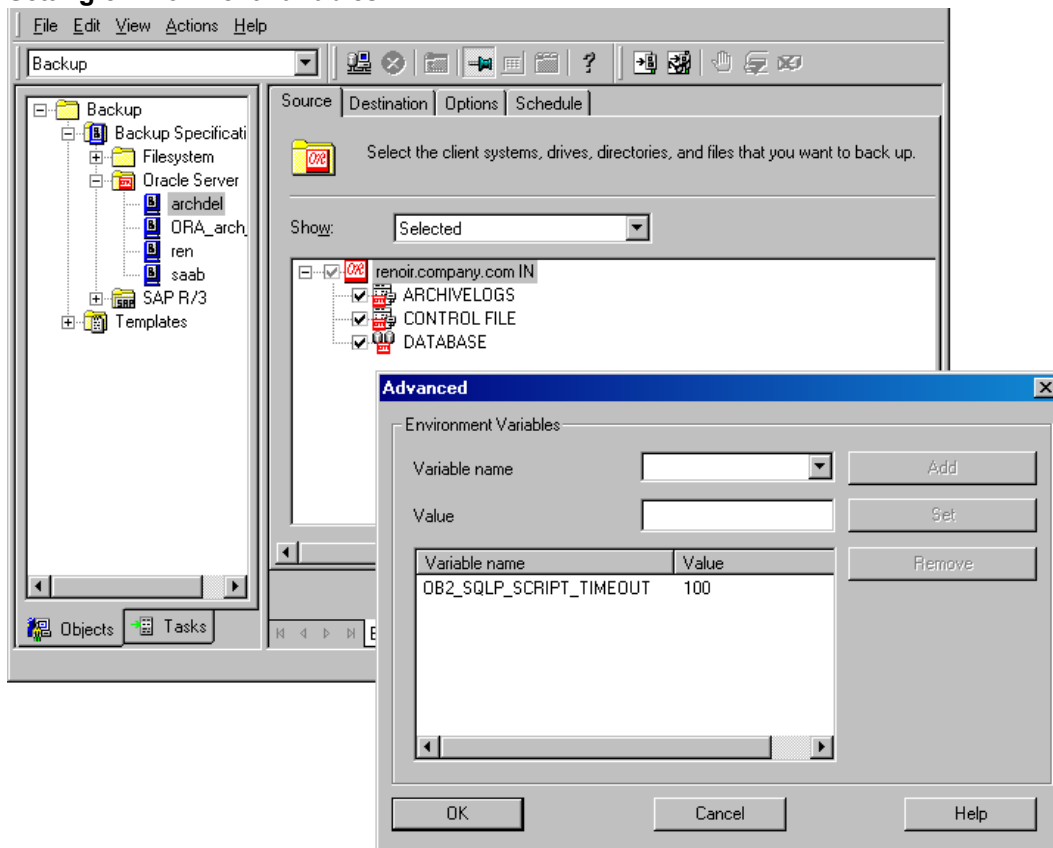
To set environment variables, use the Data Protector GUI or CLI.

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**.
 See [Setting environment variables, below](#).

Setting environment variables



3. Click **OK**.

Using the Data Protector CLI

Execute:

```
util_cmd -putopt Oracle8 DatabaseNameVariableValue -sublist Environment
```

For details, see the `util_cmd` man page or the *HPE 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
```

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 below 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, you may 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_program_data\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_program_data \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 old method through a channel allocated for the files from the old method and the restore will fail. See [Troubleshooting , on page 108](#) on how to restore such a backup.

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 *HPE Data Protector Help* index: “configuring devices”.
2. Configure media pools and media for a backup. For instructions, see the *HPE Data Protector Help* index: “creating media pools”.
3. Ensure 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. On how to create a non-ZDB backup specification, see the *HPE Data Protector Integration Guide for Oracle and SAP*.
5. Create a Data Protector Oracle ZDB backup specification.
See [Creating backup specifications , below](#).

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 only a ZDB backup specification.

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 `fs cmd` 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:

Backup set method

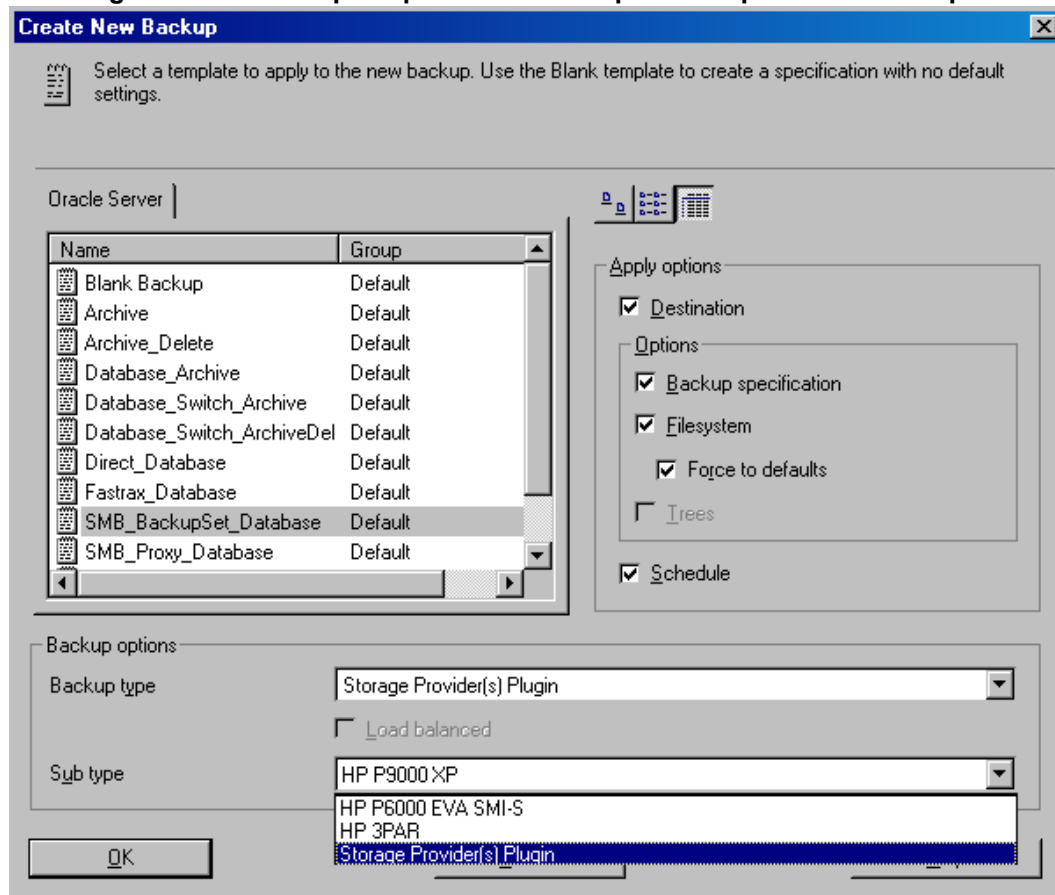
To perform a ZDB of the entire database using the backup set method, select the **SMB_BackupSet_Database** template.

Proxy-copy method

To perform a ZDB of the entire database using the proxy-copy method, select the **SMB_Proxy_Database** template.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system. See [Selecting an Oracle backup template and the snapshot or split mirror backup, below](#).

Selecting an Oracle backup template and the snapshot or split mirror backup



Click **OK**.

4. 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.

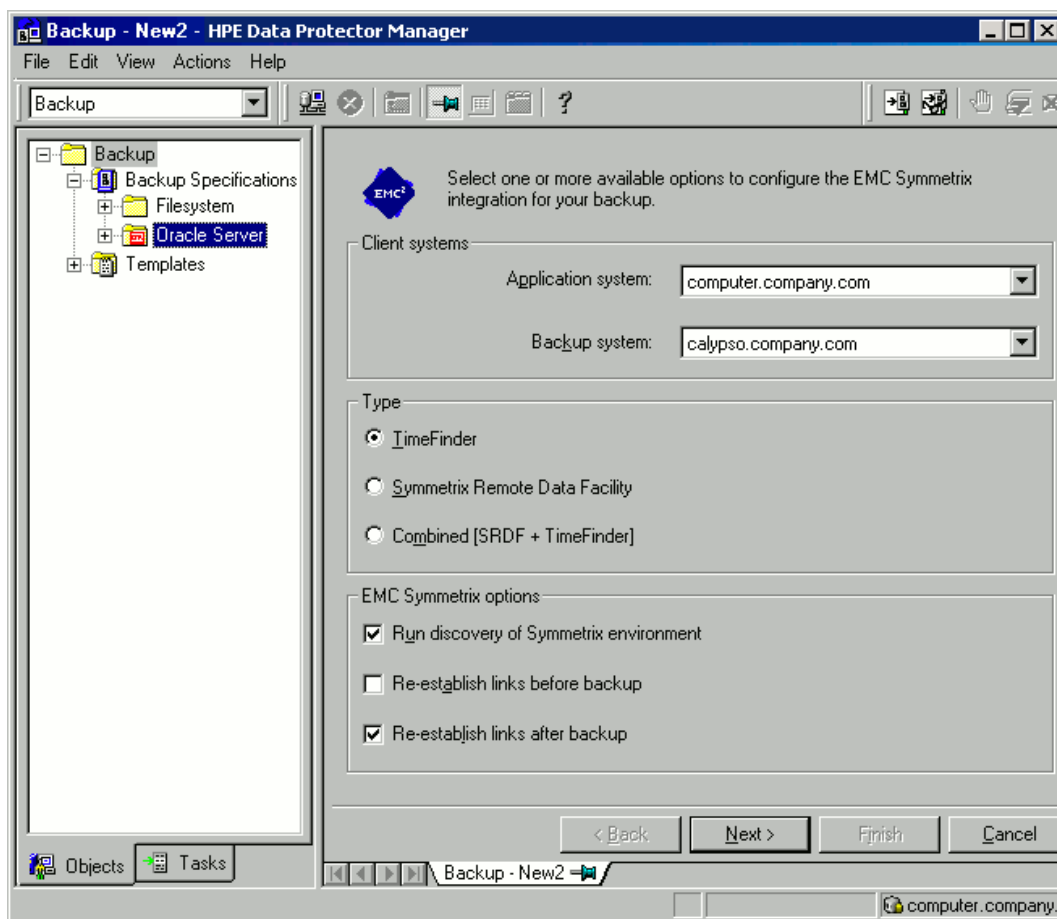
Select other disk array-specific backup options (see [Backup](#), on page 54 for EMC, [P9000 XP Array backup options - tbd, on the next page](#) for P9000 XP Array, [P6000 EVA Array backup options, on page 58](#) for P6000 EVA Array, [NetApp Storage backup options, on page 60](#) for NetApp Storage, [EMC VNX Storage backup options, on page 61](#) for EMC VNX Storage, or [EMC VMAX Storage backup options, on page 62](#) for EMC VMAX Storage. For detailed information on the backup options, press **F1**.

EMC GeoSpan specifics

In the EMC GeoSpan for Microsoft Cluster Service environment, select the backup system for the active node and specify the TimeFinder configuration.

After a failover in EMC GeoSpan for MSCS, select the backup system for the currently active node and save the backup specification.

EMC backup options

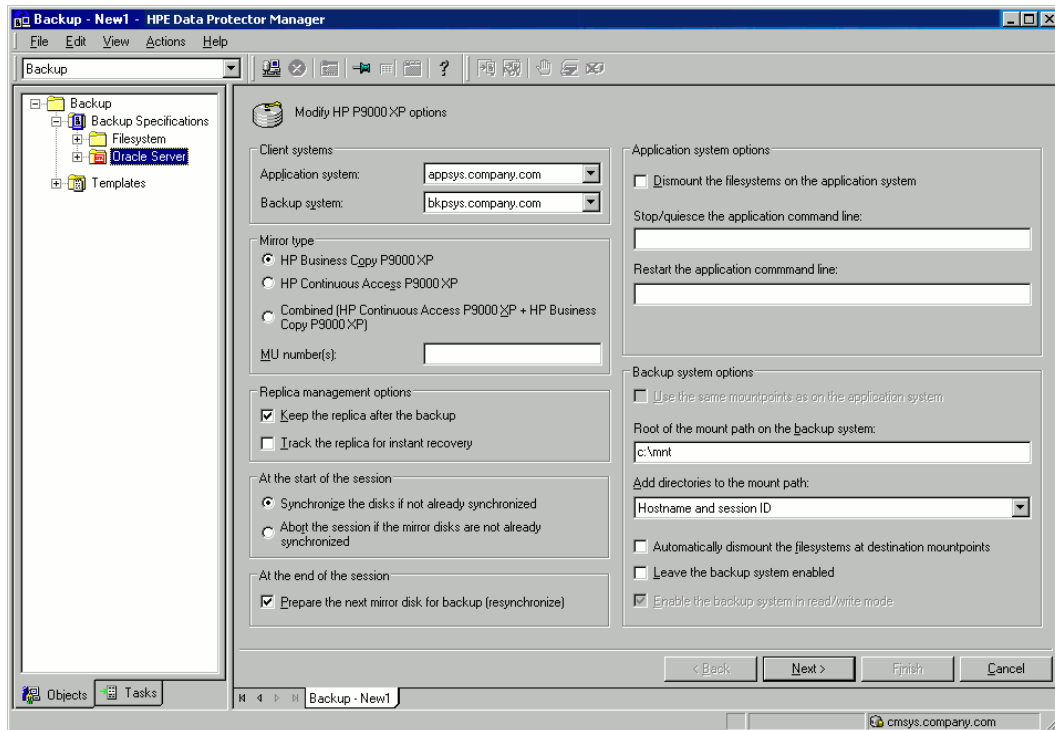


P9000 XP Array specifics

To enable instant recovery, leave the **Track the replica for instant recovery** option selected. It is not possible to run instant recovery with Data Protector if this option is cleared.

NOTE: It is recommended to select the option **Use the same mountpoints as on the application system**, during a ZDB+IR+ORACLE backup using a P9000 XP Array On Linux.


P9000 XP Array backup options - tbd



P6000 EVA Array specifics

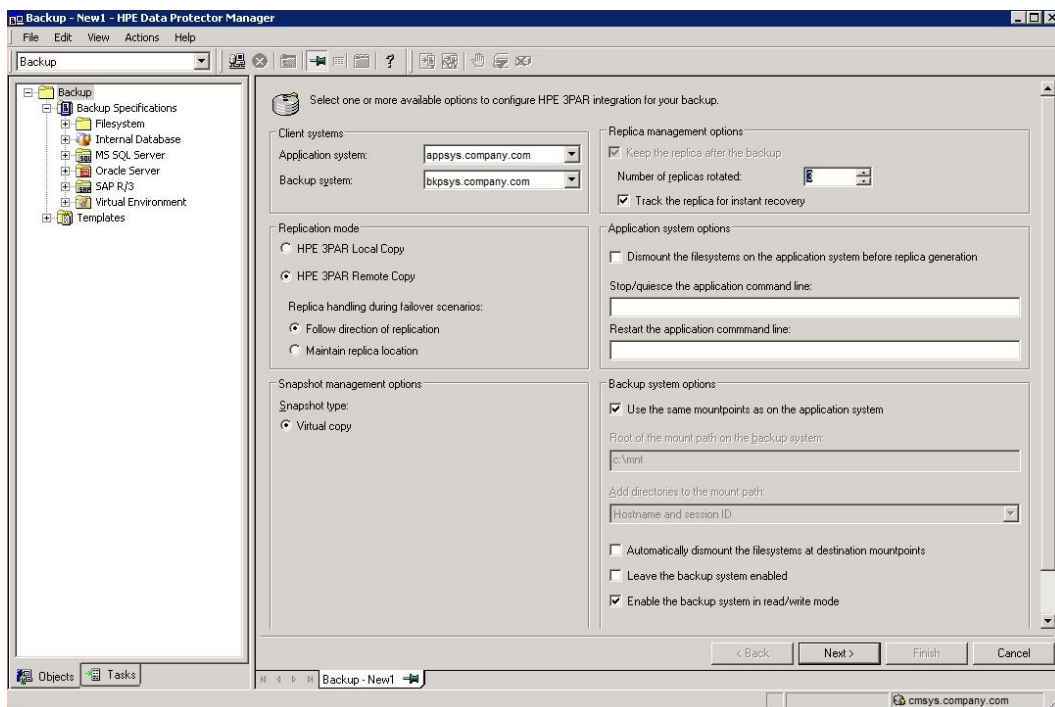
To enable instant recovery, select the **Track the replica for instant recovery** option.

P6000 EVA Array backup options




P10000 3PAR Array specifics

P10000 3PAR Array backup options




NetApp Storage specifics

NetApp Storage backup options




EMC VNX Storage specifics

EMC VNX Storage backup options



EMC VMAX Storage:

EMC VMAX Storage backup options



Click **Next**.

5. In **Application database**, type the name of the database to be backed up.

The database name can be obtained using SQL*Plus:

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

NOTE:

In a single-instance configuration, the database name is usually the same as its instance name. In this case, the instance name can be also used. The instance name can be obtained as follows:

```
SQL>select instance_name from v$instance;
```

Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 systems, as follows:

- **UNIX systems:** 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). This user must be configured as described in [Configuring Oracle user accounts, on page 38](#).
- **Windows Server 2008 systems:** It is not mandatory to specify these options and if they are not specified, the backup runs under the Local System Account.


In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP). This user must be set up for the Data Protector Inet service user impersonation.

For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: "Inet user impersonation".


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

NOTE:
If this is not your first backup specification, Data Protector fills in **Username** and **Group/Domain name** for you, providing the values of the last configured Oracle database.

Specifying an Oracle Server system (Windows)



Specifying an Oracle Server system (UNIX)



Click **Next**.

NOTE:

When you click **Next**, Data Protector performs a configuration check.

UNIX systems: The check is started under the specified OSDBA user account. If it completes successfully, the OSDBA user and group are also saved in both the Oracle database specific configuration file and Oracle system global configuration file, overriding previous values if they exist.

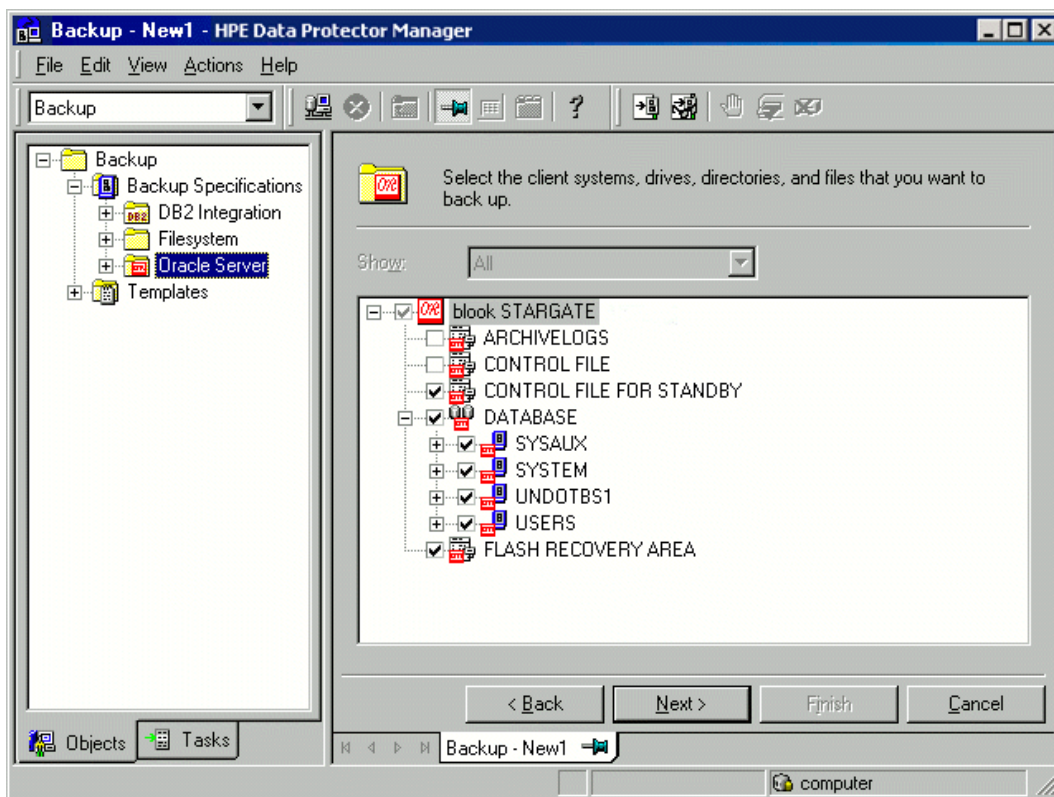
You must enter username and group name for the user for which you want to use the OS Authentication.

6. If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed. Configure the Oracle database for use with Data Protector as described in [Configuring Oracle databases, on page 41](#).
7. 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.

Selecting backup objects



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.

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

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

For detailed information on the object mirror functionality, see the *HPE Data Protector Help* index: “object mirroring”.

NOTE:

Object mirroring is not supported for ZDB to disk.

Click **Next** to proceed.

9. Set the backup options.


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

Offline ZDB

To perform an offline ZDB, select the **Backup offline** option in the Application Specific Options dialog box. This option stops the database before creating a replica, and restarts it after the replica

is created. Note that if a ZDB-to-tape or ZDB-to-disk+tape session is being performed, the database is not offline during the actual backup to tape. See [Backup offline option, below](#).

Backup offline option



For information on other Application Specific Options, see [Oracle backup options](#), on the next page or press **F1**.

Click **Next**.

10. Click **Save As** to save the backup specification, specify a name and a backup specification group. It is recommended that you save all Oracle backup specifications in the **Oracle** group. Optionally, you can click **Save and Schedule** to save, and then schedule the backup specification. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

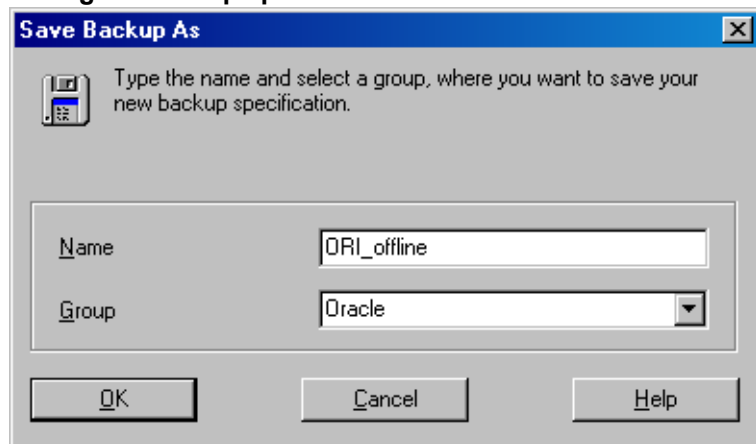
Note that only the **Full** backup type is supported.

IMPORTANT:

The word **DEFAULT** is a reserved word and therefore must not be used for backup specification names or labels of any kind. Therefore, do not use a punctuation in the names

of backup specifications, since the Oracle channel format is created from the backup specification name.

Saving the backup specification



Click **OK**.

To start the backup, see [Starting backup sessions , on page 72](#).

11. For **online backup**, also create a standard Data Protector Oracle integration backup specification for backing up the application system archived log files. See the *HPE Data Protector Integration Guide*.

TIP:
 The backup specification for the backup of archived log files can be either triggered by the **Post-Exec** command defined in the ZDB backup specification for the backup of database files (recommended), or started manually after the ZDB backup specification has been started. See the *HPE Data Protector Help* index: “pre- and post-exec commands” for more information on configuring the **Pre-Exec** and **Post-Exec** commands.

Oracle backup options

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 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 how to edit the RMAN script section, see Editing the Oracle RMAN script , on the next page.
Pre-exec, Post-exec	<p>Specify a command or RMAN script that will be started by <code>ob2rman.pl</code> on the Oracle Server system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). RMAN scripts must have the <code>.rman</code> extension. Do not use double quotes.</p> <p>For example, you can provide scripts to shut down and start up an Oracle instance. For examples of shut-downing and starting an Oracle instance on a UNIX system, see Examples of pre-exec and post-exec scripts on UNIX systems , below.</p> <p>Provide the pathname of the command or RMAN script.</p>
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. See Backup offline option, on page 66 .

Examples of pre-exec and post-exec scripts on UNIX systems

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
```

Editing the Oracle RMAN script

The RMAN script is used when the Data Protector backup specification is started to perform a backup of the Oracle objects.

The RMAN script section is not written to the backup specification until the backup specification is either saved or manually edited by clicking the **Edit** button.

You can edit the RMAN script section of only after the Data Protector Oracle backup specification has been saved.

Limitations

When editing the RMAN script sections of the Data Protector backup specifications, consider the following limitations:

- The Oracle manual configuration convention must be used and not the Oracle automatic configuration convention.
- Double quotes (") must not be used - single quotes should be used instead.
- By default, RMAN scripts created by Data Protector contain instructions for backing up one or more of the following objects:
 - Databases, tablespaces, or datafiles (the first backup command)
 - Archive logs (the second backup command)
 - Control files (the last backup command)

The RMAN scripts with all combinations of the above listed backup objects are recognized by Data Protector as its own scripts and it is possible to modify the selection of objects that will be backed up in the **Source** tab of the Results Area.

If the RMAN script contains *additional* manually entered backup commands, for example a second backup command for backing up a database that is already listed in the first backup command, the object selection is disabled and it is only possible to browse the **Source** tab.

To edit an Oracle RMAN script, click **Edit** in the **Application Specific Options** window (see [Recovery catalog settings dialog, on page 78](#)), edit the script, and then click **Save** to save the changes to the script.

See the *Oracle Recovery Manager User's Guide and References* for more information on Oracle RMAN commands.

Data Protector RMAN script structure

The RMAN script created by Data Protector consists of the following parts:

- **The Oracle channel allocation** together with the Oracle environment parameters' definition for every allocated channel.

For all backup specifications except for Oracle proxy-copy ZDB backup specifications, the number of allocated channels is the same as the sum of concurrency numbers for all devices selected for backup.

NOTE:

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

IMPORTANT:

On Windows systems, a maximum of 32 or 64 (if device is local) channels can be allocated. If the calculated number exceeds this limitation, you have to manually edit the RMAN script and reduce the number of allocated channels.

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

```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME, OB2BARLIST=Backup_
Specification_Name)';
```

Proxy-copy

For Oracle proxy-copy ZDB backup sessions, Data Protector allocates *one* channel.

For Oracle proxy-copy ZDB, the OB2SMB parameter must be set to 1. If you use the Blank Oracle Backup template, the number of concurrently running DMA (OB2DMAP) is automatically calculated as the sum of all device concurrences; for example, if there are 4 devices with concurrency set to 3 then OB2DMAP will be set to 12.

If you use the Oracle_SMB template, the OB2DMAP parameter is set to 1. To improve the backup and restore performance, you may want to increase the value of this parameter. The environment parameters must be added in the following format:

```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME, OB2BARLIST=Backup_
Specification_Name, OB2SMB=1, OB2DMAP=Concurrent_DMAs)';
```

NOTE:

The OB2DMAP parameter does not change after it has been calculated, even if you adjust the device concurrency. To change OB2DMAP, you have to manually edit the RMAN script.

- Depending on the backup objects selection, **an RMAN backup statement for the backup of the whole database instance, and/or for any combination of RMAN commands to back up tablespaces and datafile**. The backup statement consists of the following:
 - The Oracle format of the backup file in the following format:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf' database;
```

NOTE:

When an Oracle format of the backup file is manually defined or changed by editing the

RMAN script, any user-defined combination of the Oracle substitution variables can be *added* to the %s:%t:%p substitution variables and *DB_NAME*, which are obligatory.

- In case of an Oracle proxy-copy ZDB-to-disk+tape or ZDB-to-tape session, the `PROXY ONLY` option is required. Only one `BACKUP` command with the `proxy only` option is permitted and only one additional backup command for backing up the control file is permitted.
- The RMAN datafile `tablespace_name*datafile_name` command.
- If the archived redo logs were selected for a backup, **an RMAN backup statement for the backup of Oracle archive logs.**

The backup statement consists of the Oracle format of the backup file:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf'
```

NOTE:

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

- If the control file was selected for a backup, **an RMAN backup statement for the backup of Oracle control files.** The backup statement consists of the following:

- The Oracle format of the backup file in the following format:

```
format 'Backup_Specification_Name<DB_NAME_%s:%t:%p>.dbf' current controlfile;
```

NOTE:

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

- The RMAN `current controlfile` command.

For Oracle proxy-copy ZDB to disk or disk+tape, it is not possible to select only the control file. You must also select either a `DATABASE`, `TABLESPACE`, or `DATAFILE` object.

Example of the Oracle proxy-copy ZDB to disk+tape RMAN script

The following is an example of the RMAN script section as created by Data Protector based on the Oracle `SMB_Proxy_Database` template, after the whole database selection:

```
run {  
  allocate channel 'dev_0' type 'sbt_tape' parms  
    'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1,  
    OB2SMB=1,OB2DMP=1)';  
  backup incremental level <incr_level>format 'New1<DIPSI_%s:%t:%p>.dbf'  
  proxy only  
  database  
  ;  
  backup format 'New1<DIPSI_%s:%t:%p>.dbf' controlfile;  
}
```

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:

Backup methods

- Schedule a backup of an existing Oracle ZDB backup specification using the Data Protector Scheduler. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.
- Start an interactive backup of an existing Oracle ZDB backup specification using the Data Protector GUI or the Data Protector CLI. See [Running an interactive backup](#), on the next page.

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.
- On P9000 XP Array, if the LVM mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their HPE BC P9000 XP pairs assigned. This message should be ignored.
- If the control file, SPFILE, or online redo logs are on the same source volumes as the datafiles and the **Track the replica for instant recovery** option is selected, the backup session will be aborted. In this case, you need to either reconfigure the database or set the ZDB_ORA_INCLUDE_CF_OLF, ZDB_ORA_INCLUDE_SPF, and ZDB_ORA_NO_CHECKCONF_IRomnirc options. See [Reconfiguring an Oracle instance for instant recovery](#), on page 406 or [ZDB integrations omnirc options](#), on page 409.

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 how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

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. Right-click the backup specification you want to schedule and click **Edit Schedule**. The Scheduler page opens. All schedules available for this backup specification are listed in the right pane.
4. Click the schedule you want to edit, and then click the Edit icon. The Schedule wizard opens.

5. Review the options in the Options page, and click **Next**. The Recurrence page opens.

Note that 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.

NOTE:

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.

6. Set the Recurrence pattern, and click **Next**. The Summary page opens.
7. Review the options in the Summary page, and click **Finish**.

Running an interactive backup

An interactive backup can be performed any time after a backup specification has been 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** context.
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**.

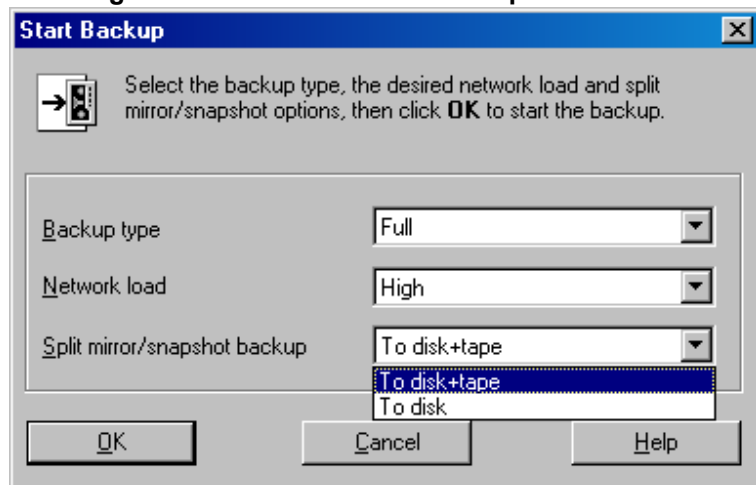
Note that only the Full backup type is supported.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. See [Selecting ZDB to disk or ZDB to disk+tape session when starting an interactive backup, on the next page](#).

NOTE:

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.

Selecting ZDB to disk or ZDB to disk+tape session when starting an interactive backup



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 *HPE 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.

Restore

You can restore the following database objects using both 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 [Duplicating an Oracle database](#), on page 84.

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 [Restoring from backup media to the application system on LAN, on the next page.](#)
- Instant recovery. See [Instant recovery and database recovery , on page 102.](#)


See also [Introduction, on page 15](#) for an overview of recovery methods depending on the backup type and type of recovery.

Microsoft Cluster Server systems

Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the Cluster Administrator utility.


See [Taking the Oracle resource group offline, below.](#)

Taking the Oracle resource group offline



Verify that you have set the **Prevent Fallback** option for the Oracle resource group and **Do not restart** for the `DB_NAME.world` resource, which is an Oracle Database resource.

Checking properties



HPE Serviceguard systems

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

```
run {  
  allocate channel dev1 type 'sbt_tape' parms  
    'SBT_LIBRARY=Path_to_Data_Protector_MML,  
    ENV=(OB2BARHOSTNAME=virtual.domain.com)';  
  restore datafile '/opt/ora10g/oradata/MAKI/example02.dbf';  
  release channel dev1;  
}
```

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.
One way of achieving this is by configuring static service information for your Oracle listener. For details, see the Oracle documentation. You can find an example of static service information configuration in the troubleshooting [Troubleshooting , on page 108](#).
- 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 [Restoring Oracle using the Data Protector GUI, below](#).
- RMAN. See [Restoring Oracle using RMAN , on page 90](#).

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 itself. You can also use the workaround described in [How to modify the RMAN restore script, on page 120](#).

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:

Required database states

Item to restore	Database state
Control file, duplicating a database	NoMount (started)
All other items ¹	Mount

To put the database into the correct state, execute:

```
sqlplus /nolog
```

```
SQL>connect user/password@service as sysdba;
```

```
SQL>shutdown immediate;
```

To put the database into NoMount state, execute:

```
SQL>startup nomount;
```

To put the database into Mount state, execute:

```
SQL>startup mount;
```

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use `as sysbackup` instead of `as sysdba`.

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`.

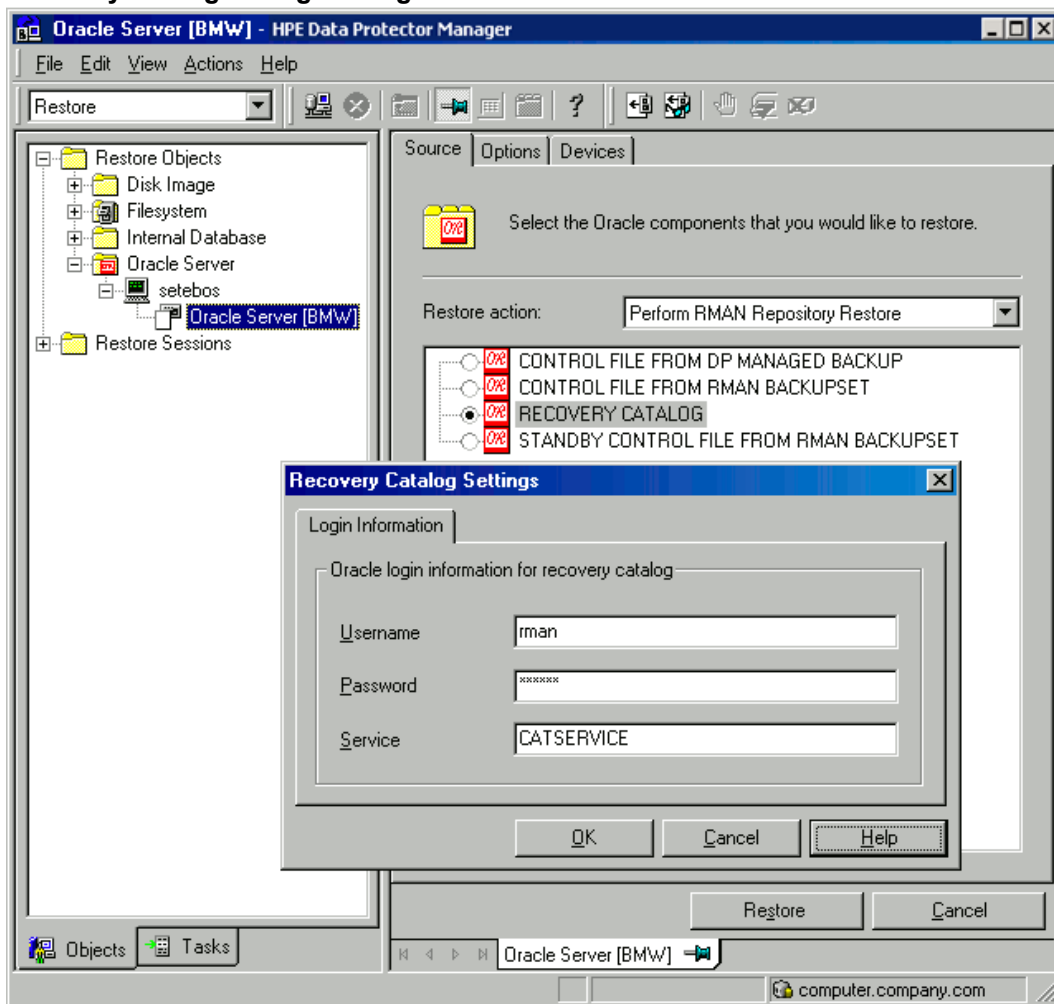
¹When restoring only a few tablespaces or datafiles, then the database can be open with the tablespaces or datafiles to be restored offline.

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**.

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.

Recovery catalog settings dialog



6. In the **Options** page:
In **User name** and **User group**, specify the user name and password to the recovery catalog database.
From the **Session ID** drop-down list, select the Session ID.
For further information, see [Restore, recovery, and duplicate options](#), on page 86.
7. 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 the default Data Protector temporary files directory.

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:

- Open the `sqlplus` window and put the database in the `nomount` state. See [Changing the database state](#), on page 77.
- In the Data Protector GUI, switch to the **Restore** context.
- 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.
- In the **Restore Action** drop-down list, select **Perform RMAN Repository Restore**.
In the Results area, select the control file for restore.
- 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 [Restore, recovery, and duplicate options](#) , on page 86.

6. Click **Restore**.

Proceed with restoring 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 [Restoring the recovery catalog database](#) , on page 77 and [Restoring the control file](#) , on the previous page.


To restore Oracle database objects:

1. Put the database in the mount state. See [Changing the database state](#) , on page 77.
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 [Restore, recovery, and duplicate options](#) , on page 86.

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 [Restoring Oracle using RMAN](#) , on page 90.

Source page



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.

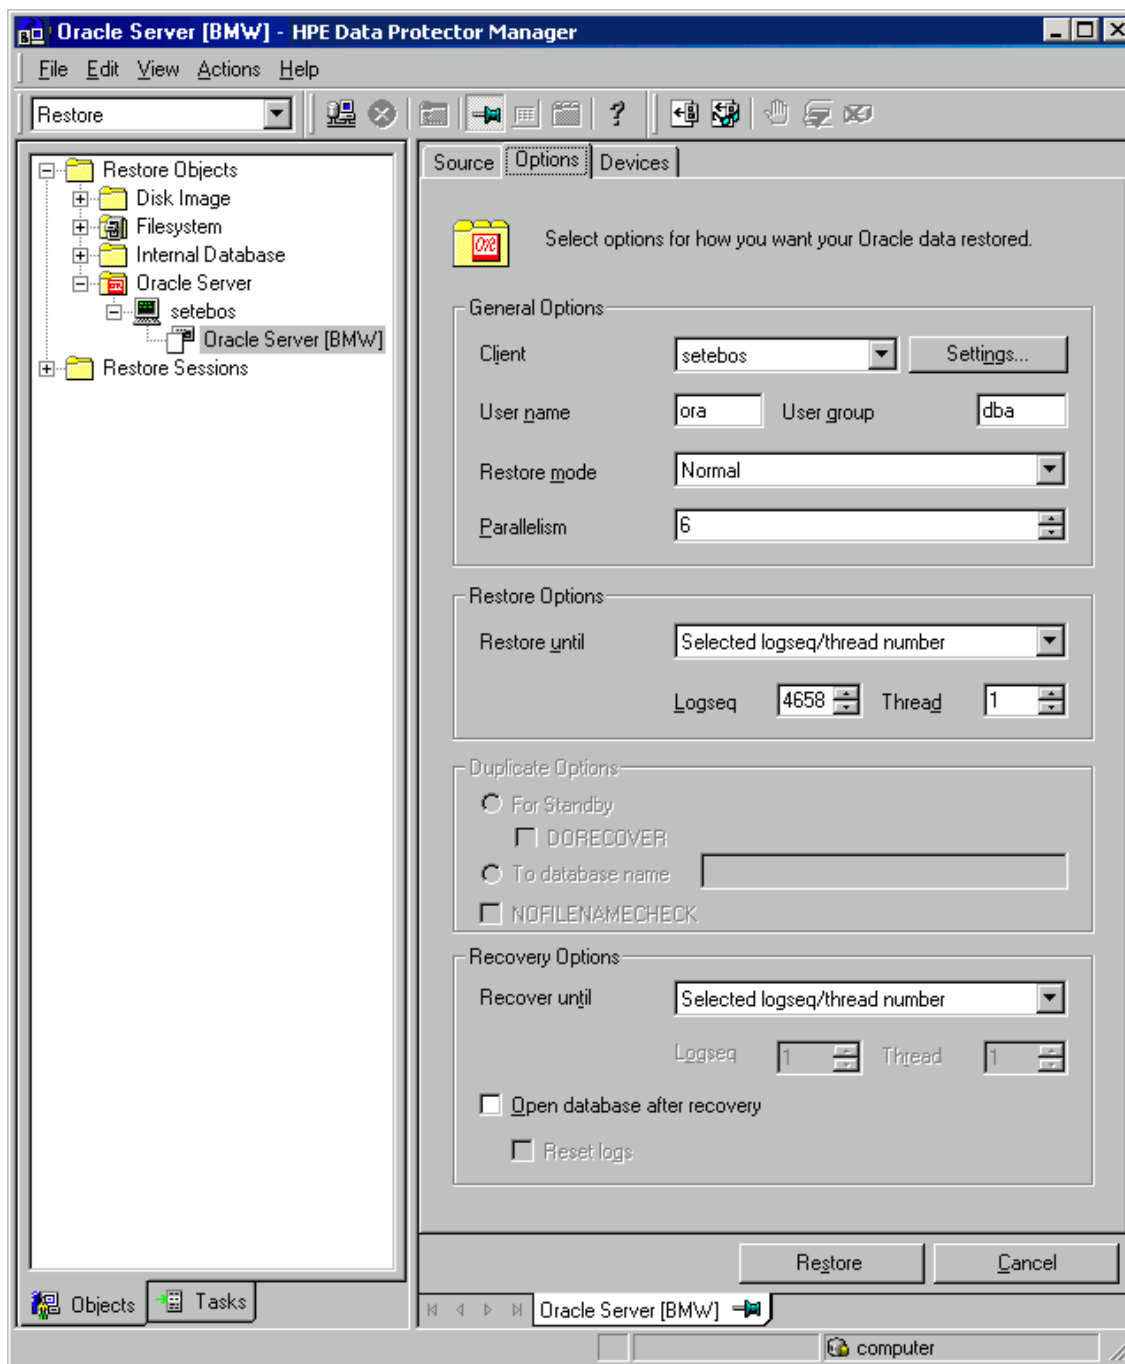
NOTE:

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.

Set the other restore options. For information, see [Restore, recovery, and duplicate options](#), on page 86.


Options page



7. 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 *HPE Data Protector Help* index: "restore, selecting devices for".

Devices page



8. Click **Restore**.

After the restore:

1. Put the database in the correct state.

If you selected **Perform Restore and Recovery** or **Perform Recovery Only** in the **Source** page, then the database is automatically put into **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_CatLog_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 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;
```

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use as sysbackup instead of as sysdba.

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

Open a command line window and enter the following commands:

```
sqlplus /nolog
SQL>connect user/password@service as sysdba
If you are restoring a datafile enter:
SQL>alter database datafile 'datafile_name' online;
```

If you are restoring a tablespace enter:

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

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use as sysbackup instead of as sysdba.

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 whole 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.

Limitations

- Database duplication is not supported using proxy copy backups of the primary database.
- If you perform duplication of a database (not for standby) on the same system on which the target or production database resides, note that you cannot use the same database name for the target and duplicate databases when the duplicate database resides in the same Oracle home directory as the target database. Note also that if the duplicate database resides in a different Oracle home directory than the target database, then the duplicate database name has to differ from other database names in that same Oracle home directory.

To duplicate a production database:

1. On the system where the selected database will be duplicated, put the Oracle auxiliary database instance in the nomount state. See [Changing the database state](#), on page 77.
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.p1) 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.p1) will be started.

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.

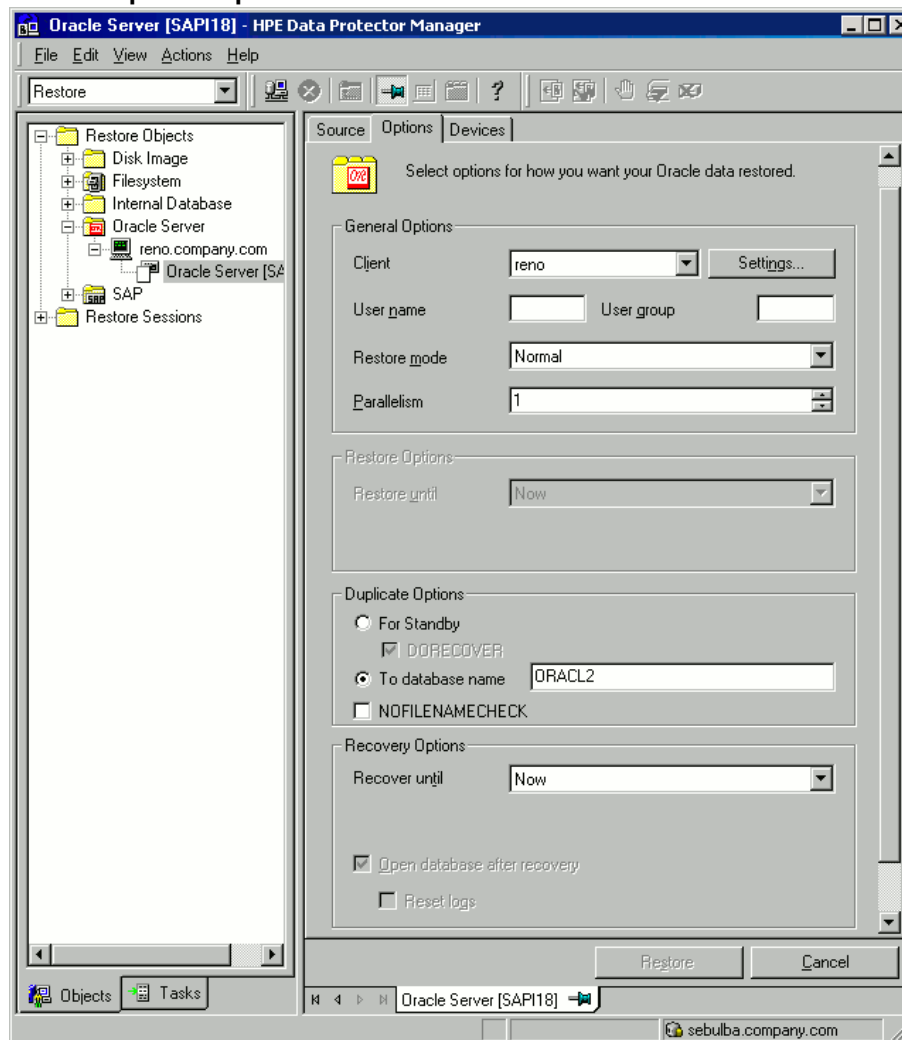
In **User name** and **User group**, specify the user name and group for the OSDBA account, which will be used by the Data Protector Oracle integration agent.

In **Parallelism**, specify the number of RMAN auxiliary channels to be allocated for database duplication.

Set duplicate options. For information, see [Duplicate options](#), on page 88 or press **F1**.

If you are creating a new database copy (not for standby), specify also the **Recover until** option to recover the duplicated database until a specified point in time.

Oracle duplicate options



6. Click **Restore**.

When the standby database is created, it is left mounted. Start the managed recovery process (log apply services) manually.

For information on how to use the RMAN commands to duplicate a database, see the Oracle documentation.

Restore, recovery, and duplicate options

Restore action options

The following describes each of the options in the **Source** page. This page is used to define the combination of restore and recovery you would like to perform using the GUI.

In the context of Data Protector, “restore” means to restore the datafiles. You can select which database, tablespace, or datafiles they would like to restore and up to which point in time they would

like them to be restored. “Recover” means applying the redo logs. You can select which redo logs to apply according to SCN number, logseq, or you can apply all the redo logs to the time of the last backup.

Perform Restore	Use this option to only restore (but not recover) the database objects using Data Protector. After restore, recover the database manually using RMAN. For information on recovering the database using RMAN, see Restoring Oracle using RMAN , on page 90 .
Perform Restore and Recovery	Use this option to perform both the restore and recovery of the database objects using Data Protector.
Perform Recovery Only	Use this option to only recover the database. This action can only be performed on the whole database.
Perform RMAN Repository Restore	Use this option to restore the recovery catalog or the control file when the database objects are not available in the Source page.
Perform Duplication	This option is used to perform duplication of a production database. This action can only be performed on the whole database.

General options

Client	This option specifies the system on which the Data Protector Oracle integration agent (ob2rman.p1) will be started.
Settings	<p>Click Settings to specify the login information (user name, password, and net service name) for the target database (in case of restore and recovery) or auxiliary database (in case of duplication) where you want the selected database objects to be restored or duplicated.</p> <p>If this is not specified in the case of restore or recovery, the login information of the selected database that resides on the selected system will be used.</p> <p>If this is not specified in the case of duplication, the duplication session will fail.</p>
User name, User group (UNIX systems only)	<p>Specify the operating system user account under which you want the restore to start.</p> <p>Ensure that this user has Oracle rights to restore the database (for example, it is in the DBA user group). The user must also be in the Data Protector admin or operator user group (actually, the Start restore and See private objects user rights suffice).</p>
Restore mode	<p>This drop-down list allows you to specify which type of restore you would like perform. The options are:</p> <ul style="list-style-type: none"> • Normal

	<p>This option should be used when a conventional backup or ZDB using the backup set method was performed.</p> <ul style="list-style-type: none"> Proxy copy <p>This option should be used when the original Oracle backup was made using the Oracle RMAN proxy-copy method.</p> <p>This option is disabled when you perform recovery only.</p>
Parallelism	<p>This field is used to specify the number of concurrent data streams that can read from the backup device. The default value is one.</p> <p>In case of Normal restore mode, to optimize restore performance, specify the same number of data streams as were used during the backup. For example, if you set the backup concurrency to 3, set the number of parallel data streams to 3 as well. Note that if a very high number of parallel data streams is specified this may result in a resource problem because too much memory is being used.</p> <p>For Oracle proxy-copy ZDB sessions, this option is disabled and Data Protector sets the number of concurrent data streams to the value that was used at backup. If you are restoring a backup created using a previous version of Data Protector, parallelism is set to the number of devices that were used for backup, regardless of the concurrency numbers for these devices.</p>

Duplicate options

Available if **Perform Duplication** was selected.

For Standby	<p>Select this option to create a standby database.</p> <p>Default: selected.</p>
DORECOVER	<p>Available if For Standby was selected.</p> <p>Select this option if you want RMAN to recover the database after creating it.</p>
To database name	<p>Select this option to create a new database copy. In the text box, specify its name. The name should match the name in the initialization parameter file that was used to start the auxiliary database instance. By default, the database name is set to the database name of the currently selected target database.</p>
NOFILENAMECHECK	<p>Select this option to disable RMAN to check whether the target datafiles share the same names with the duplicated datafiles.</p> <p>Select this option when the target datafiles and duplicated datafiles have the same names, but reside on different systems.</p> <p>Default: not selected.</p>

Restore and recovery options

Restore until	<p>The options in this drop-down list allow you to limit the selection to those backups that are suitable for an incomplete recovery to the specified time.</p> <ul style="list-style-type: none"> • Now Use this option to restore the full backup. By default, this option is selected. • Selected time Use this option to specify an exact time to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified time. • Selected logseq/thread number A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to restore. Data Protector restores the backup that can be used in recovery to the specified log sequence number. • Selected SCN number Use this option to specify the SCN number to which you wish the database to be restored. Data Protector restores the backup that can be used in recovery to the specified SCN number.
Recover until	<p>The options in this drop-down list allow you to specify to which point in time you would like the recovery to be performed.</p> <ul style="list-style-type: none"> • Now Data Protector starts RMAN to recover the database to the most recent time possible by applying all archived redo logs. By default, this option is selected. • Selected time Use this option to specify an exact time to which the archive logs are applied. • Selected logseq/thread number A logseq number is a redo log sequence number. Use this option to specify a particular redo log sequence and a thread number which will act as an upper limit of redo logs to recover. • Selected SCN number Use this option to specify the SCN number to which you perform the recovery. <p>If you reset the logs, also reset the database; otherwise, Oracle will during the next backup try to use the logs that were already reset and the backup will fail. Login to the target and recovery catalog database and execute:</p> <pre>rman target <i>Target_Database_Login</i> catalog <i>Recovery_CatLog_Login</i> RMAN> RESET DATABASE; RMAN> exit</pre>
Open database	<p>Opens the database after a recovery is performed.</p>

after recovery	
Reset logs	<p>Resets the archive logs after the database is opened.</p> <p><i>Always</i> reset the logs:</p> <ul style="list-style-type: none">• After an incomplete recovery (not Recover until now).• If a backup of a control file is used in recovery or restore and recovery. <p><i>Do not</i> reset the logs:</p> <ul style="list-style-type: none">• After a complete recovery (Recover until now) when the backup of a control file was not used in recovery or restore and recovery.• On the primary database, if the archive logs are used for a standby database. However, if you must reset the archive logs, you will need to recreate the standby database. <p>If you reset the logs when the Recover until option is set to Now, a warning is displayed, stating that you should reset the logs only if you use an older control file for restore.</p> <p>NOTE: Oracle recommends that you perform a complete backup immediately after a database was opened with the Reset Logs option.</p>

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 section only describes *examples* of how you can perform a restore. The examples provided do not apply to all situations where a restore is needed.

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.

The following examples of restore are given:

- [Example of full database restore and recovery, on page 93](#)
- [Example of point-in-time restore, on page 94](#)
- [Example of tablespace restore and recovery, on page 95](#)
- [Example of datafile restore and recovery, on page 97](#)
- [Example of archive log restore, on page 100](#)

The restore and recovery procedure of Oracle control files is a very delicate operation, which depends on the version of the Oracle database you are using. For detailed steps on how to perform the restore of control files, see the *Recovery Manager User's Guide and References*.

Preparing the Oracle database for restore

The restore of an Oracle database can be performed when the database is in 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:

- Make sure that the recovery catalog database is open. If the recovery catalog database cannot be brought online, you will probably need to restore the recovery catalog database. See [Restore , on page 74](#) for details of how to restore the recovery catalog database.
- 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. See the *Oracle Recovery Manager User's Guide and References* for more details.

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.

- Make sure that the following environment variables are set:
 - ORACLE_BASE
 - ORACLE_HOME
 - ORACLE_TERM
 - PATH
 - NLS_LANG
 - NLS_DATE_FORMAT

Windows systems example

```
ORACLE_BASE=Oracle_home  
ORACLE_HOME=Oracle_home\product\10.1.0  
ORACLE_TERM=HP  
PATH=$PATH:Oracle_home\product\10.1.0\bin  
NLS_LANG=american  
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

UNIX systems example

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

```
PATH=$PATH:/opt/oracle/product/10.1.0/bin
```

```
NLS_LANG=american
```

```
NLS_DATE_FORMAT='Mon DD YYYY HH24:MI:SS'
```

- Check 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 *HPE Data Protector Integration Guide*.

The `SBT_LIBRARY` path parameter must contain no spaces such that RMAN can load the library successfully. On a Windows system, you may run into problems configuring your Oracle backup, and the configuration may fail with the following error messages:

```
SBT_LIBRARY=C:/Program Files/OmniBack/bin/orasbt.dll
```

```
[...]
```

```
ORA-19554: error allocating device, device type: SBT_TAPE, device name:
```

```
ORA-27209: syntax error in device PARMS - unknown keyword or missing =
```

If, at this stage, `SBT_LIBRARY` parameter contains a path with space, you must check if the path `C:\Program Files` or the Data Protector installation directory, `DP_HOME` contains a shortpath (8dot3) notation with a CMD command `<dir /X>`.

If the shortpath does not exist or has been deleted, or disabled during installation, you can manually add using the following command:

```
<windows\System32\fsutil.exe file setshortname 'C:\Program Files' PROGRA~1>
```

Note that since `Program Files` directory is in use when the OS is active, the command will only work when you run the Windows repair mode CMD on your Oracle machine.

If you are not able to add the shortname because you can not take your production database offline for repair mode, use the following workaround:

1. Create a symbolic link to `<DP_HOME>\bin\orasbt.dll` using CMD `<mklink>` and save it to a path with no spaces.
2. Set the Oracle variable `SBT_LIBRARY` as described in the above procedure, and the `util_cmd` CLI reference pages.

If there are no Oracle configurations because of new installation, you must specify optional parameter to ensure that the config file is created:

```
-local
```

```
"C:\ProgramData\OmniBack\Config\Server\Integ\Config\Oracle8\hostname.domain.suffix%  
SID"
```

CAUTION:

If you use the above workaround, you are responsible for the location of the symbolic link, security, and access rights.

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:

```
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:

```
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 `restore_datafile` file in the default Data Protector temporary files 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:

```
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:

```
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 restore_PIT file in the default Data Protector temporary files 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`

Example of tablespace restore and recovery

If a table is missing or corrupted, you need to perform a restore and recovery of the entire tablespace. To restore a tablespace, you may take only a part of the database offline, so that the database does not have to be in the mount mode. You can use either a recovery catalog database or control files to perform a tablespace restore and recovery. Follow the steps below:

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 tablespace restore and recovery.

- If the database is in the open state, the script to restore and recover the tablespace should have the following format:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql 'alter tablespace TEMP offline immediate';
restore tablespace TEMP;
recover tablespace TEMP;
sql 'alter tablespace TEMP online';
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:

```
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)';
sql 'alter tablespace TEMP offline immediate';
restore tablespace TEMP;
release channel 'dev1';
recover tablespace TEMP;
sql 'alter tablespace TEMP online';
release channel 'dev2';
}
```

- If the database is in the mount state, the script to restore and recover the tablespace should have the following format:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore tablespace 'TEMP';
recover tablespace 'TEMP';
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:

```
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)';
restore tablespace 'TEMP';
release channel 'dev1';
recover tablespace 'TEMP';
release channel 'dev2';
}
```

You can also save the script into a file and perform a tablespace restore using the saved files:

1. Create a restore_TAB file in the default Data Protector temporary files directory.
2. Start the tablespace restore.

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

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

Example of datafile restore and recovery

To restore and recover a datafile, you may take only a part of the database offline.

To restore and recover a datafile:

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 datafile restore and recovery:

- If the database is in an open state, the script to restore the datafile should have the following format:

UNIX systems

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' offline";
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' online";
release channel dev1;
}
```

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:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' offline";
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
```

```
release channel dev2;
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
sql "alter database datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf' online";
release channel dev1;
}
```

Windows systems

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
sql "alter database datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf' offline";
restore datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf';
recover datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf';
sql "alter database datafile
  'C:\oracle\data\oradata\DATA\temp01.dbf' online";
release channel dev1;
}
```

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

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
sql "alter database datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf' offline";
restore datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev2;
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
sql "alter database datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf' online";
release channel dev1;
}
```

- If the database is in a mount state, the script to restore and recover the datafile should have the following format:

UNIX system

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev1;
}
```

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

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
allocate channel dev2 type 'sbt_tape' parms
  'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';
restore datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev2;
recover datafile
  '/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev1;
}
```

Windows system

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel dev1 type 'sbt_tape' parms
  'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,
  ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
restore datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
recover datafile
  'Oracle_home\data\oradata\DATA\temp01.dbf';
release channel dev1;
}
```

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

```
run{
allocate channel dev1 type 'sbt_tape' parms
```

```
'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,  
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';  
allocate channel dev2 type 'sbt_tape' parms  
'SBT_LIBRARY=Data_Protector_home\bin\orasbt.dll,  
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME, OB2PROXYCOPY=1)';  
restore datafile  
'Oracle_home\data\oradata\DATA\temp01.dbf';  
release channel dev2;  
recover datafile  
'Oracle_home\data\oradata\DATA\temp01.dbf';  
release channel dev1;  
}
```

You can also save the script into a file and perform a datafile restore using the saved files:

1. Create a restore_dbf file in the default Data Protector temporary files directory.
2. Start the datafile restore:

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

UNIX systems: ORACLE_HOME/bin/rman target sys/manager@PROD catalog
rman/rman@CATAL cmdfile=Data_Protector_home/tmp/restore_dbf

Example of archive log restore

To restore an archive log:

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 archive log restore:

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

You can also save the script into a file and perform an archive log restore using the saved files:

1. Create a restore_arch file in the default Data Protector temporary files directory.
2. Start the archive log restore:

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

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

Example of database restore using a different device (with the automatic device selection functionality disabled)

Suppose a database was backed up with the device dev1. To restore the database with the device dev2, add the line `send device type 'sbt_tape' 'CHDEV=dev1>dev2'`; to the RMAN script:

1. Log in to the Oracle RMAN:

Windows systems: `ORACLE_HOME\bin\rman target sys/manager@TIN`

UNIX systems: `ORACLE_HOME/bin/rman target sys/manager@TIN`

2. Execute:

```
run {
allocate channel 'dev_0' type 'sbt_tape'
  parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
allocate channel 'dev_1' type 'sbt_tape'
  parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
allocate channel 'dev_2' type 'sbt_tape'
  parms 'SBT_LIBRARY=C:/PROGRA~1/OmniBack/bin/orasbt.dll,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=CAN,OB2BARLIST=test)';
send device type 'sbt_tape' 'NO_AUTO_DEVICE_SELECTION=1';
send device type 'sbt_tape' 'CHDEV=dev1>dev2';
restore database;
}
```

NOTE:

- The line `device type 'sbt_tape' 'NO_AUTO_DEVICE_SELECTION=1'`; disables the automatic device selection.
- You can also specify multiple device redirections using CHDEV parameter with `'CHDEV=dev1>dev2;dev3>dev4'`; syntax.

Restoring using another device

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

Specify these devices in the `/etc/opt/omni/server/cell/restoreddev` (UNIX systems) or `Data_Protector_program_data\Config\server\Cell\restoreddev` (Windows systems) file in the following format:

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

where

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

On Windows systems, this file must be in the Unicode format.

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

Example

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

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

Restoring Oracle to different cell/client

To restore Oracle to different cell/client perform the following steps:

1. Copy the Data Protector Oracle configuration files from source to target.
2. Export the Oracle SID.
3. Create password file for the database.
4. Set the DBID and start the database in nomount state.
5. Recover the **pfile** from RMAN.
6. Start the database in nomount state with recovered **pfile**.
7. Start the control file restore from GUI and copy the Data Protector restored control file to control file location.
8. Start the database in mount state.
9. Configure the database backup with control file from target side.
10. Start the database restore only.
11. Execute the sql command "alter database open resetlogs" from the target database to recover the database.

Instant recovery and database recovery

For general information on instant recovery, see the *HPE Data Protector Concepts Guide* and the *HPE Data Protector Zero Downtime Backup Administrator's Guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), HPE Serviceguard, and Microsoft Cluster Server), see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

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.

IMPORTANT:

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

Limitations

- For ZDB-to-disk sessions, only archived redo logs can be used for a database recovery after an instant recovery.
- The recovery process will fail if the log entry with the specified logseq number or SCN number was created before the target volume.

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 instant recovery is to be performed to some other node than the one that was backed up, the following procedure must be performed before the standard instant recovery procedure:

1. Make sure that the HPE 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
```

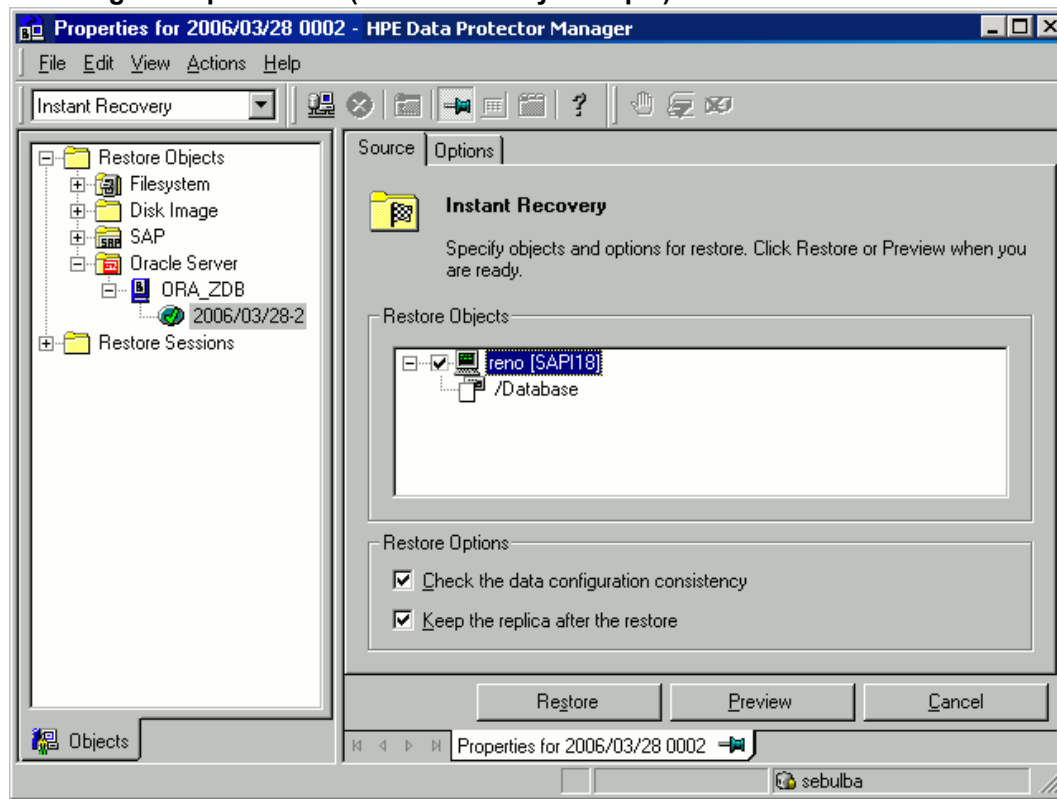
Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Put the Oracle database instance in a nomount state using `sqlplus`. In case of RAC, set all instances to nomount state.
For example:

```
sqlplus
sql> shutdown immediate
sql> startup nomount
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. With HPE P9000 XP Disk Array Family, it is recommended to leave the **Keep the replica after the restore** option set to enable a restart of an instant recovery session. With HPE P6000 EVA Disk Array Family, replica is kept on the disk array only if the **Copy replica data to the source location** is selected.
Set other HPE P6000 EVA Disk Array Family or HPE P9000 XP Disk Array Family options. For details, press **F1**.

Selecting backup sessions (P9000 XP Array example)



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. See [Oracle database recovery after the instant recovery , on the next page](#).

- To perform a database recovery immediately after an instant recovery, click on 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. See [Oracle recovery options, below](#) and [Restore, recovery, and duplicate options , on page 86](#) for details on available options.

Oracle recovery options



6. Click **Restore** or **Preview**. Note that preview 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 performing instant recovery by switching the database to a mount state, restoring the necessary incremental backups and archived redo logs from tape, and applying the redo logs.

If you reset the logs, reset the database; otherwise, Oracle will during the next backup 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_CatLog_Login  
RMAN> RESET DATABASE;  
RMAN> exit
```

Oracle database recovery after the instant recovery

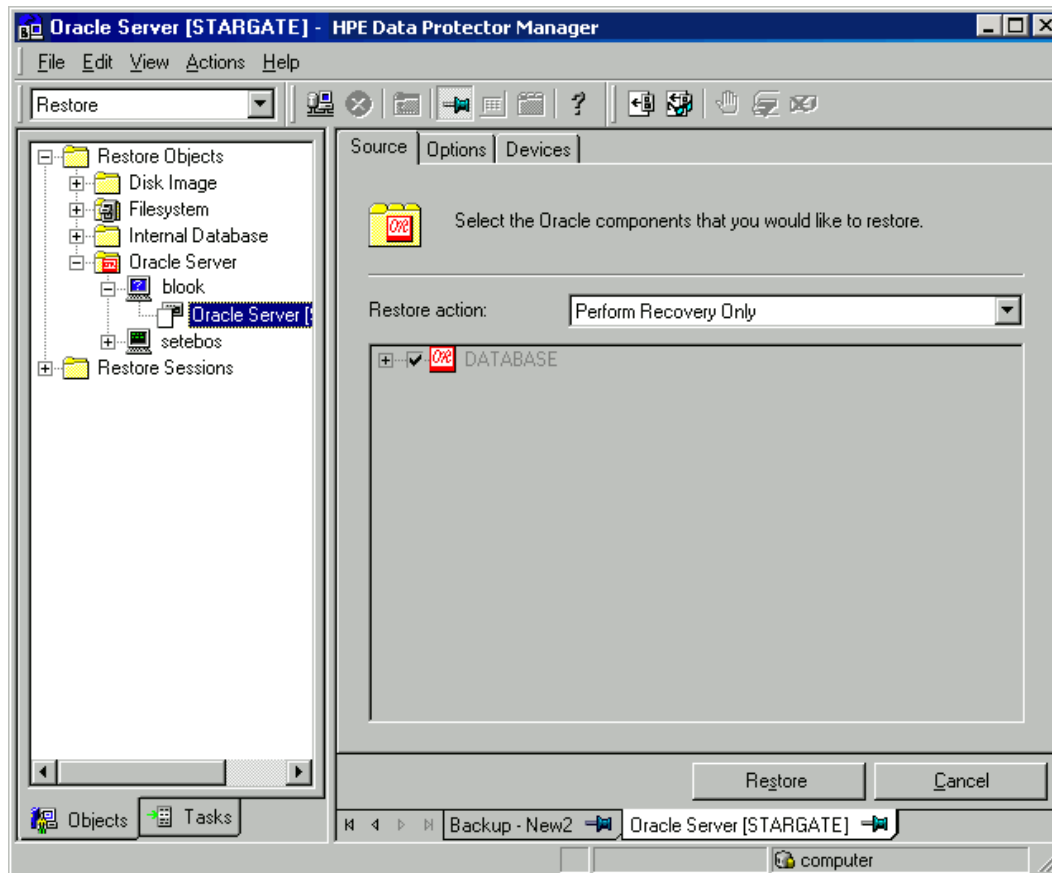
To recover the Oracle database after the instant recovery has been performed, perform the following steps:

1. Put the Oracle database in a mount state by connecting to the target database from the sqlplus and then running the following command:

```
startup mount
```
2. To recover the database, the following two options are available:

- Perform a recovery from the Data Protector Manager Restore Context:
 - a. Expand **Oracle Server** and select the database to recover. In the **Source** tab, under **Restore action**, select **Perform recovery only**.

Selecting the database for recovery



- b. In the **Options** tab, select the recovery options. For details, see [Restore, recovery, and duplicate options](#), on page 86.
 - 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 [Restoring Oracle using RMAN](#), on page 90.

Oracle in Veritas Cluster instant recovery

If Oracle on the application system runs in a Veritas Cluster, the following two Veritas Cluster resources must be disabled before instant recovery is performed, and enabled after instant recovery has finished to prevent the failover of the Oracle Veritas Cluster Service Group:

- Veritas Cluster application resource for the Oracle application and
- Veritas Cluster mountpoint resource for the Oracle database files.

Follow the steps below to perform an instant recovery to the application system with Oracle in a Veritas Cluster:

1. On the application system, enter the following commands to disable the two Veritas Cluster resources:
 - a. `hares -offline application_resource_name -sys system`
where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application and *system* is the name of the active node.
`hares -offline mountpoint_resource_name -sys system`
where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *system* is the name of the active node.
 - b. `hares -modify application_resource_name Enabled 0`
where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application.
`hares -modify mountpoint_resource_name Enabled 0`
where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.
2. Perform an instant recovery.
3. If you performed only an instant recovery without the database recovery, use RMAN as described in [Oracle database recovery after the instant recovery](#), on page 105 to bring the Oracle database to a consistent state.
4. On the application system, enter the following commands to enable the two Veritas Cluster resources:
 - a. `hares -modify mountpoint_resource_name Enabled 1`
where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.
`hares -modify application_resource_name Enabled 1`
where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application.
 - b. `hares -online application_resource_name -sys system`
where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application and *system* is the name of the active node.
`hares -online mountpoint_resource_name -sys system`
where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *system* is the name of the active node.

Aborting sessions

You can abort currently running sessions by clicking the abort button.

If, during a session, RMAN or SQL*Plus do not respond when requested, 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 of how to set environment variables, see [Setting environment variables](#), on page 51. For details of how to set the corresponding `omnirc` options, see the *HPE Data Protector Help* index: “omnirc option”. Note that environment variables override `omnirc` options.

Troubleshooting

This section contains a list of general checks and verifications and a list of problems you might encounter when using the Data Protector Oracle integration. You can start at [Problems, on page 114](#) and if you cannot find a solution there, perform general checks and verifications.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery troubleshooting information, see the troubleshooting sections in the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HPE Data Protector Help* index: “patches” on how to verify this.
- See the *HPE Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <https://softwaresupport.hpe.com/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

For more detailed information on performing any of the following procedures, see the Oracle documentation.

If your configuration, backup, or restore failed:

- On the application system, verify that you can access the Oracle target database and that it is opened:
 1. Perform the following:
 - Windows systems:** Set the `ORACLE_HOME` variable.
 - UNIX systems:** Export the `ORACLE_HOME` variable as follows:

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

```
ORACLE_HOME="ORACLE_HOME"  
export ORACLE_HOME
```
- If you are using a csh-like shell, enter the following commands:

```
setenv ORACLE_HOME "ORACLE_HOME"
```
- 2. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```
- 3. Start SQL*Plus and type:

```
connect user_name/password@service as sysdba;  
select * from dba_tablespaces;  
exit
```

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use as sysbackup instead of as sysdba.

If this fails, open the Oracle target database.

- On the application system, verify that you can access the recovery catalog (if used) and that it is opened as follows:
 1. Export or set the *ORACLE_HOME* variable as described in [Perform the following;, on the previous page](#).
 2. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```
 3. Start SQL*Plus and type:

```
connect Recovery_Catalog_Login  
select * from rcver;  
exit
```

If this fails, open the recovery catalog.

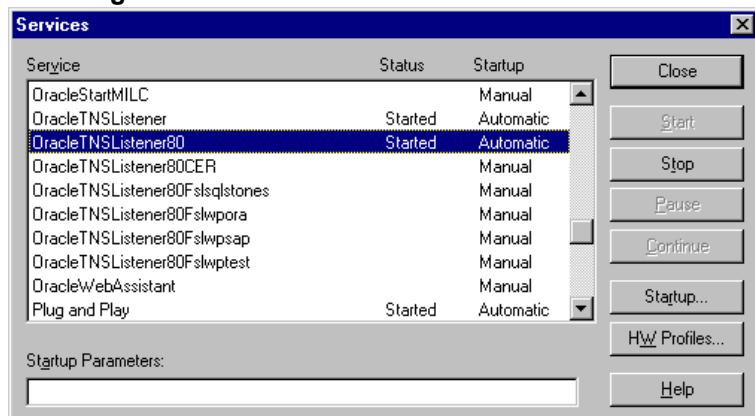
- Verify that the listener is correctly configured for the Oracle target database and the recovery catalog database. This is required to properly establish network connections:
 1. Export or set the *ORACLE_HOME* variable as described in [Perform the following;, on the previous page](#).
 2. Start the listener from the bin directory in the *ORACLE_HOME* directory:

```
lsnrctl status service
```

If this fails, startup the listener process and see the Oracle documentation for instructions on how to create a configuration file (LISTENER.ORA).

On Windows, the listener process can be started in the Control Panel > Administrative Tools > Services.

Checking the status of the Oracle listener



The status of the respective listener service in the **Services** window should be **Started**, otherwise you must start it manually.

3. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```

4. Start SQL*Plus and type:

```
connect Target_Database_Login
```

```
exit
```

and then

```
connect Recovery_Catalog_Login
```

```
exit
```

If this fails, see the Oracle documentation for instructions on how to create a configuration file (NAMES.ORA).

- Verify that the Oracle target database configured to allow remote connections with the system privileges:
 1. Export or set the *ORACLE_HOME* variable as described in [Perform the following:, on page 108](#).
 2. Start SQL*Plus from the bin directory in the *ORACLE_HOME* directory:

```
sqlplus /nolog
```
 3. Start SQL*Plus and type:

```
connect Target_Database_Login as SYSDBA
```

```
exit
```

Repeat the procedure using SYSOPER instead of SYSDBA.

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use as sysbackup instead of as sysdba.

If this fails, see the Oracle documentation for instructions on setting up the password file and any relevant parameters in the *initDB_NAME.ora* file.

- From the application system, verify that the target database and recovery catalog database are configured to allow remote connections with the system privileges and to allow backup:

- If you use the recovery catalog database:

Export or set the `ORACLE_HOME` and `DB_NAME` variables as described in [Perform the following;](#) on [page 108](#).

```
SQL> connect login_to_recovery_catalog_or_target_database as sysdba;  
> exit
```

```
ORACLE_HOME /bin/rman target login_to_target_database catalog login_to_  
Recovery_CatLog
```

- If you do not use the recovery catalog database:
- If you use the recovery catalog database, verify that the target database is registered in the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Perform the following;](#) on [page 108](#).

2. Start SQL*Plus from the `bin` directory in the `ORACLE_HOME`; directory:

```
sqlplus /nolog
```

3. Start SQL*Plus and type:

```
connect Recovery_CatLog_Login;  
select * from rc_database;  
exit
```

If this fails, start the configuration using Data Protector on the application system, or see the Oracle documentation for information on how to register an Oracle target database in the recovery catalog database.

- On the application system, verify backup and restore directly to disk using an RMAN channel type disk:

If you use the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Perform the following;](#) on [page 108](#).

2. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login catalog Recovery_CatLog_Login
```

If you do not use the recovery catalog:

1. Export or set the `ORACLE_HOME` variable as described in [Perform the following;](#) on [page 108](#).

2. Start RMAN from the `bin` directory in the `ORACLE_HOME` directory:

```
rman target Target_Database_Login nocatalog
```

An example of the RMAN backup script is presented below:

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

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

```
run {  
allocate channel 'dev0' type disk;
```

```
sql 'alter tablespace tablespace_name offline immediate';  
restore tablespace tablespace_name;  
recover tablespace tablespace_name;  
sql 'alter tablespace tablespace_name online';  
}
```

If this fails, see the Oracle documentation for details of how to execute a backup and restore directly to disk using RMAN.

Additionally, if your configuration or backup failed:

- Verify that the Data Protector software has been installed properly.
For details, see the *HPE Data Protector Installation Guide*.
- Check if the SYSDBA/SYSBACKUP privilege is granted to the Oracle administrator.
- If you have special Oracle environment settings, ensure that they are entered in the Data Protector Oracle configuration files on the Cell Manager. For information on setting the variables in the Data Protector Oracle configuration files, see the `util_cmd` man page or the *HPE Data Protector Command Line Interface Reference*.
- Perform a filesystem backup (non-ZDB) of the Oracle Server system so that you can eliminate any potential communication problems between the Oracle Server and the Data Protector Cell Manager system.

For details on performing filesystem backups, see the *HPE Data Protector Help* index: “standard backup procedure”.

Ensure that the hostname defined in the backup specification as a system to be backed up is the name of the application system.

- **Windows systems:** Check the Data Protector Inet service startup parameters on the Oracle Server system:
Go to **Control Panel > Administrative Tools > Services > Data Protector Inet**.
The service must run under a specified user account. Make sure that the same user is also added to the Data Protector `admin` or `user` group.
- Examine the system errors reported in the following file on the application system (Oracle proxy-copy ZDB method) or backup system (Oracle backup set ZDB method) into the `debug.log` file.

Additionally, if your backup or restore failed:

- Test the Data Protector internal data transfer using the `testbar2` utility:
 1. Verify that the Cell Manager name is correctly defined on the Oracle Server system. Check the `cell_server` file located in the default Data Protector client configuration directory, which contains the name of the Cell Manager system.
 2. From the `bin` directory in the `ORACLE_HOME` directory, execute:

If backup failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:backup -bar:backup_  
specification_name
```

If restore failed:

```
testbar2 -type:Oracle8 -appname:DB_NAME-perform:restore
```

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

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

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

- Check if the user under which the backup or restore session was started has appropriate Oracle permissions (for example, belongs to the DBA group). This user must also be in the Data Protector operator or admin user group.
- Check that the respective Data Protector user group has the See private objects user right enabled.
- **If backup failed:** Create an Oracle backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices. See the *HPE Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.
- **If restore failed:** Execute the omnidb command to see objects in the database.

If the test fails again, call a support representative for assistance.

Additionally, if your restore failed:

- Verify that an object exists on the backup media.

This can be done by executing the following command on the Oracle server system from the bin directory in the *ORACLE_HOME*; directory:

```
omnidb -oracle8 "object_name" -session "Session_ID" -media
```

The output of the command lists detailed information about the specified Oracle object, as well as the session IDs of the backup sessions containing this object and a list of the media used. For detailed syntax of the omnidb command, see its man page.

- Ensure that the database is in the correct state.

If you are trying to restore a database item using the Data Protector GUI and the GUI stops responding, try one of the following:

- If you are restoring the control file, the database should be in the NoMount state.

Open a command window and enter the following:

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

- If you are restoring datafiles, the database should be in the Mount state.

Open a command window and enter the following:

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

- If there is a problem you cannot resolve while you are trying to restore a database item using the

Data Protector GUI, try using the RMAN CLI to restore the database items.

For information, see [Restoring Oracle using RMAN](#) , on page 90.

- Try putting the database into the Open state manually after using the Data Protector GUI to recover and restore a backup session.

If you have used the Data Protector GUI to recover and restore a backup session and you see the following error message:

```
Oracle Error: ORA-1589: must use RESETLOGS or NORESETLOGS option for database open.
```

Open a SQLplus window and use the following command:

```
sqlplus/nolog
```

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

```
SQL>alter database open noresetlogs;
```

If this does not work, try using the following command:

```
SQL>alter database open resetlogs;
```

NOTE:

With Oracle 12c version, if the user has SYSBACKUP privilege, then you must use as sysbackup instead of as sysdba.

Problems

Problem

SQL*Plus is unable to connect to destination

Action

Check if the Oracle listener process is up and running. Check if there are any environment variables you need to enter (for example, TNS_ADMIN). Enter these variables in the Data Protector Oracle configuration files on the Cell Manager. For information, see the `util_cmd` man page.

Problem

The following error is displayed: ORA-12532: : invalid argument

If this is reported by SQL*Plus in the Data Protector monitor, the application system may be low on resources (CPU, memory, and so on).

Action

Try to configure the application system in such a way that it consumes as little resources as possible. This error can be reproduced without using Data Protector by starting SQL*Plus on the application system, and connecting to the target database on the application system.

Problem

Backup set ZDB is aborted after 10 minutes

While performing a backup set ZDB, the following warning is displayed for each database datafile:

```
RMAN-06554: WARNING: file n is in backup mode
```

The ZDB session then aborts with the following message:

```
Bar backup session was started but no client connected in 600 seconds.
```

Action

Increase the value of the following global options (by default, these options are set to 10):

- If you upgraded Data Protector from a previous version of Data Protector:

```
SmWaitForFirstClient=minutes
```

- If you performed a clean installation:

```
SmWaitForFirstBackupClient=minutes
```

For more information about the global options, see the *HPE Data Protector Troubleshooting Guide*.

Problem

Backup set ZDB fails after changing the physical schema of the Database

Backup fails after you have modified the physical schema of a database, for example, if you added or dropped a tablespace, added a new datafile, or added or dropped a rollback segment. Depending on the performed modification, different error messages are displayed, for example:

```
RMAN-06056: could not access datafile datafile
```

The problem occurs because the physical schema of target database is not updated in the recovery catalog.

Action

Manually re-synchronize the recovery catalog database with the current control file.

Problem

On UNIX systems, a backup set ZDB-to-disk+tape session fails

While performing a backup set ZDB to disk+tape, the session fails with the following error when Oracle Server attempts to start up an instance on the backup system:

```
[Major] From: ob2rman@computer.company.com DB_NAME Time: Date Time
```

```
The database reported error while performing requested operation.
```

This problem occurs when either the user ID or the group ID number of the Oracle operating system user account on the application and backup systems do not match. Under such circumstances, Oracle Server is unable to start up the instance on the backup system due to missing privileges.

Action

Configure the user accounts as described in [Configuring Oracle operating system user accounts](#), on [page 38](#), and restart the session.

Problem

Proxy copy restore fails

Proxy copy restore fails with the following error:

RMAN-10035: exception raised in RPC: ORA-27197: skgfprs: sbtpcrestore returned error

RMAN-10031: ORA-27197 occurred during call to DBMS_BACKUP_RESTORE.PROXYRESTOREDATAFILE

Action

Check the IDB for the session and the objects of the latest backup. You might check if a more recent session exists in the recovery catalog. Connect to the RMAN prompt:

```
rman target user/password@TGT_DB catalog user/password@CDB
```

At the RMAN> prompt, enter

```
list backup;
```

to display a list of the objects in the recovery catalog. Check the list of Proxy Copy sessions, listed at the end.

To synchronize the recovery catalog and the IDB, execute the RMAN command:

```
resync catalog;
```

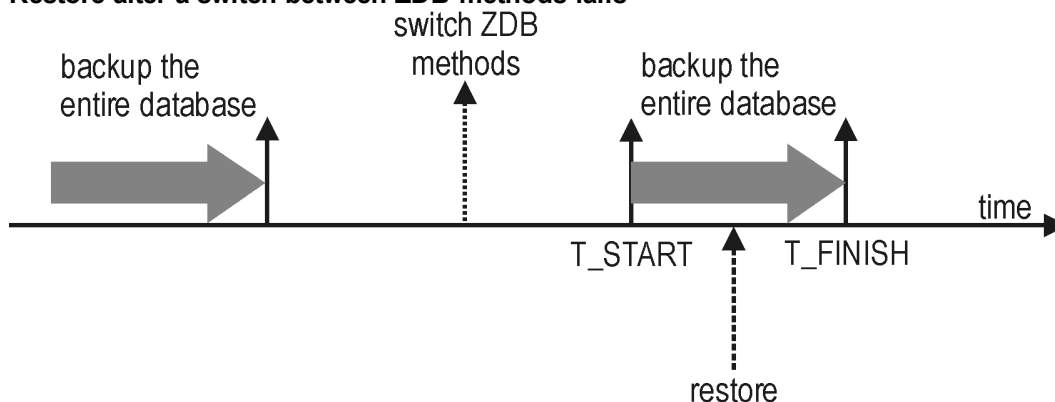
After the synchronization is performed, restore should be possible.

Problem

Restore after a switch between ZDB methods fails

If you perform a restore to a specified time ($T_RESTORE$) that lies in the time interval between the start of the first backup of the entire database using the new method (T_START), and before this backup is finished (T_FINISH), RMAN may try to restore the backup files made with the new method using a channel allocated for backup files made using the previous method. As a result, the restore procedure fails.

Restore after a switch between ZDB methods fails



Action

Restore the backup session manually using RMAN scripts. Add the required parameter to the allocated channels, that is `OB2PROXYCOPY=1` for the channel which will be used for restoring the backup made using the proxy-copy ZDB method. Then restore the backup files using the correct channels.

For example, if you switched from the backup set to the proxy-copy ZDB method, the script may look similar to the following one:

```
run {
  ALOCATE CHANNEL 'dev_0' TYPE 'sbt_tape'
  PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,
        ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
  ALOCATE CHANNEL 'dev_1' TYPE 'sbt_tape'
  PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,
        ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,
        OB2PROXYCOPY=1)';
  RESTORE DATAFILE list_of_backup_set_backups UNTIL
    T_RESTORE CHANNEL 'dev_0';
  RESTORE DATAFILE list_of_proxy-copy_backups UNTIL
    T_RESTORE CHANNEL 'dev_1';
  RELEASE 'dev_0';
  RECOVER DATABASE UNTIL T_RECOVER ...
  RELEASE 'dev_1';
}
```

Where:

T_RESTORE specifies the time to which to restore and *T_RECOVER* the time to which to apply the transactions.

list_of_backup_set_backups is a list of backups of the entire database using the backup set ZDB method.

list_of_proxy_copy_backups is a list of datafile backups completed after the start of the backup of the entire database (T_START) and before *T_RESTORE*.

Problem

Data Protector reports errors when calling SYS.LT_EXPORT_PKG.schema_inf_exp during Oracle backup

The following errors are listed in the Data Protector monitor:

EXP-00008: ORACLE error 6550 encountered

```
ORA-06550: line 1, column 13:
PLS-00201: identifier 'SYS.LT_EXPORT_PKG' must be declared
ORA-06550: line 1, column 7:
PL/SQL: Statement ignored
EXP-00083: The previous problem occurred when calling
SYS.LT_EXPORT_PKG.schema_info_exp
. exporting statistics
Export terminated successfully with warnings.
[Major] From: ob2rman.pl@machine "MAKI" Time: 10/01/01 16:07:53
Export of the Recovery Catalog Database failed.
```

Action

Start SQL*Plus and grant the execute permission to the LT_EXPORT_PKG as follows (make sure that the user sys has the SYSDBA privilege granted beforehand):

```
sqlplus 'sys/password@CDB as sysdba'  
SQL> grant execute on sys.lt_export_pkg to public;
```

Restart the failed backup session.

Problem

On a UNIX system, Data Protector reports “Cannot allocate/attach shared memory”

Backup fails and the following error message is displayed:

```
Cannot allocate/attach shared  
memory (IPC Cannot Allocate Shared Memory Segment)  
System error: [13] Permission denied) => aborting
```

Action

Set the OB2SHMEM_IPCGLOBAL omnirc option to 1 to use the memory windowing properly, and restart the failed backup session. See the *HPE Data Protector Troubleshooting Guide* for details on using the omnirc file.

Problem

Backup fails after a point in time restore and recovery

The following error is displayed:

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

Action

Connect to the target and recovery catalog database using RMAN and reset the database to register the new incarnation of database in the recovery catalog:

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

Problem

Oracle online backup fails with the following error:

```
RMAN-06004: ORACLE error from recovery catalog database: RMAN-20220: controlfile  
copy not found in the recovery catalog
```

When running an online backup, Data Protector adds the filename of the *controlfilecopy* to the RMAN backup script. This filename has to be cataloged to the RMAN catalog prior to the backup command.

Action

To catalog the *controlfilecopy* to the RMAN catalog:

1. Connect to RMAN on the application system.
2. Execute the following command:

```
RMAN> catalog controlfilecopy 'CONTROL_FILE_LOCATION/ctrlDB_NAME.ctl'
```

Problem

Backup of archive logs on RAC cannot be performed

On RAC, the archive logs are not installed on a NFS mounted disk. Backup of archive logs cannot be performed.

Action

Edit the archive logs backup specification:

- Add an additional `allocate channel` command for *each* node.
- Add a command to connect to each instance. The connection parameters should be given as `username/passwd@INSTANCE`.

For example, if you are using two nodes, the backup specification might look as follows:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,OB2BARLIST=RAC_arch)'
connect username/passwd@INSTANCE_1;
allocate channel 'dev_2' type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,OB2BARLIST=RAC_arch)'
connect username/passwd@INSTANCE_2;
backup
format 'RAC_arch<QU_%s:%t:%p>.dbf'
archivelog all;
}
```

Problem

The Recovery Catalog was lost and the control file cannot be restored from Data Protector managed backup

The Recovery Catalog was not used, the RMAN autobackup feature was not used, and the control file cannot be restored from Data Protector managed backup. A valid control file backup exists on tape.

Action

Restore the control file from RMAN backup set, mount and restore the database, and perform database recovery:

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML';
restore controlfile from 'backup piece handle';
sql 'alter database mount';
set until time 'MMM DD YY HH24:MM:SS';
}
```

```
restore database;  
recover database;  
sql 'alter database open resetlogs';  
release channel 'dev_0';  
}
```

At this point you must manually register any backups made after the control file backup that was restored. After that, continue with the restore procedure.

For the *backup piece handle* search the Data Protector Internal Database and session outputs of previous backup sessions.

Problem

How to modify the RMAN restore script

When you start a restore of an Oracle database using the Data Protector GUI or CLI, an RMAN restore script is created, which is instantly run, so you cannot edit it first.

Action

To edit the script before it is run, set the Data Protector `omnirc` option `OB2RMANSAVE` to point to an existing directory. When the variable is set and you start a restore, the RMAN restore script, which is created at run time, is saved to the specified location under the name `RMAN_restore_backup_specification_name.rman`, and the actual restore is skipped. Then you can edit the script and run it manually afterwards. On how to set the `omnirc` options, see the *HPE Data Protector Help* index: "omnirc options".

To start a restore using Data Protector again, clear the `OB2RMANSAVE` option by deleting its content or commenting or removing the whole option. If you comment or remove the option on a Windows system, restart the Data Protector `Inet` service for the settings to take effect.

Problem

Instant recovery session for an Oracle database fails

On a Windows Server 2008 x64 system, when you run an instant recovery from a backed up Oracle 11g application database that resides on a disk array of the HPE P6000 EVA Disk Array Family, the session may fail with the following error messages:

```
[Major] From: SMISA@appsystem.company.com "SMISA"  
Time: 5/17/2010 5:00:50 AM  
A filesystem could not be mounted.  
Filesystem name :  
Mount point : I:\  
  
[Critical] From: SMISA@appsystem.company.com "SMISA"  
Time: 5/17/2010 5:00:50 AM  
Failed to resume the application system.  
  
[Critical] From: SMISA@appsystem.company.com "SMISA"  
Time: 5/17/2010 5:00:50 AM  
Instant Recovery failed.
```

Such an instant recovery failure occurs when the “copy-back” instant recovery method (the default) is chosen and the following instant recovery options are selected in the Source and Options panes of the Data Protector GUI:

- Wait for the replica to complete (with the default waiting period used)
- Retain source for forensics
- Check the data configuration consistency
- Force the removal of all replica presentations
- Recovery
- Open database after recovery

Action

On the application system and the backup system, set the values of the `omnirc` options `ZDB_DELAY_BEFORE_RESCAN` and `ZDB_DELAY_AFTER_RESCAN` to `300` and restart the instant recovery session.

Problem

Instant recovery of an Oracle database fails

An instant recovery session for an Oracle database fails with a message similar to the following:

```
[Normal] From: ob2rman@x64-node1.x64ring.com "testdb" Time:
2/7/2008 10:48:19 AM
Starting target database instant recovery.
```

```
Net service name: testdb.
Instance status: .
Instance name: .
Database DBID = .
Database control file type: .
Database log mode: .
```

```
[Major] From: ob2rman@x64-node1.x64ring.com "testdb" Time:
2/7/2008 10:48:20 AM
The database reported error while performing requested operation.
```

Note that the database parameters are empty, which happens when Data Protector is not able to connect to the Oracle database.

Action

Data Protector must be able to connect to the Oracle database even when the database is in the Mount or Nomount state. One way of achieving this is by configuring static service information for your Oracle listener. For details, see the Oracle documentation.

The following example shows how static service information is configured using the Oracle Net Manager.

Suppose you have the following environment:

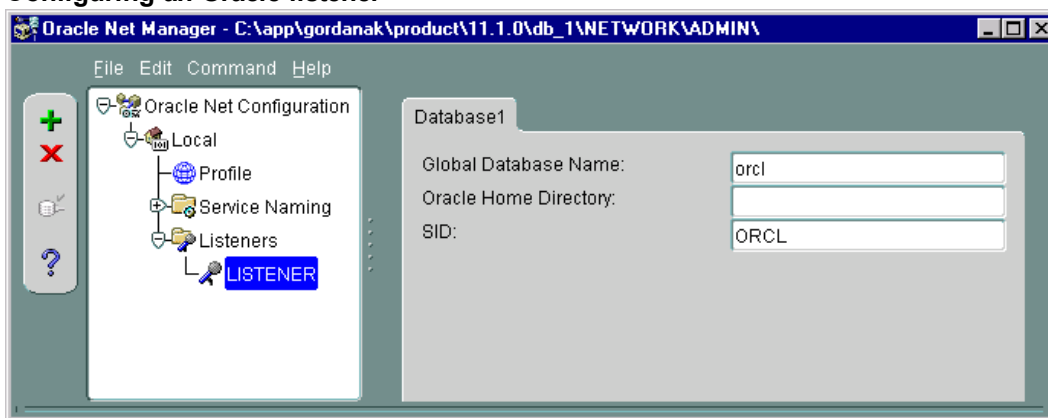
```
Listener name: LISTENER
```

Global database name: orcl

Oracle SID: ORCL

To configure static service information for the listener, open the Oracle Net Manager, select the listener, go to the **Database Services** context, add a database and specify the Oracle database parameters.

Configuring an Oracle listener



As a result, the `listener.ora` file gets updated with the `SID_LIST_LISTENER` section.

```
SID_LIST_LISTENER =  
  (SID_LIST =  
    (SID_DESC =  
      (GLOBAL_DBNAME = orcl)  
      (SID_NAME = ORCL)  
    )  
  )  
)
```

At the end, restart the Oracle listener service to apply the changes.

Problem

“IPC Invalid Hostname or IP Address” error message is displayed when browsing Oracle restore sessions

The following error message is displayed when browsing the Oracle database for restore sessions in the Data Protector GUI Restore context:

IPC Invalid Hostname or IP Address

The problem can appear in the following cases:

- When restoring database items to a different client.
- When importing Data Protector media containing backups of Oracle database from another Data Protector cell.
- When restoring 64-bit Oracle version 10.2.0.4 in the RAC environment, on HP-UX 11.23 PA-RISC systems. If `util_orarest` is present on the system, this error can mean that the `util_orarest` agent ends abnormally while trying to load the 32-bit OCI library from the `ORACLE_HOME/lib32` directory.

Actions

- To successfully restore database items to a different client, make sure that the system on which the Data Protector Oracle integration agent will be started is configured as Data Protector Oracle database instance (*ORACLE_SID*).

To verify, check if it is listed in the **Client** drop-down list in the **Options** page.

Select the system and proceed with [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., on page 81 of the Restoring Oracle database objects , on page 80 procedure.](#)

- When restoring 64-bit Oracle database in RAC environment, on HP-UX 11.23, resolve the problem as follows:

In the directory *ORACLE_HOME/lib* remove the soft link *libclntsh.sl*, which points to the 64-bit OCI library *ORACLE_HOME/lib/libclntsh.sl.10.1*.

Problem

When editing the RMAN script section of a backup specification using the Data Protector GUI, an RMAN backup script error is displayed

In the Data Protector GUI, when you edit the RMAN script section of a Data Protector backup specification, the following error message may display:

```
Cannot proceed, invalid RMAN backup script.
```

The error is displayed if you have specified an Oracle RMAN parameter, which has not been recognized by the Data Protector parser or a parsing error has occurred.

Action

Disable Oracle RMAN script parsing in the Data Protector GUI by setting the Data Protector `NoGUIRMANScriptParsing` global option to 1.

For details of how to set the option, see the *HPE Data Protector Help* index: "global options".

Problem

The Oracle Recovery Manager fails with "fatal error in recovery manager"

The Oracle RMAN fails with the following error messages:

```
RMAN-06900: WARNING: unable to generate V$RMAN_STATUS or V$RMAN_OUTPUT row
RMAN-06901: WARNING: disabling update of the V$RMAN_STATUS and V$RMAN_OUTPUT rows
RMAN-06003: ORACLE error from target database:
ORA-03113: end-of-file on communication channel
Process ID: 4222
Session ID: 29 Serial number: 51
..
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
```

RMAN-00601: fatal error in recovery manager

RMAN-03004: fatal error during execution of command

Action

In the `sqlnet.ora` file, increase the `SQLNET.RECV_TIMEOUT` value.

Chapter 2: Data Protector SAP R/3 ZDB integration

Introduction

This chapter explains how to configure and use the Data Protector SAP R/3 ZDB integration (**SAP R/3 ZDB integration**). It describes concepts and methods you need to understand to back up and restore the following files of the SAP R/3 database environment (**SAP R/3 objects**):

- Data files
- Control files
- Online redo logs
- Offline (archived) redo logs
- SAP R/3 logs and parameter files

Data Protector supports offline and online backups. During an online backup, the SAP R/3 application is actively used.

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

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

Data Protector supports only a filesystem restore. You can restore SAP R/3 files:

- To the original location
- To another client
- To another directory

[Backup and restore sessions, below](#) shows which restore methods are available, depending on the ZDB session you restore from.

Backup and restore sessions

ZDB type	Restore methods
ZDB to tape	Standard restore
ZDB to disk	Instant recovery
ZDB to disk+tape	Standard restore, instant recovery

When the instant recovery completes, you can recover the database to a specific point in time using the SAP BRTOOLS interface.

Supported disk arrays and array configurations


Supported array	Supported configurations
EMC Symmetrix (EMC)	TimeFinder, SRDF, combined SRDF+TimeFinder
HPE P9000 XP Disk Array Family (P9000 XP Array)	HPE BC P9000 XP, HPE CA P9000 XP, combined HPE CA+BC P9000 XP
HPE P6000 EVA Disk Array Family (P6000 EVA Array)	HPE BC P6000 EVA, combined HPE CA+BC P6000 EVA
Non-HPE Storage Arrays (NetApp storage)	Local replication

This chapter provides information specific to the Data Protector SAP R/3 ZDB integration. For general Data Protector procedures and options, see the *HPE Data Protector Help*. For details on ZDB terminology, ZDB types, advantages of offline and online backups, and instant recovery concepts, see the *HPE Data Protector Concepts Guide*.

Integration concepts

SAP R/3 integration architecture, below shows the architecture of the Data Protector SAP R/3 ZDB integration. The figure illustrates the preferred configuration, in which the Oracle control file, online redo log files, and Oracle SPFILE reside on a different volume group than the Oracle data files.

SAP R/3 integration architecture

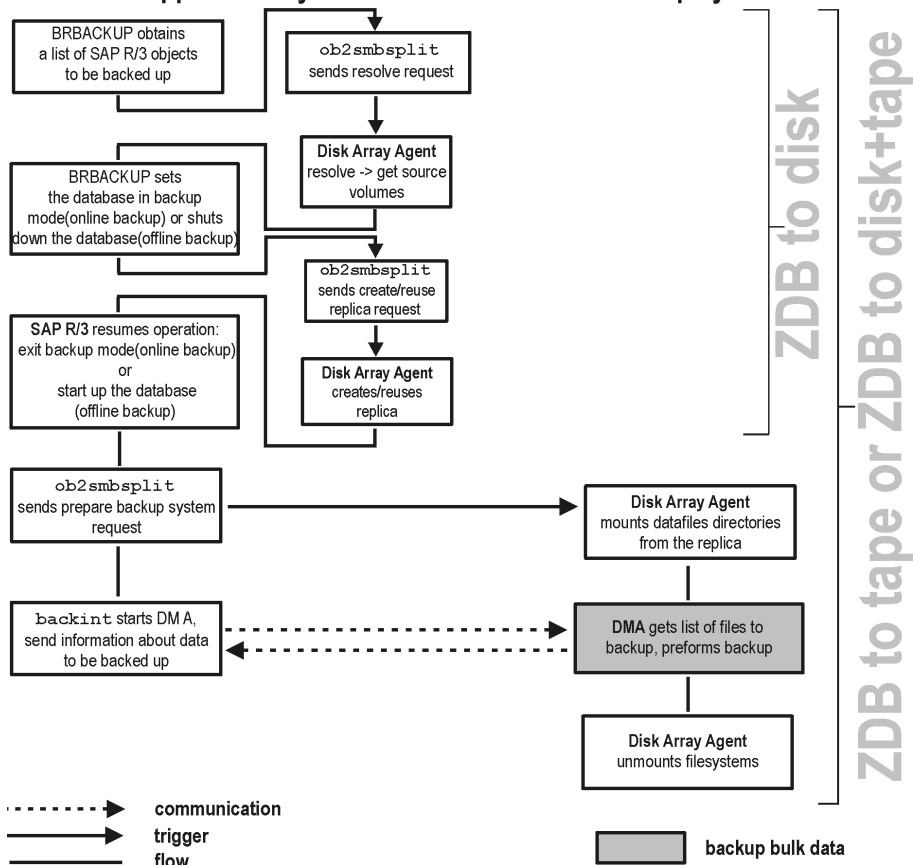


For other supported configurations, see [ZDB integrations omnirc options, on page 409](#).

ZDB flow

For details of how ZDB options affect the ZDB flow, including mounting, activating volume or disk groups and so on, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

SAP R/3 ZDB session flow (BRBACKUP started on the application system)



1. The SAP R/3 backup specification is read and the Data Protector `omnisap.exe` program is started on the application system.
2. `omnisap.exe` prepares an exclude list (Oracle control files and online redo log files)
3. `omnisap.exe` starts BRBACKUP, which switches the database to the backup mode (online backup) or shuts down the database (offline backup), and starts the split command on the application system, with the list of files to be included in the replica creation.

The database is put out of the backup mode (online backup) or is restarted (offline backup) after the replica is created.

NOTE:

Data Protector can use the `splitint` interface (if BRTOOLS supports it) to reduce the time during which the database is in the backup mode.

4. The Data Protector `ob2smbsplit` command resolves the backup configuration, creates a replica,

and preparing the replica for backup.

The mountpoints for the backed up object are created on the backup system.

The backup volume/disk groups are activated and the filesystems are mounted on the backup system.

5. **ZDB to disk:** The remaining ZDB options are processed and the details on the session are written to the ZDB database. The session then finishes.
6. **ZDB to tape, ZDB to disk+tape:** BRBACKUP starts the Data Protector `backint` program, which starts establishing a connection between the Data Protector Data Movement Agents (DMA) on the backup system and the General Media Agents (MA). The process is coordinated by the Data Protector Backup Session Manager (BSM). When the connection is established, the data specified for backup is streamed to tape.

NOTE:

You can configure SAP to use a third-party `backint` tool to perform ZDB-to-disk+tape backup sessions. To enable this feature, set the `OB2_3RD_PARTY_BACKINT` environment variable to 1 and copy the `backint` you want to use to the SAP `BRTOOLS` directory. In case of a ZDB-to-disk+tape backup session, a `disk_only` backup object is created in the Data Protector IDB. The third-party `backint` starts after the split and is responsible for backing up the needed files.

7. When the data transfer completes, the backup system is disabled (filesystems are dismounted on all platforms and volume/disk groups deactivated on UNIX).
8. **EMC and P9000 XP Array:** Links are re-established, depending on how ZDB options are specified.

Data Protector SAP R/3 configuration file

Data Protector stores the integration parameters for every configured SAP R/3 database in the following file on the Cell Manager:

Windows systems: `Data_Protector_program_data\Config\Server\Integ\Config\Sap\ClientName%ORACLE_SID`

UNIX systems: `/etc/opt/omni/server/integ/config/SAP/ClientName%ORACLE_SID`

The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- SAPDATA home directory
- user name and user group
- temporary directory used for the copy of the control file or redo logs
- list of control files and redo logs that will be copied to a safe location

- concurrency number and balancing (for each backup specification), and number of channels for RMAN backup
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

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

IMPORTANT:

To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

NOTE:

You can set up the parameters in the `Environment` section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='ORACLE_HOME';  
ConnStr='ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE';  
BR_directory='BRTOOLS_HOME';  
SAPDATA_HOME='SAPDATA_HOME';  
ORA_NLS_CHARACTERSET='CHARACTER_SET';  
OSUSER='USER_NAME';  
OSGROUP='USER_GROUP';  
Environment={  
  [ENV_var1='value1'];  
  [ENV_var2='value2'];  
  ...]  
}  
SAP_Parameters={backup_spec_name=(' -concurrency #_of_concurrency  
' | '-time_balance' | '-load_balance' | '-manual_balance');  
}  
speed={  
  AVERAGE=1;  
  'filename'=#_of_seconds_needed_to_back_up_this_file;  
}  
compression={'filename'=size_of_the_file_in_bytes_after_the  
_compression;  
}  
manual_balance={backup_specification_name={  
  'filename'=device_number;  
  }  
}
```

The `ORA_NLS_CHARACTERSET` parameter is set automatically by Data Protector during SAP R/3 database configuration. For details of how to configure SAP R/3 database for use with Data Protector, see [Configuring SAP R/3 databases, on page 140](#).

Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBBHBB
QDBBOFBBCFBPFBCFBIFBFBGFBBDGBBBFBCFBDFBFCFBB';
BR_directory='/usr/sap/ABA/SYS/exe/run';
SAPDATA_HOME='/sap';
ORA_NLS_CHARACTERSET='USASCII7';
OSUSER='oraid';
OSGROUP='dba';

Environment={
  SAP_Parameters={
    sap_weekly_offline=('-concurrency 1','-no_balance');
    sap_daily_online=('-concurrency 3','-load_balance');
    sap_daily_manual=('-concurrency 3','-manual_balance');
  }
  speed={
    AVERAGE=203971;
    '/file1'=138186;
    '/file2'=269756;
  }
  compression={
    '/file1'=1234;
    '/file2'=5678;
  }
  manual_balance={
    sap_daily_manual={
    '/file1'=1; /* file 1 is backed up by the first sapback */
    '/file2'=2; /* file 2 is backed up by the second sapback */
    '/file3'=1; /* file 3 is backed up by the first sapback */
    '/file4'=1;
    }
  }
}
```

Setting, retrieving, listing, and deleting Data Protector SAP R/3 configuration file parameters using the CLI

The Data Protector SAP R/3 configuration file parameters are normally written to the Data Protector SAP R/3 configuration file after:

- the Data Protector configuration of the Oracle instance that is run by SAP R/3 is completed.
- a new backup specification is created.

- a backup that uses balancing by time algorithm is completed.

The `util_cmd` command

You can set, retrieve, list, or delete the Data Protector SAP R/3 configuration file parameters using the `util_cmd -putopt` (setting a parameter), `util_cmd -getopt` (retrieving a parameter), or `util_cmd -getconf` (listing all parameters) command on the Data Protector SAP R/3 client.

Cluster-aware clients

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

Windows systems: `set OB2BARHOSTNAME=virtual_hostname`

UNIX systems: `export OB2BARHOSTNAME=virtual_hostname`

The `util_cmd` synopsis

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

```
util_cmd -getconf[ig] SAP oracle_instance [-local filename]
```

```
util_cmd -getopt[ion] [SAP oracle_instance] option_name [-sub[list] sublist_name] [-local filename]
```

```
util_cmd -putopt[ion] [SAP oracle_instance] option_name [option_value] [-sub[list] sublist_name] [-local filename]
```

where:

option_name is the name of the parameter

option_value is the value for the parameter

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

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

- When it is used with the `-getconf[ig]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the standard output.
- When it is used with the `-getopt[ion]`, it specifies the filename of the file from which the parameter and its value are to be taken and then written to the standard output. If the `-local` option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output.
- When it is used with the `-putopt[ion]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the Data Protector SAP R/3 configuration file.

NOTE:

If you are setting the *option_value* parameter as a number, the number must be put in single quotes, surrounded by double quotes.

Return values

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

- Configuration read/write operation successful.
This message is displayed when all the requested operations have been completed successfully.
- Configuration option/file not found.
This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.
- Configuration read/write operation failed.
This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, and so on.

Setting parameters

To set the Data Protector `OB2OPTS` and the Oracle `BR_TRACE` parameters for the Oracle instance `ICE` that is run by SAP R/3, use the following commands on the Data Protector SAP R/3 client:

Windows, HP-UX, Solaris systems

```
util_cmd -putopt SAP ICE OB2OPTS '-debug 1-200 debug.txt' -sublist Environment
util_cmd -putopt SAP ICE BR_TRACE "'10'" -sublist Environment
```

Retrieving parameters

To retrieve the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -getopt SAP ICE OB2OPTS -sublist Environment
```

Listing parameters

To list all the Data Protector SAP R/3 configuration file parameters for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -getconf SAP ICE
```

Deleting parameters

To remove the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

```
util_cmd -putopt SAP ICE OB2OPTS "" -sublist Environment
```

Configuring the integration

To configure the integration:

1. Configure the required user accounts. See [Configuring user accounts, on page 134](#).
2. Configure SQL*Net V2 or Net8 TNS listener. See [Configuring SQL*Net V2 or Net8 TNS listener, on page 135](#).

3. Check the connection to the Oracle database from the application system. See [Checking the connection, on page 136](#).
4. Enable the use of the authentication password file. See [Authentication password file, on page 137](#).
5. Optionally, set the archived logging mode to enable online backups. See [Enabling archived logging, on page 137](#).
6. Share directories on the application system. See [Sharing directories on the application system, on page 138](#).
7. Configure every SAP R/3 database you intend to back up from or restore to. See [Configuring SAP R/3 databases, on page 140](#).
8. Configure the SAP R/3 parameter file. See [Configuring the SAP R/3 parameter file, on page 147](#).

Prerequisites

- Ensure that you have correctly installed and configured the SAP R/3 application. The database used by the SAP R/3 application must be an Oracle database. If any other database is used, you can back it up using the corresponding Data Protector integration. It is assumed that you are familiar with the SAP R/3 application and Oracle database administration.
 - For supported versions, platforms, devices, and other information, see the latest support matrices at <https://softwaresupport.hpe.com/manuals>.
 - For information on installing, configuring, and using the SAP R/3 application and the SAP backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE), see the SAP R/3 application documentation.
- Ensure that you have a license to use the Data Protector SAP R/3 ZDB integration. For information, see the *HPE Data Protector Installation Guide*.
- Ensure that you have correctly installed Data Protector.
 - For information on how to install the Data Protector disk array integration (P6000 EVA Array, P9000 XP Array, EMC, or non-HPE Storage Arrays) with SAP R/3 in various architectures, see the *HPE Data Protector Installation Guide*.
 - On how to configure the Data Protector ZDB integration (EMC, P9000 XP Array, P6000 EVA Array, or NetApp Storage), see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.
 - For information on the Data Protector Cell Manager package configuration in the HPE SG cluster, see the *HPE Data Protector Help* index: "HPE Serviceguard integration".

NOTE:

You cannot run ZDB sessions in the RMAN mode.

- The SAP R/3 directories SAPBACKUP, SAPARCH, SAPREORG, SAPCHECK, and SAPTRACE must not reside on the same disk array source volumes as the data files. Otherwise, the BRTOOLS data needed for complete recovery of a database is overwritten during instant recovery. You can set the locations for these directories in the `initDBSID.sap` file.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP R/3 system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore.
- **Windows systems:**

On supported Windows operating systems, configure the Data Protector Inet service user impersonation for the user that has the appropriate SAP R/3 permissions for running backups and restores.

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

If there are several SAP R/3 instances running on the same system with different SAP administrator accounts configured for each instance, create an additional, common SAP administrator account. Configure the Data Protector Inet service to use this account as the service startup account.

Cluster-aware clients

- Configure SAP R/3 databases only on one cluster node, since the configuration files reside on the Cell Manager.

Windows systems: During the configuration, Data Protector copies the Data Protector `backint` and `ob2smbsplit.exe` programs (the latter only if `splitint` is supported by BRTOOLS) from `Data_Protector_home\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`. This is done only on the currently active node. On the other node, do it manually.

UNIX systems: During the configuration, Data Protector creates a link to the Data Protector `backint` and `splitint` programs on the currently active node. On all the other nodes, do it manually. Execute:

```
ln -s /opt/omni/1bin/backint \ /usr/sap/ORACLE_SID/sys/exe/run
```

If `splitint` is supported by BRTOOLS, also execute:

```
ln -s /opt/omni/1bin/ob2smbsplit \ /usr/sap/ORACLE_SID/sys/exe/run/splitint
```

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

Windows systems: set `OB2BARHOSTNAME=virtual_server_name`

UNIX systems: export `OB2BARHOSTNAME=virtual_server_name`

NOTE:

SAP recommends installing SAP backup utilities on all cluster nodes.

Configuring user accounts

To enable backup and restore of SAP R/3 database files, you need to configure or create several user accounts.

Oracle operating system user account	Operating system user account that is added to the following user groups: Windows systems: ORA_DBA and ORA_SID_DBA local groups UNIX systems: dba and sapsys For example, user oraSID. UNIX systems: Ensure that this user is the owner of the filesystem or of the raw logical volume on which the database is mounted. The minimum permissions should be 740.
User account root (UNIX systems only)	Default operating system administrator's user account added to the dba user group.
Oracle database user account	Database user account granted at least the following Oracle roles: <ul style="list-style-type: none"> • sysdba • sysoper For example, user system. Do not configure the Oracle SYS user for backing up SAP R/3 objects. When backing up using the SYS user account, the SAP backup fails with the error ORA-28009: connection as SYS should be as SYSDBA or SYSOPER.

Add the following user accounts to the Data Protector admin or operator user group:

- Oracle operating system user account
 (if you are using backup set method, add this user on the application as well as on backup system)
- **UNIX systems:** User account root (for both the application system and backup system)

In cluster environments, add these user accounts to the Data Protector admin or operator user group for the following clients:

- virtual server
- every node in the cluster

For information on adding Data Protector users, see the *HPE Data Protector Help* index: "adding users".

Configuring SQL*Net V2 or Net8 TNS listener

1. Ensure that the listener.ora and tnsnames.ora files on the application system are configured as shown in the following example. The files are located in:

UNIX systems: ORACLE_HOME/network/admin

Windows systems: ORACLE_HOME\network\admin

Example

Oracle instance: PRO

Application system: alpha.hp.com

listener.ora	<pre>LISTENER = (DESCRIPTION_LIST = (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP) (HOST = alpha.hp.com) (PORT = 1522)))) SID_LIST_LISTENER = (SID_LIST = (SID_DESC = (GLOBAL_DBNAME = PRO) (SID_NAME = PRO) (ORACLE_HOME = /app/oracle815/product)))</pre>
tnsnames.ora	<pre>PRO = (DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP) (HOST = alpha.hp.com) (PORT = 1522))) (CONNECT_DATA = (SERVICE_NAME = PRO)))</pre>

2. Start the SQL*Net V2 or Net8 TNS listener by executing the following on the Oracle Server system:

UNIX systems: `ORACLE_HOME/bin/lsnrctl start`

Windows systems: `ORACLE_HOME\bin\lsnrctl start`

Checking the connection

To check the connection to the Oracle instance from the application system:

1. Log in to the application system as the Oracle OS user.
2. Export/set the `ORACLE_HOME` and `ORACLE_SID` variables.
3. Start `sqlplus`.
4. Connect to the Oracle target database as the Oracle database user, first with the `sysdba` role and then with the `sysoper` role.

Example

For the following configuration:

Oracle instance: `PROORACLE_HOME: /app/oracle816/product`

execute:

```
id
uid=102(oracle) gid=101(dba)
export ORACLE_SID=PRO
export ORACLE_HOME=/app/oracle816/product
export SHLIB_PATH=/app/oracle816/product/lib:/opt/omni/lib
sqlplus /nolog
SQLPLUS> connect system/manager@PRO as sysdba;
Connected.
SQLPLUS> connect system/manager@PRO as sysoper;
Connected.
```

Authentication password file

Enable the use of the authentication password file for the database administrator:

1. Shut down the Oracle target database on the application system.
2. In the `initORACLE_SID.ora` file, specify:

```
remote_login_passwordfile = exclusive
```

For instructions on how to set up the password file, see the Oracle documentation.

Enabling archived logging

When you set the database to the archived logging mode, you protect the unsaved online redo logs from being overwritten. Online backup of data files is useless without the related redo logs because you cannot recover the database to a consistent state.

TIP:

Archive the redo log files generated during the online backup immediately after BRBACKUP completes.

To protect the archive directory from overflowing, clear the directory regularly.

To enable archived logging:

1. In the `initORACLE_SID.ora` file, set
`log_archive_start = true`
and specify the `log_archive_dest` option.

Example

This is an example of the `initORACLE_SID.ora` file for the Oracle instance PRO:

```
# @(#)initSID.ora 20.4.6.1 SAP 13/03/30
#####
# (c)Copyright SAP AG, Walldorf
#####
. . . .
. . . . . . . . . .
```

```
. . . . .  
. . . . .  
### ORACLE Authentication Password File  
remote_login_passwordfile = exclusive  
### ORACLE archiving  
log_archive_dest = /oracle/PRO/saparch/PROarch  
log_archive_start = true  
. . . . .
```

2. Mount the Oracle database and start the archived logging mode using the Oracle Server Manager.
Execute:

```
startup mount  
alter database archivelog;  
archive log start;  
alter database open;
```

Example

For the Oracle instance PRO, execute:

Windows systems: set ORACLE_SID=PRO

UNIX systems: export ORACLE_SID=PRO

Any operating system:

```
sqlplus /nolog  
SQLPLUS> connect user/passwd@PRO;  
Connected.  
SQLPLUS> startup mount  
ORACLE instance started.  
Total System Global Area          6060224 bytes  
Fixed Size                        47296 bytes  
Variable Size                     4292608 bytes  
Database Buffers                  1638400 bytes  
Redo Buffers                       81920 bytes  
Database mounted.  
SQLPLUS> alter database archivelog;  
Statement processed.  
SQLPLUS> archive log start;  
Statement processed.  
SQLPLUS> alter database open;
```

Sharing directories on the application system

The following directories on the application system must be accessible:

- sapbackup
- sapreorg
- Oracle home directory
- BR*Tools home directory

NOTE:

The `sapreorg` and `BR*Tools` home directories must be accessible only if you want to run SAP compliant ZDB sessions (`BRBACKUP` is started on the backup system and not on the application system).

UNIX application system

1. Share the directories on the application system through NFS with root permissions.

For example, suppose that the `sapbackup` directory on the application system points to `/oracle/SID/sapbackup` and the backup system is `backup.company.com`. To share the `sapbackup` directory:

HP-UX systems: In the file `/etc/exports` on the application system, add the line:

```
/oracle/SID/sapbackup -root=backup.company.com
```

Solaris systems: In the file `/etc/dfs/dfstab` on the application system, add the line:

```
share -F nfs -o root=backup.company.com/oracle/SID/sapbackup
```

2. Mount the directories on the backup system. Ensure that you have the same directory structure on both the application and backup system.

For example, suppose that you have an HP-UX application system `app.company.com`. To mount the `/oracle/SID/sapbackup` directory on the backup system, add the following line to the file `/etc/fstab` on the backup system:

```
app.company.com:/oracle/SID/sapbackup  
/oracle/SID/sapbackup nfs defaults 0 0
```

Windows application system

- On the application system, find the location of the `sapbackup` and `sapreorg` directories. If they reside inside the SAP data home directory, share this directory under any name you want. Then, on the backup system, create the `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\ORACLE_SID\Environment\SAPDATA_HOME` Windows registry key, specifying the SAP data home directory path as seen from the backup system.

For example, suppose your application system is `mycomputer.company.com` and the SAP data home is `K:\oracle\my_instance`, which you share under the name `my_SAPinstance`. Then, the `SAPDATA_HOME` Windows registry key on the backup system must have the value `\\mycomputer.company.com\my_SAPinstance`.

If the `sapbackup` and `sapreorg` directories do not reside together, share each directory separately and also create a separate Windows registry key on the backup system (`SAPBACKUP`, `SAPREORG`).

- Share the Oracle home directory on the application system and specify its path in the `ORACLE_HOME` Windows registry key on the backup system, similarly as described above.
- Ensure that the `BR*Tools` home directory on the application system is accessible from the backup system as follows:

```
\\application_system\sapmnt\SAP_SID\SYS\exe\run
```

TIP:

Instead of creating registry keys, you can also set the Data Protector SAP R/3 integration environment variables (ORACLE_HOME, SAPDATA_HOME, SAPBACKUP, SAPREORG).

Choosing authentication mode

Data Protector SAP R/3 ZDB integration supports two authentication modes for accessing Oracle databases that are used by SAP R/3:

- database authentication mode
- operating system authentication mode

With database authentication mode, you need to re-configure the SAP R/3 integration for an SAP R/3 database with the new Oracle login information each time the corresponding Oracle database user account changes. Such a reconfiguration is not needed if operating system authentication mode is used.

You select the preferred authentication mode when you configure a particular SAP R/3 database.

Configuring SAP R/3 databases

You need to provide Data Protector with the following configuration parameters:

- Oracle Server home directory
- SAP R/3 data home directory
- Optionally, if you choose database authentication mode, Oracle database user account. The user account is used by BRBACKUP and BRARCHIVE during backup.
- Directory in which the SAP backup utilities are stored

Data Protector then creates the configuration file for the SAP R/3 database on the Cell Manager and verifies the connection to the database. On UNIX systems, Data Protector also creates a soft link for the `backint` program from the directory that stores the SAP backup utilities to `/opt/omni/1bin`.

On Windows systems, Data Protector copies the `backint` and `programs` (the latter only if `splitint` is supported by BR*Tools) from `Data_Protector_home\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`.

To configure an SAP R/3 database, use the Data Protector GUI or CLI.

Before you begin

- Ensure that the SAP R/3 database is open.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the template.

From the **Backup type** drop-down list, select the **Snapshot or split mirror backup** option, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system.

Click **OK**.

4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

In **Backup system**, select the backup system.

Specify other disk array specific backup options (see [EMC backup options, on page 151](#) for EMC, [P9000 XP Array backup options, on page 151](#) for P9000 XP Array, [P6000 EVA Array backup options, on page 152](#) for P6000 EVA Array, or [NetApp Storage backup options, on page 153](#) for NetApp Storage). For details on the options, press **F1**.

NOTE:

P9000 XP Array and P6000 EVA Array: To enable instant recovery, select **Track the replica for instant recovery**.

5. In **Application database**, type the Oracle instance name (ORACLE_SID).

Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 clients, as follows:


Windows Server 2008: In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

UNIX systems: In **Username**, type the Oracle OS user described in [Configuring user accounts, on page 134](#). In **Group/Domain name**, type dba.

Ensure that this user has been added to the Data Protector admin or operator user group, has the SAP R/3 backup rights, and has been set up for the Data Protector Inet service user impersonation. This user becomes the backup owner.

For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: "Inet user impersonation".

Specifying an SAP R/3 system and Oracle instance



Click **Next**.

6. In the **Configure SAP** dialog box, specify the pathname of the Oracle Server home directory and SAP R/3 data home directory. If you leave the fields empty, the default `ORACLE_HOME` directory is used.

Under **Oracle login information to target database**, specify the following:

- For the database authentication mode, specify **Username**, **Password**, and **Service**.
- For the local operating system authentication mode, leave **Username**, **Password**, and **Service** empty.
- For the remote operating system authentication mode, specify only **Service** (leave **Username** and **Password** empty).

The following are the option descriptions:

- **Username** and **Password**: Specify the user name and password of the Oracle database user account described in [Configuring user accounts, on page 134](#).
- **Service** : Specify the Oracle service name.

In **Backup and restore executables directory**, specify the pathname of the directory in which the SAP backup utilities reside. By default, the utilities reside in:

Windows systems: `\\SAP_system\sapmnt\ORACLE_SID\sys\exe\run`

UNIX systems: `/usr/sap/ORACLE_SID/SYS/exe/run`


Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)

The screenshot shows a 'Configure SAP' dialog box with the following fields and values:

- Client:** vasgora.company.com
- Oracle SID:** AIR
- Oracle Server home directory:** /app/oracle/product/81
- SAP Data home directory:** /app/oracle/product/AIR
- Oracle login information to target database:**
 - Username:** (empty)
 - Password:** (empty)
 - Service:** AIR
- Backup and restore executables directory:** /usr/sap/CER/SYS/exe/run

Buttons at the bottom: OK, Cancel, Help.

Configuring an SAP R/3 database on a Windows system (database authentication mode)



Click **OK**.

7. The SAP R/3 database is configured. Exit the GUI or proceed with creating the backup specification at [Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs., on page 153.](#)

Using the Data Protector CLI

1. Log in to the SAP R/3 system using the Oracle operating system user account.
2. At the command prompt, change current directory to the following directory:

Windows systems: `Data_Protector_home\bin`

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

3. Execute:

```
util_sap.exe -CONFIG ORACLE_SIDORACLE_HOMEtargetdb_connection_stringSAPTOOLS_
DIR [SAPDATA_HOME][SQL_PATH]
```

Parameter description

<code>ORACLE_SID</code>	Oracle instance name.
<code>ORACLE_HOME</code>	Pathname of the Oracle Server home directory.

<i>targetdb_connection_string</i>	This argument value determines the authentication mode used for accessing the Oracle database: <ul style="list-style-type: none">• To select the database authentication mode, specify the login information to the target database in the format <i>user_name/password@Oracle_service</i>.• To select the local operating system authentication mode, specify only the character <i>/</i>.• To select the remote operating system authentication mode, specify the login information to the target database in the format <i>/@Oracle_service</i>.
<i>SAPTOOLS_DIR</i>	Pathname of the directory that stores the SAP backup utilities.
<i>SAPDATA_HOME</i>	Pathname of the directory where the SAP R/3 data files are installed. By default, this parameter is set to <i>ORACLE_HOME</i> .

The message **RETVAl*0* indicates successful configuration.

Handling errors

If you receive the message **RETVAl*error_number* where *error_number* is different than zero, an error occurred.

To get the error description, execute:

Windows systems:

```
Data_Protector_home\bin\omnigetmsg 12 error_number
```

HP-UX systems:

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

TIP:

To get a list of Oracle instances that are used by the SAP R/3 application, execute:

```
util_sap.exe -APP
```

To get a list of tablespaces of an Oracle instance, execute:

```
util_sap.exe -OBJS0ORACLE_SID
```

To get a list of database files of a tablespace, execute:

```
util_sap.exe -OBJS1ORACLE_SID TABLESPACE
```

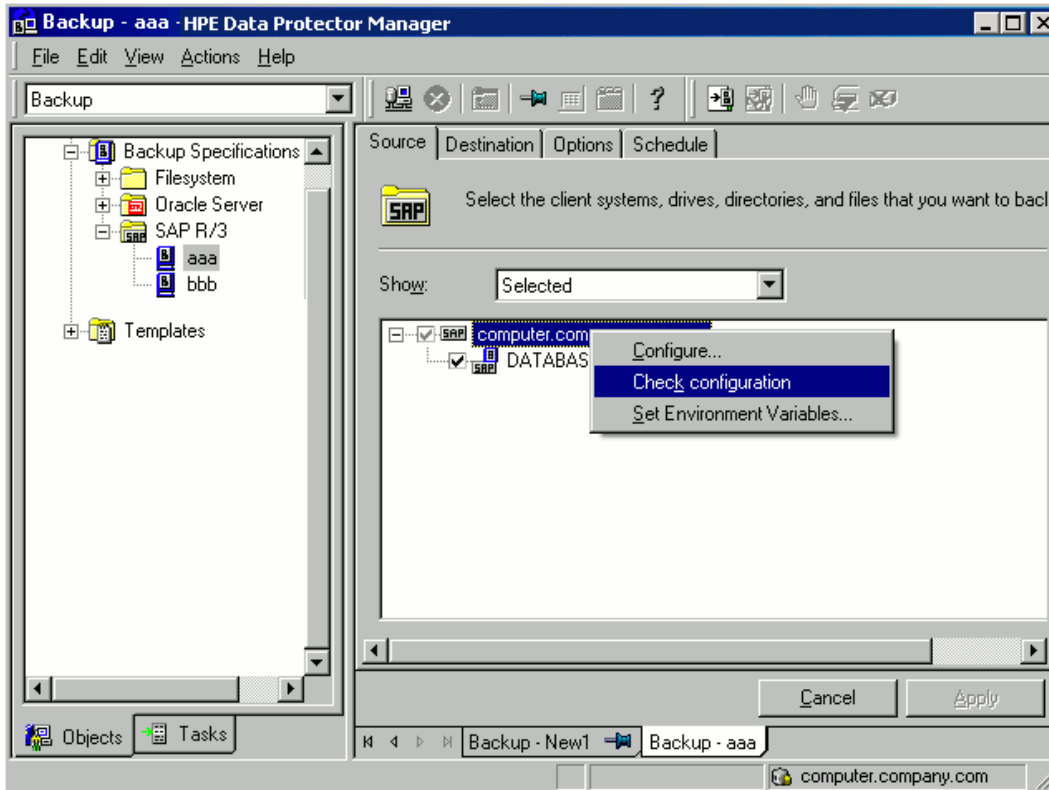
Checking the configuration

You can check the configuration of an SAP R/3 database after you have created at least one backup specification for this database. Use the Data Protector GUI or CLI.

Using the Data Protector GUI

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

Checking the SAP R/3 configuration



Using the Data Protector CLI

Log in to the SAP R/3 system as the Oracle OS user and execute:

```
util_sap.exe -CHKCONF ORACLE_SID
```

where *ORACLE_SID* is the name of the Oracle instance.

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

If you receive the message **RETVL*error_number* where *error_number* is different than zero, an error occurred. On how to get the error description, see [Handling errors, on the previous page](#).

To check if the SAP R/3 configuration is suitable for instant recovery, execute:

```
util_sap.exe -CHKCONF_IR ORACLE_SID [-verbose]
```

The `-verbose` option creates a file with a list of control files and redo log files that are on the same source volumes as the database files. If this list is not empty, a warning is displayed, stating that instant recovery is impossible.

Configuring the SAP R/3 parameter file

To configure the integration, you need to set some parameters in the SAP R/3 parameter file on both the application system and backup system. The file template is located in:

UNIX systems: `ORACLE_HOME/dbs/initORACLE_SID.sap`

Windows systems: `ORACLE_HOME\database\initORACLE_SID.sap`

SAP parameter file settings

Parameter	Value/Description
<code>split_cmd</code>	<p>On the application system:</p> <p>UNIX systems: <code>"/opt/omni/lbin/ob2smbsplit \$"</code></p> <p>Windows systems: <code>"Data_Protector_home\bin\ob2smbsplit \$"</code></p> <p>On the backup system, you do not need to set the parameter.</p> <p>BRBACKUP uses this parameter to trigger the replica creation. At run time, the optional sign "\$" is replaced with the name of the text file containing the names of files to be backed up.</p> <p>Windows systems: If the pathname contains spaces, use Windows short names instead.</p>
<code>primary_db</code>	<p>On the application system: LOCAL</p> <p>On the backup system: name of the service used for connecting to the Oracle database.</p> <p>This parameter defines the service name of the Oracle database to link the backup system to the application system.</p>

Backup

The integration provides online and offline database backups of the following types:

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

To configure a backup, create a ZDB backup specification.

Archived logs can only be backed up together with the database in a ZDB to disk+tape, ZDB to tape, or non-ZDB (standard backup) session. If you try to back up archived logs in a ZDB to disk session, or archived logs only in a ZDB session, the session fails.

What is backed up depends on your selection in the backup specification. For details, see [What is backed up, below](#).

What is backed up

Selected items	Backed up files
ARCHIVELOGS	<ul style="list-style-type: none">• offline (archived) redo logs• control files
DATABASE or individual tablespaces	<ul style="list-style-type: none">• data files• control files• SAP R/3 logs and parameter files• online redo logs (only during offline backups)

You can specify SAP R/3 backup options in two different ways:

- Using the BRBACKUP options.
- Using the SAP parameter file.

NOTE:

The BRBACKUP options override the settings in the SAP parameter file.

You can specify BRBACKUP options when you create a backup specification. If no options are specified, the SAP R/3 application refers to the current settings in the SAP parameter file. In such a case, before running a backup, ensure that the SAP parameter file is correctly configured.

TIP:

When you create a backup specification, select a backup template that already contains the desired BRBACKUP options.

Considerations

- Before you start a backup, ensure that the SAP R/3 database is in the open or shutdown mode.
- Before you start a backup, ensure the `primary_db` parameter is set to LOCAL in the SAP R/3 parameter file. For more information on setting the SAP R/3 parameter file, see [Configuring the SAP R/3 parameter file , on the previous page](#).
- ZDB, restore, and instant recovery sessions that use the same source volume on the application system cannot run simultaneously.
- You cannot start a ZDB to disk session if another session is backing up the archived logs, even if the Oracle data files and the archived logs reside on different source volumes.
- **P9000 XP Array:** If the LVM Mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their BC pairs assigned. The message should be ignored.

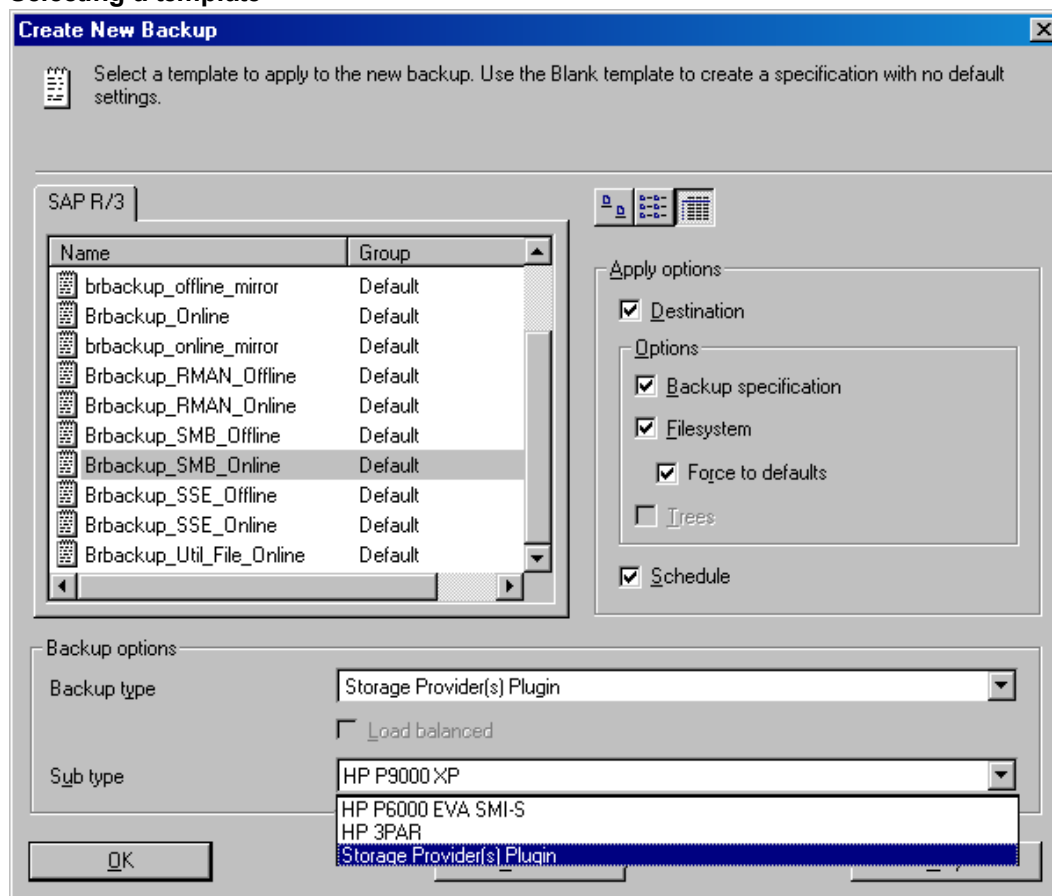
- **ZDB to disk:** Archived logs cannot be backed up. To back up archived logs, create a non-ZDB (standard) SAP R/3 backup specification. For information, see the *HPE Data Protector Integration Guide*.
- Configurable backup modes are supported only by using templates.
- By default, Data Protector supports all BRTOOL options, except -a and -b. To enable support for -a and -b, set the OB2BRTN0SECUomnirc option to 1. On how to set the omnirc options, see the *HPE Data Protector Help* index: “omnirc options”.

Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select a template (Selecting a template, below) and click **OK**.

Selecting a template



Backup templates available for zero downtime backup

Blank SAP Backup	No predefined options.
Brbackup_offline_mirror	Used for an offline ZDB (split mirror or snapshot backup) using <code>splitint</code> . The database is stopped during the creation of a replica. The database is offline for a shorter period of time than if <code>Brbackup_SMB_Offline</code> was used, but <code>splitint</code> must be supported by BRTOOLS.
Brbackup_online_mirror	Used for an online ZDB (split mirror or snapshot backup) using <code>splitint</code> . The database is active during the creation of a replica. The database is in the backup mode for a shorter period of time than if <code>Brbackup_SMB_Online</code> was used, but <code>splitint</code> must be supported by BRTOOLS.
Brbackup_SMB_Offline	Used for an offline ZDB (split mirror or snapshot backup). The database is stopped during the creation of a replica.
Brbackup_SMB_Online	Used for an online ZDB (split mirror or snapshot backup). The database is active during the creation of a replica.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**.

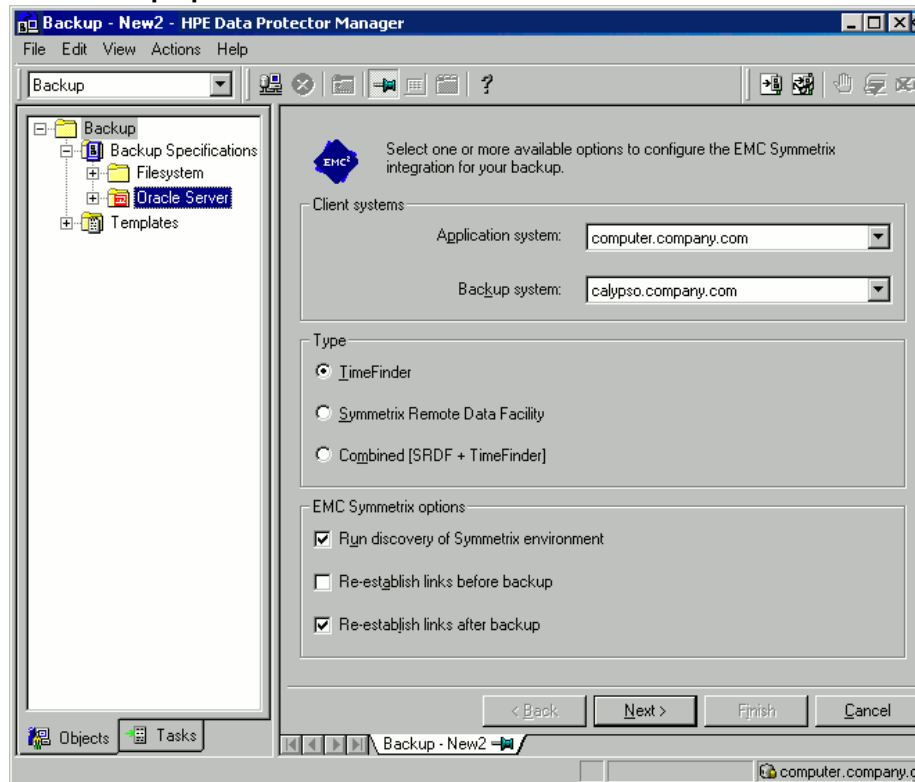
From the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system.

4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

In **Backup system**, select the backup system.

Select other disk array-specific backup options (see [EMC backup options, on the next page](#) for EMC, [P9000 XP Array backup options, on the next page](#) for P9000 XP Array, [P6000 EVA Array backup options, on page 152](#) for P6000 EVA Array, [NetApp Storage backup options, on page 153](#) for NetApp Storage). For detailed information on the backup options, press **F1**.


EMC backup options



P9000 XP Array specifics

To enable instant recovery, leave the **Track the replica for instant recovery** option selected. It is not possible to run instant recovery with Data Protector if this option is cleared.


P9000 XP Array backup options



P6000 EVA Array specifics

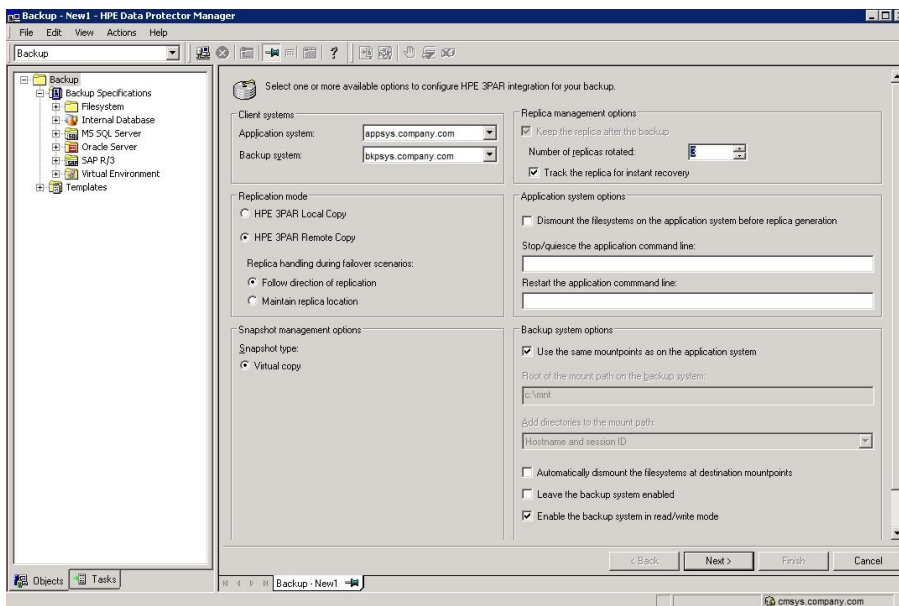
To enable instant recovery, select the **Track the replica for instant recovery** option.

P6000 EVA Array backup options

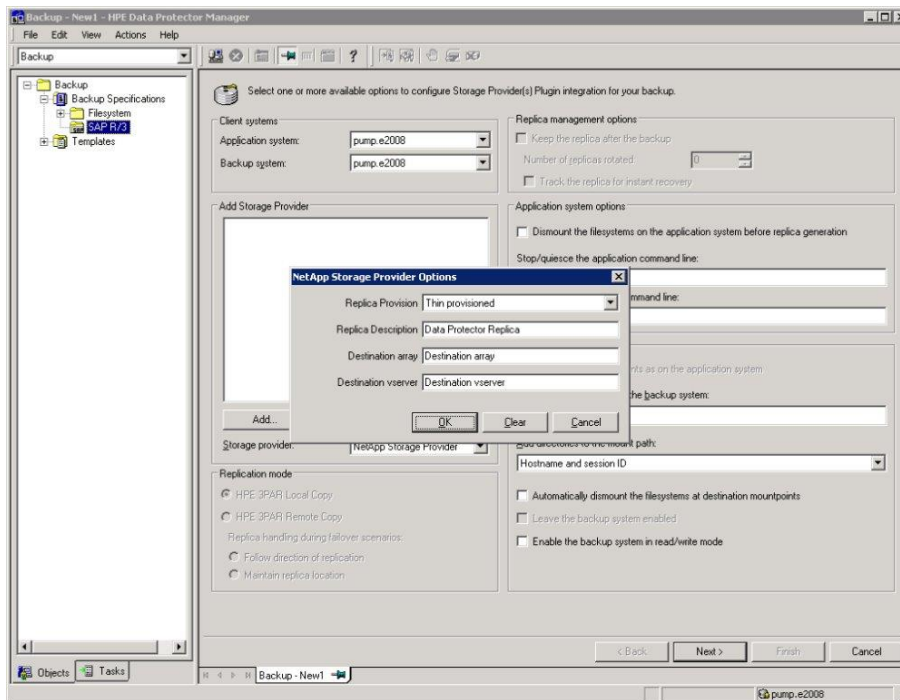


P10000 3PAR Array specifics

P10000 3PAR Array backup options



NetApp Storage backup options



Click **Next**.

5. In **Application database**, select the Oracle instance (ORACLE_SID) to be backed up. Specify the **User and group/domain** options, which are available on UNIX and Windows Server 2008 clients, as follows:

Windows Server 2008: In **Username** and **Group/Domain name**, specify the operating system user account under which you want the backup session to run (for example, the user name Administrator, domain DP).

UNIX systems: In **Username**, type the Oracle OS user described in [Configuring user accounts, on page 134](#). In **Group/Domain name**, type dba.

Ensure that this user has been added to the Data Protector admin or operator user group, has the SAP R/3 backup rights, and has been set up for the Data Protector Inet service user impersonation. This user becomes the backup owner.

For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: "Inet user impersonation".

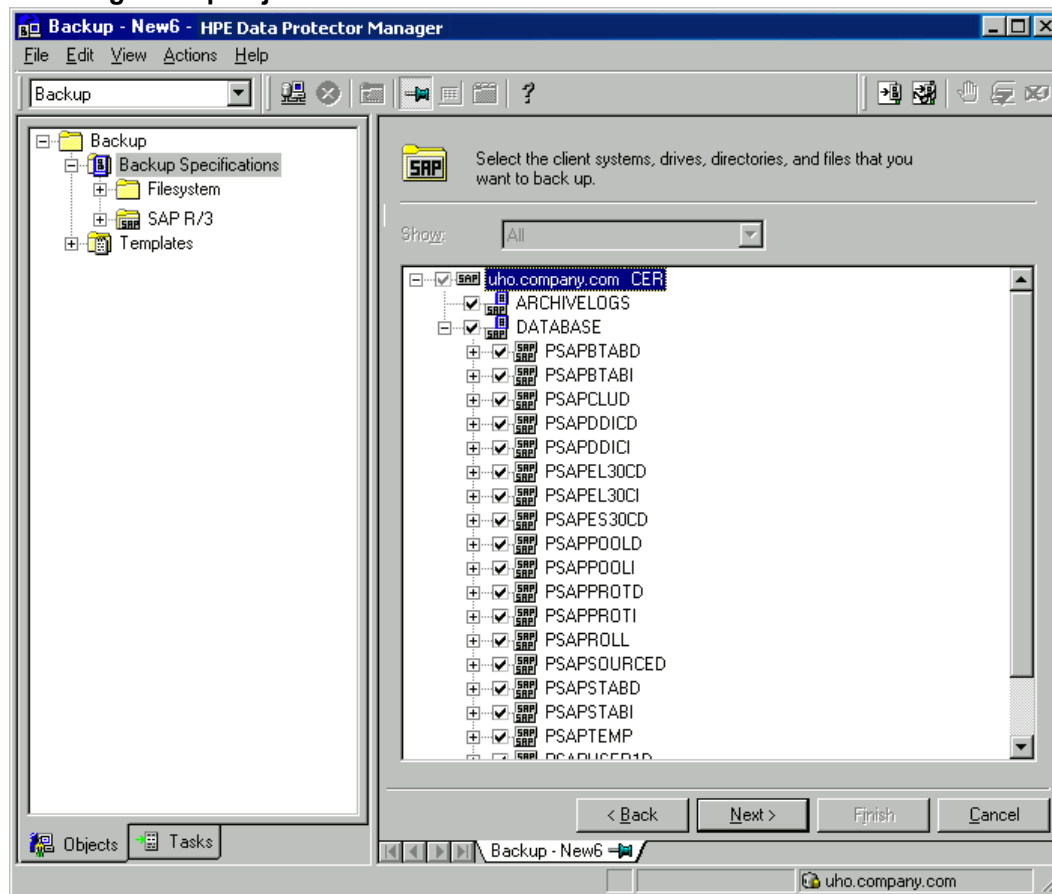
Click **Next**.

6. If the SAP R/3 database is not configured yet for use with Data Protector, the **Configure SAP** dialog box is displayed. Configure it as described in [Configuring SAP R/3 databases, on page 140](#).
7. Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs.

NOTE:

If you plan to do instant recovery, select the whole **DATABASE** item.

Selecting backup objects



Click **Next**.

8. Select devices to use for the backup.

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

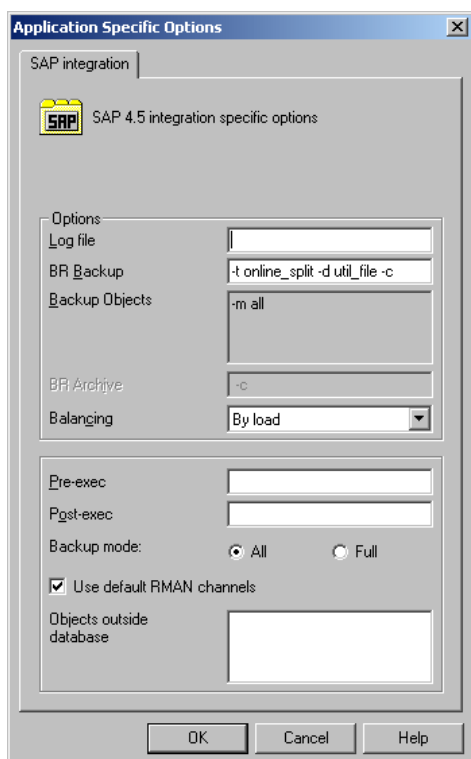
NOTE:

Parallelism (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the number of streams equals the sum of concurrencies of the selected devices.

Click **Next**.

9. Set backup options. For information on the application-specific options, see [SAP R/3 backup options, on the next page](#).

Application-specific options



Click **Next**.

- Click **Save As** to save the backup specification, specify a name and a backup specification group. Optionally, click **Save and Schedule** to save, and then schedule the backup specification. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

TIP:
 Preview your backup specification before using it for real. See [Previewing backup sessions](#), on page 157.

SAP R/3 backup options

Option	Description
Log file	If you want to create a backupint log file during backup, specify a pathname for the file. By default, this file is not created because Data Protector stores all relevant information about backup sessions in the database.
BR Backup	Specifies BRBACKUP options. For example, for online backup using the splitint interface, type <code>-t online_mirror</code> . If <code>splitint</code> is not supported by BRTOOLS, type <code>-t online_split</code> . To run BRBACKUP under a different Oracle database user than the

Option	Description
	one specified during the configuration, type <code>-u user_name</code> .
Backup Objects	Lists BRBACKUP options passed by <code>omnisap.exe</code> . The list is displayed after you save the backup specification.
BR Archive	Specifies BRARCHIVE options. Not applicable for ZDB to disk.
Balancing: By Load	Groups files into subsets of approximately equal sizes. The subsets are then backed up concurrently by Data Protector <code>sapback</code> programs. If your backup devices use hardware compression, the sizes of the original and backed up files differ. To inform Data Protector of this, specify the original sizes of the backed up files in the <code>compression</code> section of the Data Protector SAP R/3 configuration file. See Data Protector SAP R/3 configuration file , on page 128.
Balancing: By Time	Groups files into subsets that are backed up in approximately equal periods of time. The duration depends on the file types, speed of the backup devices, and external influences (such as mount prompts). This option is best for environments with large libraries of the same quality. The subsets are backed up concurrently by Data Protector <code>sapback</code> programs. Data Protector automatically stores backup speed information in the <code>speed</code> section of the Data Protector SAP R/3 configuration file. It uses this information to optimize the backup time. This type of balancing may lead to non-optimal grouping of files in the case of an online backup or if the speed of backup devices varies significantly.
Balancing: Manual	Groups files into subsets as specified in the manual balancing section of the Data Protector SAP R/3 configuration file. For more information, see Manual balancing , on page 161. Not applicable for ZDB to disk.
Balancing: None	No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order, use the Oracle Server Manager SQL command: <code>select * from dba_data_files</code>
Pre-exec , Post-exec	The command specified here is started by <code>omnisap.exe</code> on the SAP R/3 system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes. Provide only the name. The command must reside in the directory: Windows systems: <code>Data_Protector_home\bin</code>

Option	Description
	HP-UX, Solaris: /opt/omni/bin
Backup mode	Not applicable for ZDB.
Use default RMAN channels	Not applicable for ZDB.
Objects outside database	Specifies non-database files of the Oracle SAP R/3 environment to be saved. Save these files in a separate backup session.

NOTE:

The total number of sapback processes started in one session using Data Protector is limited to 256.

Modifying backup specifications

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

Scheduling backup sessions

You can run unattended backups at specific times or periodically. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

Execute:

```
omnib -sap_list backup_specification_name -test_bar
```

What happens during the preview?

The `omnisap.exe` command is started, which starts the Data Protector `testbar` command to test the following:

- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

Starting backup sessions

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

Backup methods

Start a backup of SAP R/3 objects in any of the following ways:

- Using the Data Protector GUI.
- Using the Data Protector CLI.
- Using the SAP BR*Tools.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP R/3**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify Click **OK**.

NOTE:
Only the Full backup type is supported.

ZDB to disk, ZDB to disk+tape: Specify the **Split mirror/snapshot backup** option.

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

Using the Data Protector CLI

Execute:

ZDB to tape, ZDB to disk+tape:

```
omnib -sap_list backup_specification_name
```

ZDB to disk:

```
omnib -sap_list backup_specification_name -disk_only
```

For details, see the `omnib` man page or the *HPE Data Protector Command Line Interface Reference*.

Using the SAP BRTOOLS

1. Log in to the SAP R/3 backup system or SAP R/3 application system as the Oracle OS user.
2. Export/set the following environment variables:

ORACLE_SID= <i>SAP_instance_name</i>	
ORACLE_HOME= <i>Oracle_software_home_directory</i>	
[SAPBACKUP_TYPE=OFFLINE]	Default is ONLINE.
SAPDATA_HOME= <i>database_files_directory</i>	
SAPBACKUP= <i>BRTOOLS_logs_and_control_file_copy_directory</i>	
SAPREORG= <i>BRSPACE_logs_directory</i>	
OB2BARLIST= <i>backup_specification_name</i>	The backup specification is needed only to specify which Data Protector devices should be used for backup. Other information from the backup specification, like SAP R/3 objects to be backed up or the BRBACKUP options, is ignored and has to be specified manually at run time.
[OB2_3RD_PARTY_BACKINT=1]	Specifies usage of a third-party <code>backint</code> tool to perform ZDB-to-disk+tape backup sessions. After setting the variable, copy the <code>backint</code> you want to use to the SAP BRTOOLS directory.
[OB2BARHOSTNAME= <i>application_system_name</i>]	Required if you are logged in to the backup system. Optional if you want to specify a virtual server name in cluster environments.
[OB2BACKUPHOSTNAME= <i>backup_system_name</i>]	Required if you are logged in to the application system.
OB2SMB=1	Specifies a ZDB session.
[OB2SMBIR=1]	Specifies tracking replica for instant recovery.
[ZDB_ORA_INCLUDE_CF_OLF=1]	Required if you are logged in to the backup system.
[OB2DISKONLY=1]	Specifies a ZDB-to-disk session.

If you are logged on to the backup system, ensure that the `NLS_LANG` environment variable is set to the same value as the `NLS_LANG` environment variable on the application system.

Alternatively, these variables can be specified in the `backint` parameter file. If this is required, the location of the file must be specified in the SAP configuration file using the `util_par_file` parameter:

```
util_par_file = path\filename
```

If you do not supply the path, the system searches for the parameter file in the directory:

Windows systems: `SAPDATA_HOME\database`

UNIX systems: `ORACLE_HOME/dbs`

3. Execute the `BRBACKUP` command. The command syntax depends on whether you are logged in to the application system or backup system:

Application system:

```
brbackup -t {online_split | offline_split | online_mirror | \ offline_mirror}  
[-q split] -d \ util_file -m all -c -u user/password
```

Backup system:

```
brbackup -t {online_mirror | offline_mirror} [-q split] -d util_file -m all -c  
-u user/password
```

The `-q split` option is required if `OB2DISKONLY` is set to 1.

Configuring SAP compliant ZDB sessions

SAP R/3 standards recommend that, in ZDB sessions that use the `splitint` backup interface, `BRBACKUP` is started on the backup system and not on the application system. You can configure Data Protector to comply with these standards by setting the Data Protector `OB2_MIRROR_COMP` environment variable to 1. The variable is saved in the Data Protector SAP R/3 instance configuration file. Consequently, in all `splitint` ZDB sessions for this SAP R/3 instance, `BRBACKUP` will be started on the backup system. By default, `BRBACKUP` is started on the application system.

Set the `OB2_MIRROR_COMP` environment variable using the Data Protector GUI or CLI.

NOTE:

If no backup specification for the related SAP R/3 instance exists, you cannot use the Data Protector CLI to set the `OB2_MIRROR_COMP` variable.

Using the Data Protector GUI

You can set the `OB2_MIRROR_COMP` variable when you create a backup specification or modify an existing one:

1. Proceed to the Source page of the backup specification.

NOTE:

In environments in which the control file and datafiles reside on the same source disk, Data Protector does not let you proceed to the Source page if the **Track the replica for instant recovery** option is selected. Specifically, the Data Protector instant recovery check fails.

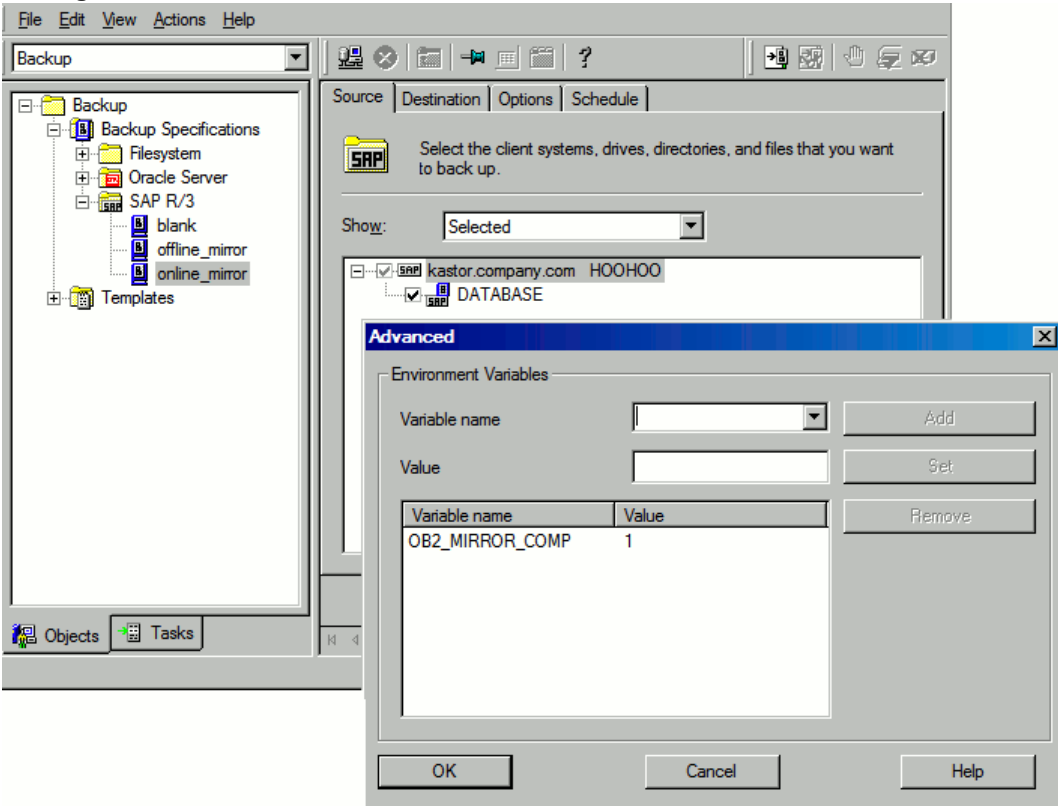
In such a case, clear the option first. You can select the option later if needed, when the OB2_MIRROR_COMP variable is already set and, consequently, the instant recovery check is no longer performed.

Right click the SAP R/3 instance at the top and click **Set Environment Variables**.

- 2. In the Advanced dialog box, set OB2_MIRROR_COMP to 1. See [Setting environment variables, below](#).

Click **OK**.

Setting environment variables



Using the Data Protector CLI

Execute the following command:

```
util_cmd -putopt SAP instance_name OB2_MIRROR_COMP 1 --sublist Environment
```

Manual balancing

Manual balancing means that you manually group files into subsets, which are then backed up in parallel. To group files into subsets, add the manual_balance section to the Data Protector SAP R/3 configuration file as described in the following example.

Example

Suppose that we have a backup specification named SAP-R3 with the following files to be backed up: fileA, fileB, fileC, fileD. To group the files into three subsets (0={fileA, fileC}, 1={fileB}, 2={fileD}), add the following lines to the Data Protector SAP R/3 configuration file:

```
manual_balance={
    SAP-R3={
        fileA=0;
        fileB=1;
        fileC=0;fileD=2;}}}
```

When you group files into subsets, consider the following:

- Use only one file from the same hard disk at a time.
- The number of files in a subset must be equal to or smaller than the sum of the concurrencies of all devices specified for backup.
- If the backup specification contains files that are not allocated to any subset, Data Protector automatically adds these files to the list of files to be backed up using the load balancing principle. Before the backup, this list is logged in:

Windows systems: *SAPDATA_HOME\sapbackup*.lst*

UNIX systems: *ORACLE_HOME/sapbackup/*.lst*

Restore

You can restore SAP R/3 objects using any of the following methods:

- **Standard restore:** Data is restored from backup media created in ZDB to tape, ZDB to disk+tape, and non-ZDB (standard backup) sessions. See [Standard restore, on page 164](#).
- **Instant recovery:** Data is restored from a replica created in *online* ZDB-to-disk or ZDB-to-disk+tape sessions. See [Instant recovery , on page 167](#).

After the restore, you can recover the database to a specific point in time using the SAP BRTOOLS interface. Instant recovery method enables you to restore and recover the database within the same session. However, you can only restore (and recover) the whole database. To restore only a part of the database or the archived logs, use the standard restore method.

[SAP recovery methods, below](#) shows which restore methods are available, depending on the backup session you restore from.

SAP recovery methods

Disk array	Backup types	Recovery of the whole database		Recovery of a part of the database
		Until now	To a point in time, logseq/thread or SCN number	
P9000 XP, P6000 EVA, EMC, NetApp Storage	ZDB to tape - online	Restore	Restore	Restore
	ZDB to tape - offline	Restore	Restore	Restore
P9000 XP, P6000 EVA	ZDB to disk - online	Instant recovery + database recovery	Instant recovery + database recovery	N/A
	ZDB to disk - offline	N/A	N/A	N/A
	ZDB to disk+tape - online	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	Restore
	ZDB to disk+tape - offline	Restore	Restore	Restore

Legend

Restore	You can do a standard restore from the Data Protector media using the Data Protector GUI or the SAP BRTOOLS. After the restore, you can recover the database using the SAP BRTOOLS.
Instant recovery + database recovery	You can do an instant recovery. You can include database recovery in the instant recovery session or do it afterwards, using the SAP BRTOOLS.
N/A	Not available.

Considerations

- SAP R/3 tablespaces located on raw partitions cannot be restored using the Data Protector GUI. Workaround: Use SAP restore commands (for example, brrestore).
- If your Oracle database is localized, you may need to set the appropriate Data Protector encoding before you start a restore. For details, see [Localized SAP R/3 objects, on page 172](#).
- Restore preview is not supported.

Standard restore


Restore SAP R/3 objects using the Data Protector Manager.

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **SAP R/3**, expand the client (backup system) from which the data was backed up, and then click the Oracle instance you want to restore.
3. In the **Source** page, select SAP R/3 files to be restored.

To restore a file under a different name or to a different directory, right-click the file and click **Restore As/Into**.

To restore a file from a specific backup session, right-click the file and click **Restore Version**.

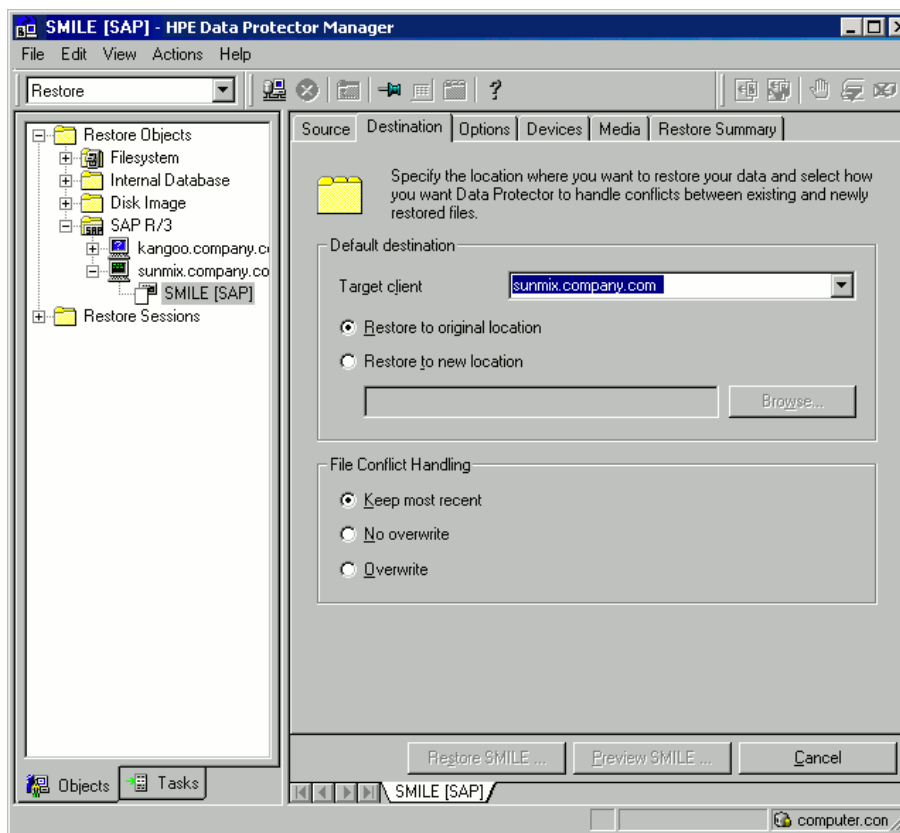
Selecting objects for restore



4. In the **Destination** tab, select the client to restore to (**Target client**). By default, this is the application system. See [Selecting the target client, below](#).

For details on options, press **F1**.

Selecting the target client




5. In the **Options** page, set the restore options. For information, press **F1**.
6. In the **Devices** page, select the devices to be used for the restore.
For more information on how to select devices for a restore, see the *HPE Data Protector Help* index: “restore, selecting devices for”.
7. Click **Restore**.
8. In the **Start Restore Session** dialog box, click **Next**.
9. Specify **Report level** and **Network load**.

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

10. **EMC and P9000 XP Array:** This step is relevant only if you have both the EMC Symmetrix Agent and HPEP9000 XP Agent components installed on the application system.
EMC: Select EMC Symmetrix restore.
P9000 XP Array: Select P9000 XP restore. See [Selecting the P9000 XP restore, on the next page](#).

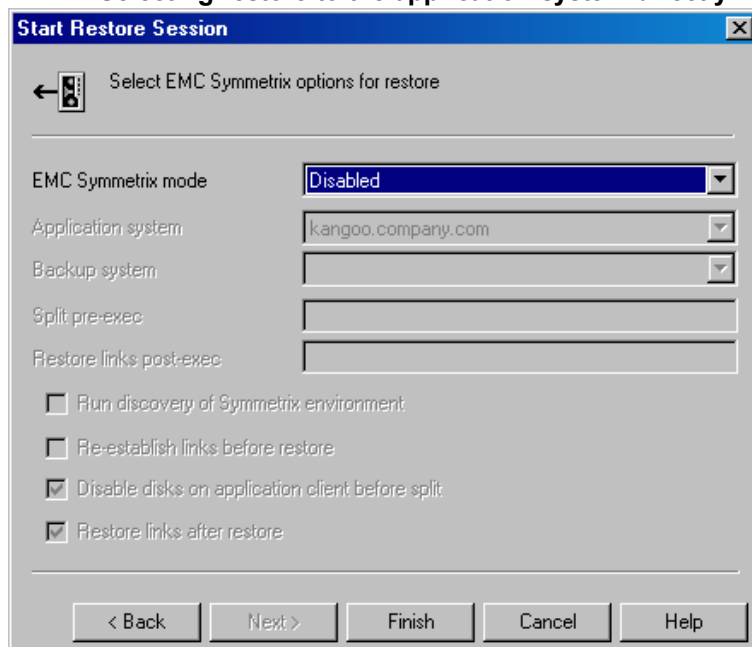
Selecting the P9000 XP restore




Click **Next**.

11. **EMC and P9000 XP Array:** From the EMC Symmetrix mode or Mirror mode drop-down list, select **Disabled**. This sets the restore from backup media to the application system directly. See [EMC – Selecting restore to the application system directly, below](#).

EMC – Selecting restore to the application system directly



P9000 XP Array – Selecting restore to the application system directly



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

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

Instant recovery

For general information on instant recovery, see the *HPE Data Protector Concepts Guide* and *HPE Data Protector Zero Downtime Backup Administrator's Guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), HPE Serviceguard, and Microsoft Cluster Server), see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

You can start an instant recovery using the Data Protector GUI or CLI.

Considerations

- The database recovery part is performed after the instant recovery procedure. During database recovery, archive log backups performed after the ZDB are restored from tape by the SAP BR*Tools utilities. If selected, the logs are reset and the database is opened.
- If the replica to be used for instant recovery contains the control file, first see [Instant recovery from replicas containing the control file, on page 171](#).

Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Shut down the Oracle database using `sqlplus`:
For example:


```
sqlplus  
sqlplus> shutdown immediate  
sqlplus> exit
```

2. In the Data Protector Manager Context List, select **Instant Recovery**.
3. Expand **SAP R/3** 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.

For HPE P9000 XP Disk Array Family, it is recommended to leave the **Keep the replica after the restore** option selected to enable a restart of an instant recovery session. The option is selected by default, except for an offline backup where the database was in NOARCHIVELOG mode during the backup. With HPE P6000 EVA Disk Array Family, replica is kept on the disk array only if the **Copy replica data to the source location** is selected.

Set other HPE P6000 EVA Disk Array Family or HPE P9000 XP Disk Array Family options. For details, press **F1**.

SAP R/3 source options



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**.
 - To automatically perform a database recovery after an instant recovery, select the recovery

options. For details, see [Database recovery options](#) , below.


Click **Restore**.

Data Protector recovers the database after performing instant recovery by switching the database to mount state, restoring the necessary archive redo logs from tape, and applying the redo logs.

Database recovery options

User name (UNIX systems only)	Specifies the user name under which the instant recovery is performed. The user needs to be a member of the DBA group.
User group (UNIX systems only)	Specifies the user group the user in the User name field belongs to. NOTE: The User name and the User group must be the same as defined in the backup ownership.
Recovery	Enables database recovery after instant recovery.
Recover until	The options in this drop-down list enables you to specify to which point in time you would like the recovery to be performed. The following options are available: Now : All existing archive logs are applied. Selected time : Only archive logs until the specified time are applied. Selected logseq/ thread number : Specifies an incomplete recovery. Only archive logs with a lower or equal number than the specified log sequence or thread number are applied. Selected SCN number : Only archive logs until the specified SCN number are applied.
Open database after recovery	Opens the database after the recovery was performed.
Reset logs	Resets the archive logs after the database is opened. This option is by default not selected if the Recover until option is set to Now . The following are Oracle recommendations on when to reset the logs: <i>Always</i> reset the logs: <ul style="list-style-type: none"> • After an incomplete recovery, that is, if not all archive redo logs are applied. • If a backup of the control file is used for recovery. <i>Do not</i> reset the logs: <ul style="list-style-type: none"> • After a complete recovery, when the control file is not used. • If the archive logs are used for a standby database. However, if you must reset the archive logs, recreate the standby database.

SAP R/3 recovery options



Instant recovery using the Data Protector CLI

Execute:

```
omnir
-host ClientName
-session SessionID
-instant_restore
[P9000_DISK_ARRAY_XP_OPTIONS | P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS]
-sap
-user UserName -group GroupName
-recover {now | time MM/DD/YY hh:mm:ss | logseq LogSeqNumber thread ThreadNumber |
SCN Number} [-open [-resetlogs]]
-appname ApplicationDatabaseName
```

The order of options is important. On Windows clients, the user name and group name options are not required. For a detailed description of the options, see the *HPE Data Protector Command Line Interface Reference*.

Instant recovery from replicas containing the control file

During an instant recovery from a replica that contains the control file, the current control file and, possibly, online redo logs get overwritten. Therefore, before you start the session, copy the current control file and online redo logs to a safe location to be able to do a database recovery afterwards.

A replica contains the control file if it is created in any of the following sessions:

- *Online* ZDB session with the `omnirc` option `ZDB_ORA_INCLUDE_CF_OLF` set to 1
- *Online* SAP compliant ZDB session
- *Offline* ZDB session (any configuration)

NOTE:

An *offline* ZDB session also contains online redo logs. You can use such a session to restore the SAP R/3 database to a point in time at which the backup was performed. In this case, you do not need to follow the steps below.

To restore and recover the database:

1. Copy the current control files and online redo logs to a safe location.
2. Perform instant recovery (without database recovery). Use the Data Protector GUI or CLI.
3. Copy the current control files and online redo logs back to their original location.
4. Mount the target database.
5. Restore missing archived redo logs required for database recovery.

Example:

```
# sqlplus user/password@net_service_name
SQL> select SEQUENCE#, NAME from V$ARCHIVED_LOG where
(NEXT_TIME>to_date('2010/10/03','YYYY/MM/DD') and (FIRST_CHANGE#<='1000'));
# brrestore -a log_no,... -d util_file -c force -u user/password
```

6. Recover the target database.

Example:

```
# rman target user/password@net_service_name
RMAN> run{
  2> allocate channel dbrec type disk;
  3> recover database until scn 1000;
  4> release channel dbrec;
  5> }
```

Restoring using another device

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

Using the Data Protector GUI

For information on how to select another device for a restore using the Data Protector GUI, see the *HPE Data Protector Help* index: “restore, selecting devices for”.

Using the Data Protector CLI or SAP commands

If you are restoring using the Data Protector CLI or SAP R/3 commands, specify the new device in the file:

Windows systems: `Data_Protector_program_data\Config\Server\cell\restoredev`

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

Use the format:

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

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

IMPORTANT:

Delete this file after use.

On Windows systems, use the Unicode format for the file.

Localized SAP R/3 objects

Oracle Server uses its own encoding, which may differ from the encoding used by the filesystem. In the Backup context, Data Protector displays the logical structure of the Oracle database (with Oracle names) and in the Restore context, the filesystem structure of the Oracle database. Therefore, to display non-ASCII characters correctly, ensure that the Data Protector encoding matches with the Oracle Server encoding during backup and with the filesystem encoding during restore. However, the incorrect display does not impact the restore.

Windows systems: If the current values of DBCS and the default Windows character set for non-Unicode programs do not match, problems arise. See [Restore problems, on page 193](#).

UNIX systems: To be able to switch between the Data Protector encodings, start the GUI in UTF-8 locale.

Monitoring sessions

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

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

On how to monitor a session, see the *HPE Data Protector Help* index: “viewing currently running sessions”.

System messages generated during backups are sent to both the SAP R/3 and the Data Protector monitor. However, mount requests are sent only to the Data Protector monitor.

Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector SAP R/3 integration.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

See also the troubleshooting section in the SAP R/3 chapter of the *HPE Data Protector Integration Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HPE Data Protector Help* index: “patches” on how to verify this.
- See the *HPE Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- For an up-to-date list of supported versions, platforms, and other information, see the latest support matrices at <https://softwaresupport.hpe.com/manuals>.

General troubleshooting

Problem

Configuration fails due to a database operation failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

```
Integration cannot be configured.
```

The database reported error while performing requested operation.

Action

Review user group membership for the user account which is used in Oracle database access authentication. For details, see [Configuring user accounts, on page 134](#).

Prerequisites on the SAP side of the integration

The following verification steps must be performed in order to verify that SAP is installed as required for the integration to work. These steps do not include Data Protector components.

1. **Verify backup directly to disk as follows:**

```
brbackup -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

2. **Verify restore directly to disk as follows:**

```
brrestore -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

3. **If you are running backups in RMAN mode, verify backup and restore directly to disk using Recovery Manager channel type disk as follows:**

- a. You must define the parameter `init` in the initialization file `initORACLE_SID.ora`. Execute the following commands:

```
brrestore -d pipe -u user/password -t online -m all
```

```
brrestore -d disk -u user/password
```

- b. If this fails, see the *SAP Online Help* to learn how to execute backup and restore directly to disk using the SAP backup utility.

Check the error message and resolve these problems before you continue.

4. **Verify that the SAP backup tools correctly start backint (which is provided by Data Protector):**

Move the original backint and create a test script `namedbackint.bat` in the directory where the SAP backup utility resides, with the following entries:

```
echo "Test backint called as follows:"  
echo "%0%1%2%3%4%5%6%7%8%9"  
exit
```

Then start the following commands:

```
brbackup -t offline -d util_file -u user/password -c
```

If you receive backint arguments, this means that SAP is properly configured for backup using backint; otherwise you have to reconfigure SAP.

See [Configuring SAP R/3 databases, on page 140](#).

Configuration problems

IMPORTANT:

The procedure described in the previous sections must be performed before you start checking the Data Protector configuration.

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

For details, see the *HPE Data Protector Installation Guide*.

2. **Perform a filesystem backup of the SAP Database Server.**

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

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

See the *HPE Data Protector Help* index “standard backup procedure” for details about how to do a filesystem backup.

3. If the SAP backup utilities are installed in a shared directory, then the inet startup parameter must be specified as described in [If you use the command line to start the Data Protector commands, verify the inet startup parameters., below](#), or the Windows permissions must be set correctly.

Execute the following command (if you use the default directory):

```
dir \\client_name\sapmnt\ORACLE_SID\SYS\exe\run\brbackup
```

or

```
dir \\client_name\SAPEXE\brbackup
```

If this fails, set the inet startup parameters, or set the correct permissions to access a Windows network directory.

4. If you use the command line to start the Data Protector commands, verify the inet startup parameters.

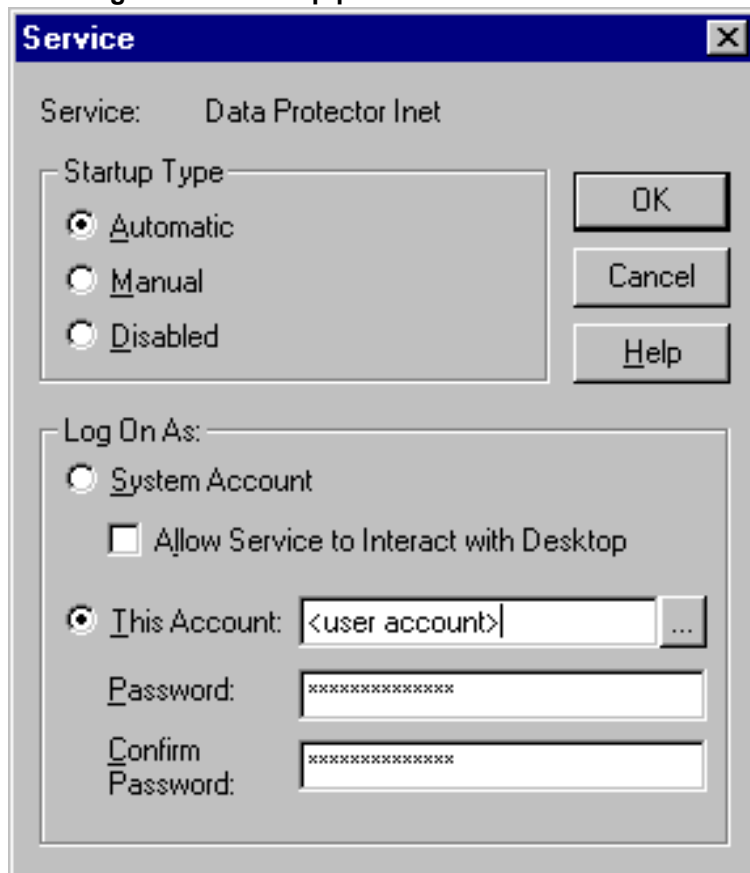
Check the Data Protector Inet service startup parameters on the SAP Database Server system. Proceed as follows:

- a. In the **Control Panel**, go to **Administrative Tools, Services**.
- b. Select **Data Protector Inet**.

In the **Services** window, select **Data Protector Inet, Startup**.

The service must run under a specified user account. Make sure that the same user is also added to the Data Protector admin user group.

Checking the Inet start-up parameters



5. Examine the environment variables.

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the Environment section of the Data Protector SAP configuration file on the Cell Manager. See [Data Protector SAP R/3 configuration file , on page 128](#).

6. Examine system errors.

System errors are reported in the `debug.log` file on the SAP Server.

Problem

Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

Integration cannot be configured.

Script failed. Cannot get information from remote host.

Action

Check the environment settings and ensure Data Protector Inet is running under a user account which has the required privileges. For details, see [Before you begin, on page 134](#).

Backup problems

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

1. Check your SAP Server configuration:

To check the configuration, start the following command on the SAP Server system:

```
Data_Protector_home\bin\util_sap.exe -CHKCONF ORACLE_SID
```

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

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

Before you run the `testbar2` utility, verify that the Cell Manager name is correctly defined on the SAP Database Server. In the default Data Protector client configuration directory, check the `cell_server` file, which contains the name of the Cell Manager system. Then execute the following command:

```
Data_Protector_home\bin\testbar2 -type:SAP -appname:ORACLE_SID -bar:backup_specification_name -perform:backup
```

Examine the errors reported by the `testbar2` utility by clicking the **Details** button in the Data Protector **Monitor** context.

If the messages indicate problems concerning the Data Protector side of the integration, create an SAP backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices. For instructions on troubleshooting devices, see the *HPE Data Protector Troubleshooting Guide*. If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barList_name
```

```
export OB2APPNAME=ORACLE_SID
```

```
Data_Protector_home\bin\backint.exe -f backup -t file -u ORACLE_SID -i input_file
```

where *input_file* is a file with a list of full pathnames for backup.

Backint anticipates a list of files in the following format: *pathName_1pathName_2pathName_3*

Problem

Backup fails with “Connect to database instance failed”

If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2  
ORA-01033: ORACLE initialization or shutdown in progress  
BR0310E Connect to database instance H00H00 failed
```

Action

Before you start a backup, ensure that the database instance is in the open or shutdown mode.

Restore problems

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

1. **Verify that a backup object exists on the backup media and in the IDB:**

This can be done by executing the command

```
omnidb -sap "object_name" -session "Session_ID" -media
```

on the SAP Database Server system.

The output of the command lists detailed information about the specified backup object, session IDs of the backup sessions containing this object, and a list of the media used.

For detailed syntax of the `omnidb` command, execute:

```
omnidb -help
```

You can also do this using the SAP tools:

Use `backint`, so that SAP tools also use this command to query:

```
Data_Protector_home\bin\backint.exe -f inquiry -u ORACLE_SID -i input_file
```

where the specified *input_file* is queried.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

Backint anticipates a list of files of the following format:

```
backup_ID_1 pathName_1 [targetDirectory_1]  
backup_ID_2 pathName_2 [targetDirectory_2]  
backup_ID_3 pathName_3 [targetDirectory_3]
```

To retrieve the *backup_ID* numbers, enter the following command:

```
echo #NULL #NULL | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify #NULL as *backup_ID_1* in the *input_file*. In this case, the latest backup session for the file is used for the restore.

2. Verify the restore using the Data Protector User Interface

This test is possible if the objects have been backed up by *backint*.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

3. Simulate a Restore Session

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

Before you execute *testbar2*, verify that the Cell Manager name is correctly defined on the SAP Database Server.

In the default Data Protector client configuration directory, check the *cell_server* file, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the *testbar2* utility:

```
Data_Protector_home\bin\testbar2 -type:SAP
  -apname:ORACLE_SID
  -perform:restore
  -object:object_name
  -version:object_version
  -bar:backup_specification_name
```

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

4. Verify the restore using *backint*

Execute the following command:

```
Data_Protector_home\bin\backint.exe -f restore -u ORACLE_SID -i input_file
```

where the contents of the *input_file* will be restored.

If this fails, check if the session was performed successfully and if the restore was started under the appropriate user account.

Backint anticipates a list of files in the following format:
backup_ID_1 *pathName_1*
[*targetDirectory_1*]*backup_ID_2* *pathName_2* [*targetDirectory_2*]*backup_ID_3*
pathName_3 [*targetDirectory_3*]

To retrieve the *backup_ID* numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

Problem

Restore sessions fail due to invalid characters in filenames

On Windows systems, where the Oracle Database Character Set (DBCS) is not set to the same value as the default Windows character set for non-Unicode programs, and where SAP tools are used to create Oracle datafiles, restore fails if the datafiles contain non-ASCII or non-Latin 1 characters.

Actions

Use any of the following solutions:

- For new Oracle installations, set the DBCS to UTF-8.
- If you do not use other non-Unicode programs, set the language for non-Unicode programs to the same value as DBCS.
- Do not use non-ASCII or non-Latin 1 characters for filenames.

Prerequisites on the SAP side of the integration

The following verification steps must be performed in order to verify that SAP is installed as required for the integration to work. These steps do not include Data Protector components.

1. Verify backup directly to disk as follows:

```
brbackup -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

2. Verify restore directly to disk as follows:

```
brrestore -d disk -u user/password
```

If this fails, check the error messages and resolve possible problems before you continue.

3. If you are running backups in RMAN mode, verify backup and restore directly to disk using Recovery Manager channel type disk as follows:

- a. Re-link the Oracle Server with the Database Library provided by SAP (`libobk.sl`).

For each RMAN channel, set the `SBT_LIBRARY` parameter to point to the `libobk.sl` file.

IMPORTANT:

Before you can use Data Protector again in the RMAN mode, you have to re-link the Oracle again with the Data Protector Database Library.

- b. You have to define the parameter `init` in the initialization file `initORACLE_SID.ora`.

Execute the following commands:

```
brrestore -d pipe -u user/password -t online -m all
```

```
brrestore -d disk -u user/password
```

If this fails, see the *SAP Online Help* to learn how to execute backup and restore directly to disk using the SAP backup utility. Check the error message and resolve this issues before you continue.

4. Verify that the SAP backup tools correctly start backint (which is provided by Data Protector):

Move the original backint and create a test script named `backint` in the directory where the SAP backup utility resides, with the following entries:

```
#!/usr/bin/sh  
echo "Test backint called as follows:"  
echo "$0 $*"  
echo "exiting 3 for a failure"  
exit 3
```

Then start the following commands as the Oracle database user described in [Configuring user accounts, on page 134](#):

```
brbackup -t offline -d util_file -u user/password -c
```

If you receive backint arguments, this means that SAP is properly configured for backup using backint; otherwise you have to reconfigure SAP.

See [Configuring SAP R/3 databases, on page 140](#).

Configuration problems

IMPORTANT:

The procedure described in the previous sections must be performed before you start checking the Data Protector configuration.

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

For details, see the *HPE Data Protector Installation Guide*.

2. **Perform a filesystem backup of the SAP R/3 Database Server:**

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

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

See the *HPE Data Protector Help* index “standard backup procedure” for details about how to do a filesystem backup.

3. **Examine the environment variables:**

If you need to export some variables before starting the Oracle Server Manager, TNS listener, or other Oracle utility, these variables must be defined in the `Environment` section of the Data Protector SAP configuration file on the Cell Manager. See [Data Protector SAP R/3 configuration file, on page 128](#).

4. **Verify the permissions of the currently used user account:**

Your user account has to enable you to perform backup or restore using Data Protector. Use the `testbar2` utility to check the permissions:

```
/opt/omni/bin/utilns/testbar2 -perform:checkuser
```

If the user account holds all required permissions, you will receive only NORMAL messages displayed on the screen.

See also [Configuring user accounts, on page 134](#).

5. **Examine system errors:**

System errors are reported in the `/var/opt/omni/log/debug.log` (HP-UX, Solaris, and Linux systems) or `/usr/omni/log/debug.log` (other UNIX systems) file on the SAP Server.

Problem

Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

Integration cannot be configured.

Script failed. Cannot get information from remote host.

Action

Resolve the problem by reviewing the user account configuration. For details, see [Configuring user accounts, on page 134](#).

Backup problems

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

1. Check your SAP Server configuration:

To check the configuration, start the following command on the SAP Server system:

```
/opt/omni/sbin/util_sap.exe -CHKCONF ORACLE_SID (HP-UX, Solaris, and Linux systems) or  
/usr/omni/bin/util_sap.exe -CHKCONF ORACLE_SID (other UNIX systems)
```

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

To get the error description, start the command:

```
/opt/omni/sbin/omnigetmsg 12 Error_number(HP-UX, Solaris, and Linux systems) or  
/usr/omni/bin/omnigetmsg 12 Error_number (other UNIX systems)
```

The message *RETVAL*0 indicates successful configuration.

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

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

```
/opt/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID -bar:backup_  
specification_name -perform:backup (HP-UX, Solaris, and Linux systems)  
/usr/omni/bin/utilns/testbar2 -type:SAP -appname:ORACLE_SID -bar:backup_  
specification_name -perform:backup (other UNIX systems)
```

Examine the errors reported by the testbar2 utility by clicking the **Details** button in the Data Protector **Monitor** context.

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

- a. Check that the owner of the backup specification is the Oracle OS user described in [Configuring user accounts, on page 134](#)
- b. Check that the respective Data Protector user group has the `See private objects` user right enabled.
- c. Create an SAP backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices.

For instructions on troubleshooting devices, see the *HPE Data Protector Troubleshooting Guide*.

If the test fails again, call support.

3. Verify the backup using backint

```
export OB2BARLIST=barList_name
export OB2APPNAME=ORACLE_SID
/opt/omni/lbin/backint -f backup -t file -u ORACLE_SID -i input_file(HP-UX,
Solaris, and Linux systems)
/usr/omni/bin/backint -f backup -t file -u ORACLE_SID -i input_file(other UNIX
systems)
```

where *input_file* is a file with a list of full pathnames for backup.

Backint expects the list of files in the following format:*pathName_1pathName_2pathName_3*

Problem

Util_File_Online SAP backup fails with “semop() error”

When the *util_file_online* option is used with BRBACKUP (for example, if you select the *Brbackup_Util_File_Online* template), the tablespaces are switched into/from backup mode individually. As there can be only one process communicating with BRBACKUP, several *sapback* processes are using a semaphore to synchronize their interaction with BRBACKUP.

The number of *sapback* processes is calculated as the sum of concurrencies of all devices used for backup. With a large number of *sapback* processes, the maximum number of processes that can have undo operations pending on any given IPC semaphore on the system may be exceeded. In such case, several *sapback* agents will fail with the following error:

```
[28] No space left on device.
```

Action

Perform any of the following actions to resolve the problem:

- Reduce the number of backup devices or their concurrency.
- Increase the value of the *semnmu* kernel parameter. After you increase the value, rebuild the kernel and restart the system.

Problem

Backup fails with “Connect to database instance failed”

If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2
ORA-01033: ORACLE initialization or shutdown in progress
BR0310E Connect to database instance H00H00 failed
```

Action

Before you start a backup, ensure that the database instance is in the open or shutdown mode.

Restore problems

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

1. Verify a user for the restore:

Verify that user specified for the restore session is the user of backup session and that he/she belongs to the Data Protector operator or admin group.

See [Configuring user accounts, on page 134](#).

2. Verify that a backup object exists on the backup media and in the IDB:

This can be done by executing the command

```
omnidb -sap "object_name" -session "Session_ID" -media (HP-UX, Solaris, and Linux systems) or
```

```
omnidb -sap "object_name" -session "Session_ID" -media (other UNIX systems)
```

on the SAP Database Server system.

The output of the command lists detailed information about the specified backup object, session IDs of the backup sessions containing this object, and a list of the media used.

For detailed syntax of the `omnidb` command, execute:

```
omnidb -help (HP-UX, Solaris, and Linux systems)
```

```
omnidb -help (other UNIX systems)
```

You can also do this using the SAP tools:

Use `backint`, so that SAP tools will also use this command to query:

```
/opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i input_file (HP-UX, Solaris, and Linux systems)
```

```
/usr/omni/bin/backint -f inquiry -u ORACLE_SID -i input_file (other UNIX systems)
```

where the specified `input_file` is queried.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

`Backint` anticipates a list of files of the following format:

```
backup_ID_1 pathName_1 [targetDirectory_1]
```

```
backup_ID_2 pathName_2 [targetDirectory_2]
```

```
backup_ID_3 pathName_3 [targetDirectory_3]
```

To retrieve the `backup_ID` numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify `#NULL` as `backup_ID_1` in the `input_file`. In this case, the latest backup session for the file is used for the restore.

3. Verify the restore using the Data Protector user interface

This test is possible if the objects have been backed up by `backint`.

If this fails, check if the backup session was performed successfully and if the query was started under the appropriate user account.

4. Simulate a restore session

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

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the SAP Database Server.

Check the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux systems) or `/usr/omni/config/cell/cell_server` (other UNIX systems) file, which contains the name of the Cell Manager system.

Then, test the Data Protector internal data transfer using the `testbar2` utility:

```
/opt/omni/bin/utilns/testbar2 -type:SAP
  -appliance:ORACLE_SID
  -perform:restore
  -object:object_name
  -version:object_version
  -bar:backup_specification_name (HP-UX, Solaris, and Linux systems) or
/usr/omni/bin/utilns/testbar2 -type:SAP
  -appliance:ORACLE_SID
  -perform:restore
  -object:object_name
  -version:object_version
  -bar:backup_specification_name (other UNIX systems)
```

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

5. Verify the restore using backint

Execute the following command:

HP-UX, Solaris, and Linux systems: `/opt/omni/sbin/backint -f restore -u ORACLE_SID -i input_file`

Other UNIX systems: `/usr/omni/bin/backint -f restore -u ORACLE_SID -i input_file`
where the contents of the `input_file` will be restored.

If this fails, check if the session was performed successfully and if the restore was started under the appropriate user account.

Backint anticipates a list of files in the following format:`backup_ID_1pathName_1`
`[targetDirectory_1]backup_ID_2pathName_2``[targetDirectory_2]backup_ID_3pathName_3`
`[targetDirectory_3]`

To retrieve the `backup_ID` numbers, enter the following command:

```
echo #NULL #NULL | backint -f inquiry -u ORACLE_SID
```

Problem

Restore of SAP R/3 tablespaces located on raw partitions fails

When restoring SAP tablespaces that are located on raw partitions using the Data Protector GUI, the restore fails with a message similar to the following:

```
[Major] From: VRDA@joca.company.com "SAP" Time: 5/9/06 3:33:51 PM
        /dev/sapdata/rsapdata Cannot restore -> rawdisk section !
[Warning] From: VRDA@joca.company.com "SAP"
Time: 5/9/06 3:42:45 PM Nothing restored.
```

Action

Use SAP commands (for example, `brrestore`) to restore these tablespaces.

Verifying the prerequisites (Oracle side)

Perform the following verification steps, in numerical order, to verify that Oracle is installed properly:

1. On the application system, verify that the target database is online, as follows:

UNIX systems:

```
export ORACLE_SID=Oracle_SID
export ORACLE_HOME=Oracle_home_path
$ORACLE_HOME/bin/sqlplus
```

Windows systems:

```
set ORACLE_SID=Oracle_SID
set ORACLE_HOME=Oracle_home_path
%ORACLE_HOME%\bin\sqlplus
```

At the SQLPlus prompt, type:

```
connect user/passwd@service
select * from dba_tablespaces
exit;
```

Try starting the target database.

2. In order to establish the TNS network connection, verify that Net8 software is configured correctly for the target database, as follows:

- On the application system, perform the following:

UNIX systems:

```
$ORACLE_HOME/bin/lsnrctl status service
```

Windows systems:

```
%ORACLE_HOME%\bin\lsnrctl status service
```

If it fails, either start the TNS listener process or see the Oracle documentation on how to create the TNS configuration file (LISTENER.ORA).

- On the application system, perform the following. Use `sqlplus`:

UNIX systems:

```
export ORACLE_SID=Oracle_SID
export ORACLE_HOME=Oracle_home_path
$ORACLE_HOME/bin/sqlplus
```

Windows systems:

```
set ORACLE_SID=Oracle_SID
set ORACLE_HOME=Oracle_home_path
```

```
%ORACLE_HOME%\bin\sqlplus
```

At the SQLPlus prompt, type:

```
connect user/passwd@service;
```

```
exit;
```

If it fails, see the Oracle documentation on how to create the TNS configuration file (TSNAMES.ORA).

Verifying the prerequisites (SAP side)

Before you begin the steps in this section, be sure you have completed all the steps in [Verifying the prerequisites \(Oracle side\)](#), on the previous page.

Perform the following verification steps, in numerical order, to verify that SAP is installed properly:

1. On the application system, verify a backup directly to disk, as follows:

```
brbackup -d disk -u user/password
```

If it fails, see the *SAP Online Help* for instructions on how to execute a backup to disk using the SAP backup utility.

2. On the application system, verify a restore from the disk, as follows:

```
brrestore -d disk -u user/password
```

If it fails, see the *SAP Online Help* for instructions on how to execute a restore to disk using the SAP restore utility.

3. On the application system, verify that SAP is configured properly, as follows:

Move the original `backint`. Create a test script with the name `backint` in the directory with the SAP backup utility, with the following entries:

```
#!/usr/bin/sh
echo "Test backint called as follows:"
echo "$0 $*"
echo "exiting 3 for a failure"
exit 3
```

Export all environment variables required by the SAP (SAPDATA_HOME, SAPBACKUP...) and then start the command with the backup owner user:

```
brbackup -t offline_split -d util_file -u user/password -c
```

or, if Data Protector uses `splitint`:

```
brbackup -t offline_mirror -d util_file -u user/password -c
```

If you receive arguments from `backint`, that means SAP is properly configured for backup using `backint`. Otherwise, you should reconfigure SAP.

Verifying the configuration

Before you begin this section, be sure that you completed all the steps provided in the sections [Verifying the prerequisites \(Oracle side\)](#), on the previous page and [Verifying the prerequisites \(SAP side\)](#), above.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. On the application system, verify a Data Protector filesystem backup of the SAP Database Server:
Perform a filesystem backup of the Oracle Server system so that you can eliminate any potential communication problems between the Oracle Server and the Data Protector Cell Manager system.
See the *HPE Data Protector Help* index “standard backup procedure” for details about how to do a filesystem backup.
If it fails, see the *HPE Data Protector Troubleshooting Guide* for help with troubleshooting a filesystem backup.
2. Verify the environment variable on the application system:
If you have to export some variables before starting the SAP backup utilities, Oracle Server Manager, or the TNS listener, set these environment variables using the Data Protector GUI.
3. Verify the permissions of the SAP user on application system:
SAP user permissions must be set to enable you to perform an SAP backup or restore with Data Protector. See [Configuring user accounts, on page 134](#). Use the `testbar2` to check the permissions:
 - Login in as an SAP user
 - Execute `/opt/omni/bin/testbar2 -perform:checkuser`If the user account has all the required permissions, you will see only the usual messages displayed on the screen.
4. Examine the system errors:
System errors are reported in the following file on the Oracle Server:
`/var/opt/omni/log/debug.log`

Verifying the backup configuration

Before you begin this section, be sure that you completed all the steps provided in the sections [Verifying the prerequisites \(Oracle side\)](#), on page 185 and [Verifying the prerequisites \(SAP side\)](#), on the previous page.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the Data Protector SAP ZDB configuration on the application system:
Execute the following command:
HP-UX and Solaris systems: `/opt/omni/lbin/util_sap -CHKCONF ORACLE_SID`
Windows systems: `Data_Protector_home\bin\util_sap.exe -CHKCONF ORACLE_SID`
If an error occurs, the error number is displayed in the form `*RETVL*error_number`.
To get the error description, on the Cell Manager, execute:
Windows systems: `Data_Protector_home\bin\omnigetmsg 12 error_number`

which is located on the Cell Manager.

HP-UX and Linux systems: `/opt/omni/lbin/omnigetmsg 12 error_number`

2. Verify the SAP user.

Check that the respective user group has the `See Private Objects` user right selected. See also [Configuring user accounts, on page 134](#).

3. On the application system, verify the backup using `testbar2`:

Execute the following to ensure that communication within Data Protector is established:

- Create a non-ZDB backup specification on the application system.

- Execute:

```
/opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID \ -perform:backup -  
file:file_name -bar barlist_name
```

If it fails, check the errors and try to fix them or call a support representative for assistance.

4. On the application system, verify the backup using `backint`:

Execute the following command to ensure that communication within Data Protector is established and that a backup of files can be performed:

- Create a non-ZDB backup specification on the backup system.

- `export OB2BARLIST=barlist_name`

```
export OB2APPNAME=ORACLE_SID
```

```
/opt/omni/lbin/backint -f backup -t file -u ORACLE_SID -i \ input_file
```

where `input_file` is the file containing the full pathnames for backup.

If it fails, check the errors and try to fix them or call a support representative for assistance.

Verifying restore

Before you begin this section, be sure that you completed all the steps provided in the sections [Verifying the prerequisites \(Oracle side\)](#), on page 185 and [Verifying the prerequisites \(SAP side\)](#), on page 186.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the user for the restore

Verify that the user specified for the restore session is the user of the backup session and that they belong to the Data Protector operator or admin group. Check that the respective user group has the `See private objects` user right selected.

2. Verify that files are backed up and in the Data Protector database:

- Using the `omnidb` command;
See the appropriate man page on using the `omnidb` command.
- Using `backint`;

SAP tools also use this command to make a query.

```
/opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i input_file
```

where *input_file* is what will be queried. Backint expects a list of files in the following format:

```
backint_ID_1 pathName_1  
backint_ID_2 pathName_2  
backint_ID_3 pathName_3
```

To retrieve the *backint_ID* numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify #NULL as *backint_ID_1* in the *input_file*. In this case, the latest backup session for the file is used for the restore.

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check the user rights. Was the query started under the correct SAP user account?
- Call a support representative for assistance.

3. Verify the restore using Data Protector or CLI:

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the restore started under the correct SAP user account?
- Call a support representative for assistance.

4. Verify the restore using *testbar2*:

Execute the following to ensure that restore is possible:

```
/opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID \ -perform:restore -  
file:file_name -bar barlist_name -object objectName
```

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the a restore started under the correct SAP user account
- Call a support representative for assistance.

5. Verify the restore using *backint*:

backint is the same command used by the SAP backup utility.

```
/opt/omni/lbin/backint -f restore -u ORACLE_SID -i input_file
```

where *input_file* specifies what will be restored; *backint* expects the list of files in the following format:

```
                backint_ID_1
                pathName_1 [targetDirectory_1]
backint_ID_2pathName_2 [targetDirectory_2]
backint_ID_3pathName_3 [targetDirectory_3]
```

To retrieve the *backint_ID* numbers, enter the following command:

```
echo "#NULL #NULL" | backint -f inquiry -u ORACLE_SID
```

or, alternatively, you can just specify #NULL as *backint_ID_1* in the *input_file*. In this case, the latest backup session for the file is used for the restore.

If it fails, proceed as follows:

- Check the backup session - was it successful?
- Check that the files are in the Data Protector database.
- Check the user rights. Was the restore started under the correct SAP user account?
- Call a support representative for assistance.

Configuration and backup problems

The following list gives a description of problems and actions to be taken to resolve them:

- **The Server Manager is unable to connect to the destination**

Check whether the Oracle TNS listener process is up and running. Check whether there are any environment variables required for a successful remote connection to the target database; for example, *TNS_ADMIN* and *SHLIB_PATH*. Set these environment variables using the Data Protector GUI.

For more information on the Data Protector SAP configuration file, see [Data Protector SAP R/3 configuration file](#), on page 128.

- **Configuration procedure fails**

Check whether the Oracle Server is up and running.

Check the login information for the target from the application system using Oracle Server Manager. If you cannot log in, then perform the following actions:

Check whether *sysoper* and *sysdba* rights are set for the Oracle administrator user.

Examine system errors reported in the *debug.log*, *sap.log* and *oracle8.log* files, located in the default Data Protector log files directory.

If you have special Oracle environment settings, ensure that they are registered in the Environment sublist of the Data Protector SAP configuration file:

```
/etc/opt/omni/server/integ/config/SAP/ client_name%ORACLE_SID (UNIX Cell Manager),  
or Data_Protector_program_data\Config\server\integ\config\ sap\client_name%ORACLE_  
SID (Windows Cell Manager).
```

For more information on the Data Protector SAP configuration file, [Data Protector SAP R/3 configuration file](#), on page 128.

- **Starting the backup fails**

On UNIX systems, check the output of the following command on the application system:

```
/opt/omni/sbin/util_sap.exe -CHKCONF ORACLE_SID
```

In case of an error, the error number is displayed in the form:

```
*RETVAl *Error_number
```

To get the error description, start the following command on the application system:

```
/opt/omni/sbin/omnigetmsg 12 Error_number
```

On Windows systems, perform the following procedure using the Data Protector GUI:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup**, **Backup Specifications**, and then **SAP R/3**. A list of SAP backup specifications is displayed.
3. In the Scoping Pane, select the failed backup specification and right-click on the SAP R/3 server item in the Results Pane to display a pop-up menu.
4. From the pop-up menu, select **Check Configuration**.

A short description of the problems and how to resolve them is displayed.

- **Backup does not work**

- Check whether the Cell Manager is correctly set on the application system. The file `/etc/opt/omni/client/cell_server` (UNIX systems) or `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard Enterprise\OpenView\OmniBackII\Site\CellServer` (Windows systems) must contain the name of the Cell Manager.
- Check that the `primary_db` parameter in the `initORACLE_SID.sap` file on the application system is set to `LOCAL`.
- On UNIX systems, check whether the users are properly configured in user groups. Both the UNIX Oracle administrator (`oraORACLE_SID`) and UNIX SAP administrator (`ORACLE_SIDadm`) have to be in the Data Protector operator class.
- On UNIX systems, check whether the permissions of the `SAPDATA_HOME/sapbackup/` directory are set to `755`.
- On Windows systems, check that the user account that started the Data Protector Inet service is added in the Data Protector operator class.

- **Backup fails with “Connect to database instance failed”**

If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2  
ORA-01033: ORACLE initialization or shutdown in progress  
BR0310E Connect to database instance H00H00 failed
```

Before you start a backup, ensure that the database instance is the open or shutdown mode.

Problem

Configuration fails due to a script failure

During configuration of an SAP R/3 database, Data Protector reports the following error:

```
Integration cannot be configured.
```

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

Action

If the SAP R/3 database is located on a Windows system, check the environment settings and ensure Data Protector Inet is running under a user account which has the required privileges. For details, see [Before you begin, on page 134](#).

If the SAP R/3 database is located on a UNIX system, resolve the problem by reviewing the user account configuration. For details, see [Configuring user accounts, on page 134](#).

Problem

Backup using backint fails on Solaris and HP-UX

A backup using backint fails on a Solaris and HP-UX systems with the following error:

```
[Major] From: OB2BAR_DMA@computer.company.com "SAP" Time: 4/29/09 3:55:52 PM  
Cannot open file '/saphome/SAP/sapbackup/cntrlSAP.dbf'. Error: 13
```

Action

Share the directories on the application system through NFS with root permissions.

On the application system, add the following line in the file `/etc/dfs/dfstab`:

Solaris systems:

```
share -F nfs -o anon=0 /usr/src
```

HP-UX systems:

```
share -F nfs -o anon=0,rw -d "" SAPHOME
```

On the application system, add the following line in the file `/etc/exports`:

RHEL/SUSE systems:

```
SAPHOME backuphost(rw, sync)
```

Problem

A ZDB session fails after reporting connections errors

A ZDB session reports the following critical errors and completes with failures:

```
Connection with DMA was reset
```

```
Signal SIGABRT (6) received from BRtools
```

Action

In the SAP parameter file on the application system, set the parameter `primary_db` to LOCAL and restart the session.

For details, see [Configuring the SAP R/3 parameter file, on page 147](#).

Restore problems

Problem

ZDB, restore, or instant recovery sessions fail due to invalid characters in filenames

On Windows systems, where the Oracle Database Character Set (DBCS) is not set to the same value as the default Windows character set for non-Unicode programs, and where SAP tools are used to create Oracle datafiles, ZDB, restore, and instant recovery fail if the datafiles contain non-ASCII or non-Latin 1 characters.

Actions

Use any of the following solutions:

- For new Oracle installations, set the DBCS to UTF-8.
- If you do not use other non-Unicode programs, set the language for non-Unicode programs to the same value as DBCS.
- Do not use non-ASCII or non-Latin 1 characters for filenames.

Problem

Restore of SAP R/3 tablespaces located on raw partitions fails

When restoring SAP tablespaces that are located on raw partitions using the Data Protector GUI, the restore fails with a message similar to the following:

```
[Major] From: VRDA@joca.company.com "SAP" Time: 5/9/06 3:33:51 PM  
/dev/sapdata/rsapdata Cannot restore -> rawdisk section ![Warning] From:  
VRDA@joca.company.com "SAP" Time: 5/9/06 3:42:45 PM Nothing restored.
```

Action

Use SAP commands (for example, `brrestore`) to restore these tablespaces.

Chapter 3: Data Protector Microsoft SQL Server ZDB integration

Introduction

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

During the backup, an SQL Server snapshot is made (the database files are frozen and any transactions to them are cached), so the database is highly available (*online* backup). The I/O to it is suspended during the time it takes to create a **replica** (split the mirror disks or create snapshots).

NOTE:

SQL Server snapshot is an SQL Server related term and does not mean the same as a disk array snapshot.

The following disk arrays and array configurations are supported:

Supported array	Supported configurations
EMC Symmetrix (EMC)	Dual-host TimeFinder
HPE P9000 XP Disk Array Family (P9000 XP Array)	HPE BC P9000 XP, HPE CA P9000 XP, combined HPE CA+BC P9000 XP
HPE P6000 EVA Disk Array Family (P6000 EVA Array)	HPE BC P6000 EVA, combined HPE CA+BC P6000 EVA
Non-HPE Storage Arrays (NetApp, EMC VNX, EMC VMAX)	Local replication

All ZDB types (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are supported by this integration. For ZDB types description, see the *HPE Data Protector Concepts Guide* and the *HPE Data Protector Help*.

Using Data Protector, you can restore your SQL Server data:

- From backup media to the application system on LAN (standard restore).
- Using the instant recovery functionality.

The following table gives an overview of SQL Server recovery methods:

ZDB type	Recovery method
ZDB to tape	Standard restore
ZDB to disk	Instant recovery
ZDB to disk+tape	Standard restore, instant recovery

For a description of ZDB and instant recovery (IR) concepts, see the *HPE Data Protector Concepts Guide*.

ZDB and IR do not support an SQL Server availability group configuration.

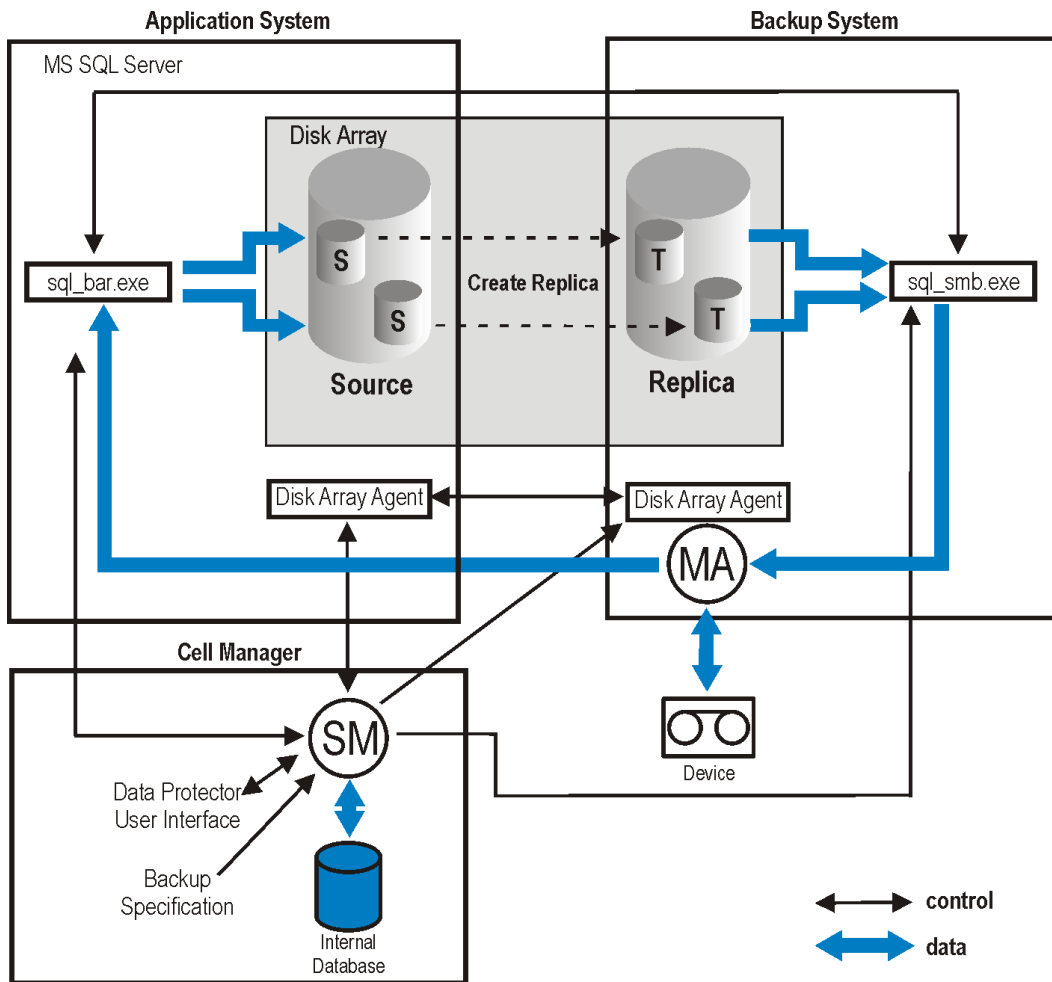
Integration concepts

Data Protector integrates with SQL Server through the Data Protector `sql_bar.exe` executable, installed on SQL Server. During backup, `sql_bar.exe`, started on the application system, connects to SQL Server to find the locations of the database files. The integration then backs up SQL Server database(s), which are replicated within a disk array.

During restore, `sql_bar.exe` connects to SQL Server to receive the restore data, which is then written to disks.

ZDB process depends on whether you replicate your data in a split mirror or snapshot replication and on the selected ZDB type. Restore process depends on the restore type - standard restore or instant recovery. See the *HPE Data Protector Concepts Guide* for a detailed description of replication techniques and ZDB and restore processes.

Backup and restore concepts



Configuring the integration

Prerequisites

- You need a license to use the SQL Server ZDB integration. For information, see the *HPE Data Protector Installation Guide*.
- Make sure that you correctly installed and configured SQL Server.
 - For supported versions, platforms, devices, and other information, see the latest support matrices at <https://softwaresupport.hpe.com/manuals>.
 - For information on installing, configuring, and using SQL Server, see the SQL Server documentation.
- Make sure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing a Data Protector disk array integration (P6000 EVA Array, P9000 XP Array, EMC, or NetApp Storage) with SQL Server, see the *HPE Data Protector Installation Guide*.

Every SQL Server to be used with Data Protector must have the MS SQL Integration component installed.

- Install SQL Server on the application system. Install user databases on the disk array source volumes (system databases can be installed anywhere). If the system databases are also installed on a disk array, they *must* reside on *different* source volumes than user databases.

If SQL Server is installed on the backup system as well, its databases *must* reside on source volumes *different* from the source volumes used for this integration. Drive letters/mount points assigned to those volumes must also be different from the drive letters/mount points assigned to the volumes used for this integration.

Before you begin

- Configure devices and media for use with Data Protector. For instructions, see the *HPE Data Protector Help* index: “configuring devices” and “creating media pools”.
- If you plan to use **Integrated authentication** to connect to an SQL Server instance, you need to restart the Data Protector Inet service under a Windows domain user account that has the appropriate SQL Server permissions for running backup and restore sessions. For supported Windows operating systems, use user impersonation. For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: “Inet user impersonation”.
- Using the SQL Server Management Studio, add the user account which you will use for backing up and restoring SQL Server data to the fixed server role `sysadmin`. For instructions, see the SQL Server documentation.
- To test whether SQL Server and Cell Manager communicate properly, configure and run a Data Protector filesystem ZDB and restore. For instructions, see the *HPE Data Protector Help*.

Data Protector SQL Server configuration file

Data Protector stores integration parameters for every configured SQL Server on the Cell Manager in:

HP-UX and Linux systems:

```
/etc/opt/omni/server/integ/config/MSSQL/ClientName%InstanceName
```

Windows systems:

```
Data_Protector_program_data\Config\Server\Integ\Config\MSSQL\ClientName%InstanceName
```

Configuration parameters are the username and password of the SQL Server user, who must have permissions to run backups and restores within SQL Server (assuming the standard security is used). They are written to the Data Protector SQL Server configuration file during configuration of the integration.

The content of the configuration file is:

```
Login='user';  
Password='encoded_password';  
Domain='domain';  
Port='PortNumber';
```

IMPORTANT:

To avoid backup problems, make sure that the syntax of your configuration file matches the examples.

Examples

- **SQL Server authentication:**

```
Login='sa';  
Domain='';  
Password='jsk74yh80fh43kdf';
```

- **Windows authentication:**

```
Login='Administrator';  
Domain='IPR';  
Password='dsjf08m80fh43kdf';
```

- **Integrated authentication:**

```
Login='';  
Domain='';  
Password='kf8u3hdgtfh43kdf';
```

Configuring an SQL Server cluster

In a cluster, all the nodes must be installed as Data Protector cluster-aware clients and the Data Protector Inet service on all nodes must run under a Windows domain user account that has also cluster administrator rights.

You must configure the Data Protector Inet service user impersonation for all cluster nodes. The Windows domain user account that is used must be given the following Windows operating system Security Policy privileges:

- Impersonate a client after authentication
- Replace a process level token

For more information, see the *HPE Data Protector Help* index: “cluster-aware client”, “Inet user impersonation”, and the SQL Server cluster documentation.

Configuring SQL Server instances

An SQL Server instance is configured during the creation of the first backup specification. The configuration consists of setting the user account that Data Protector should use to connect to the SQL Server instance. The specified login information is saved to the Data Protector SQL Server instance configuration file on the Cell Manager.

NOTE:

Make sure that the user account to be used has appropriate SQL Server permissions for running backups and restores. Check the permissions using SQL Server Enterprise Manager.

You can change configuration by following instructions described in [Changing and checking configuration, on page 200](#).

Prerequisites

- SQL Server must be online during configuration.
- Make sure that the SQL Server Browser service is running.
- Configuration must be performed for every SQL Server instance separately.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank Microsoft SQL Server Backup** template and specify backup type. For details, see [In the Create New Backup dialog box, select the Blank Microsoft SQL Server Backup template. , on page 202](#).

Click **OK**.

4. Specify the ZDB-specific options. For details, see [Under Client systems, select the SQL Server system. In cluster environments, select the virtual server of the SQL Server resource group., on page 203](#).

5. In **Application database**, select or specify the name of the SQL Server instance.

Windows Server 2008: If you intend to use **Integrated authentication** and you want that the backup session to run under the specified operating system user account, specify the **Specify OS user** option. For information on the **User and group/domain** options, press **F1**.

Click **Next**.

6. In the Configure MS SQL Server dialog box, specify the user account that Data Protector should

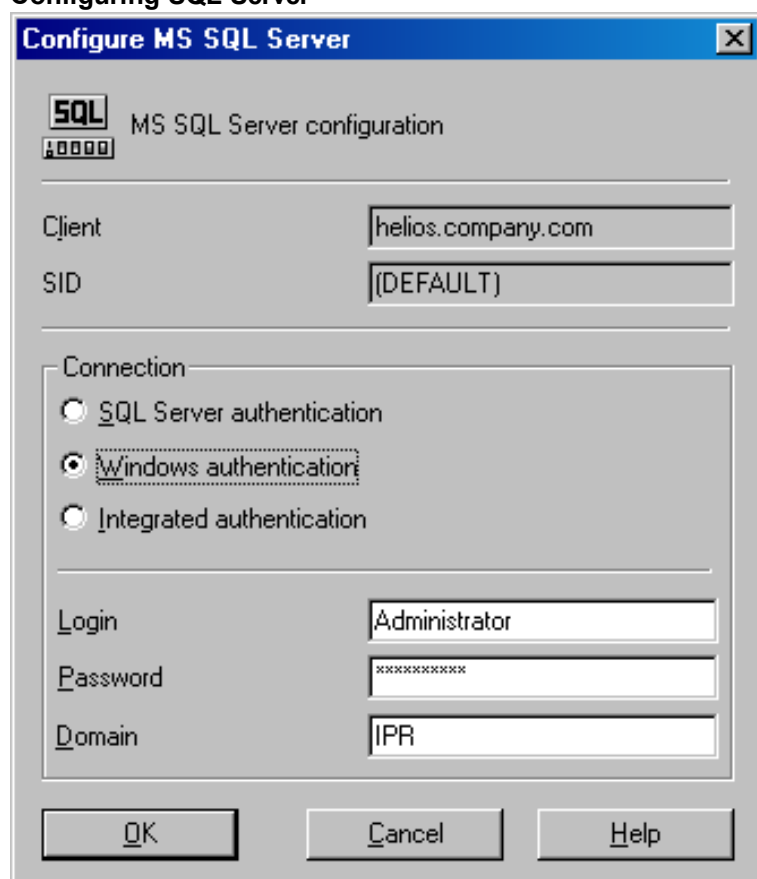
use to connect to the SQL Server instance.

- **SQL Server authentication:** SQL Server user account. Specify a username and password.
- **Windows authentication:** Windows domain user account (preferred option). Specify a username, password, and the domain.
- **Integrated authentication:** Select this option to enable Data Protector to connect to the SQL Server instance with the following Windows domain user account:
 - **Windows Server 2008:** The account specified in the **User and group/domain** options in the previous step or in the Client selection page.
 - **Other Windows systems:** The account under which the Data Protector Inet service on the SQL Server system is running.

Make sure that the user account you specify has the appropriate permissions for backing up and restoring the SQL Server databases.

See [Configuring SQL Server](#), below.

Configuring SQL Server



NOTE:

It is recommended that the SQL Server system administrator configures the integration.

For details about security, see the SQL Server documentation.

Click **OK** to confirm the configuration.

7. The SQL Server instance is configured. Exit the GUI or proceed with creating the backup specification at [Creating ZDB specifications, on page 202](#).

Using the Data Protector CLI

Execute:

```
sql_bar config [-appsrv:SQLServerClient] [-instance:InstanceName] [-  
dbuser:SQLServerUser -password:password | -dbuser:WindowsUser -password:password -  
domain:domain]
```

Parameter description

<code>-appsrv:SQLServerClient</code>	The client system on which the SQL Server instance is running. This option is not required if you execute the command locally.
<code>-instance:InstanceName</code>	The SQL Server instance name. If you omit this option, the default SQL Server instance is configured.
<code>-dbuser:SQLServerUser - password:password</code>	The SQL Server user account (SQL Server authentication)
<code>-dbuser:WindowsUser - password:password - domain:domain</code>	The Windows domain user account (Windows authentication)

NOTE:

If no user account is specified, Data Protector uses **Integrated authentication**.

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


Changing and checking configuration

You can check and change configuration using the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS SQL Server**. Click a backup specification for which you want to change the configuration.
3. In the **Source** property page, right-click the SQL Server name and select **Configure**.
4. Configure SQL Server as described in [Configuring SQL Server instances, on page 198](#).
5. Right-click SQL Server and select **Check Configuration**. See [Checking configuration, below](#).

Checking configuration



Using the Data Protector CLI

To change the configuration, execute the command for configuring SQL Server instances again, entering different data.

To check configuration, execute:

```
sql_bar chkconf [-instance:InstanceName]
```

If the optional parameter `-instance:InstanceName` is not specified, the default instance is checked.

If the integration is not properly configured, the command returns:

```
*RETVAl *8523
```

To get the information about the existing configuration, execute:

```
sql_bar getconf [-instance:InstanceName]
```

If `-instance:InstanceName` is not specified, Data Protector returns configuration for the default instance.

Backup

To run ZDB of an existing SQL Server ZDB specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

Prerequisites

- In case of nested mount points, the same drive letters, on which the source volumes to be replicated reside on the application system, must exist on the backup system to enable successful mounting of the target volumes. If the same drive letters do not exist on the backup system, the backup fails.

Considerations

Your session will fail if:

- You start ZDB, restore, or instant recovery using the same source volume on the application system at the same time. A session must be started only after the preceding session using the same source volume on the application system finishes.
- SQL Server services are not running when the backup starts.

To configure a ZDB, create a Data Protector SQL Server ZDB specification.

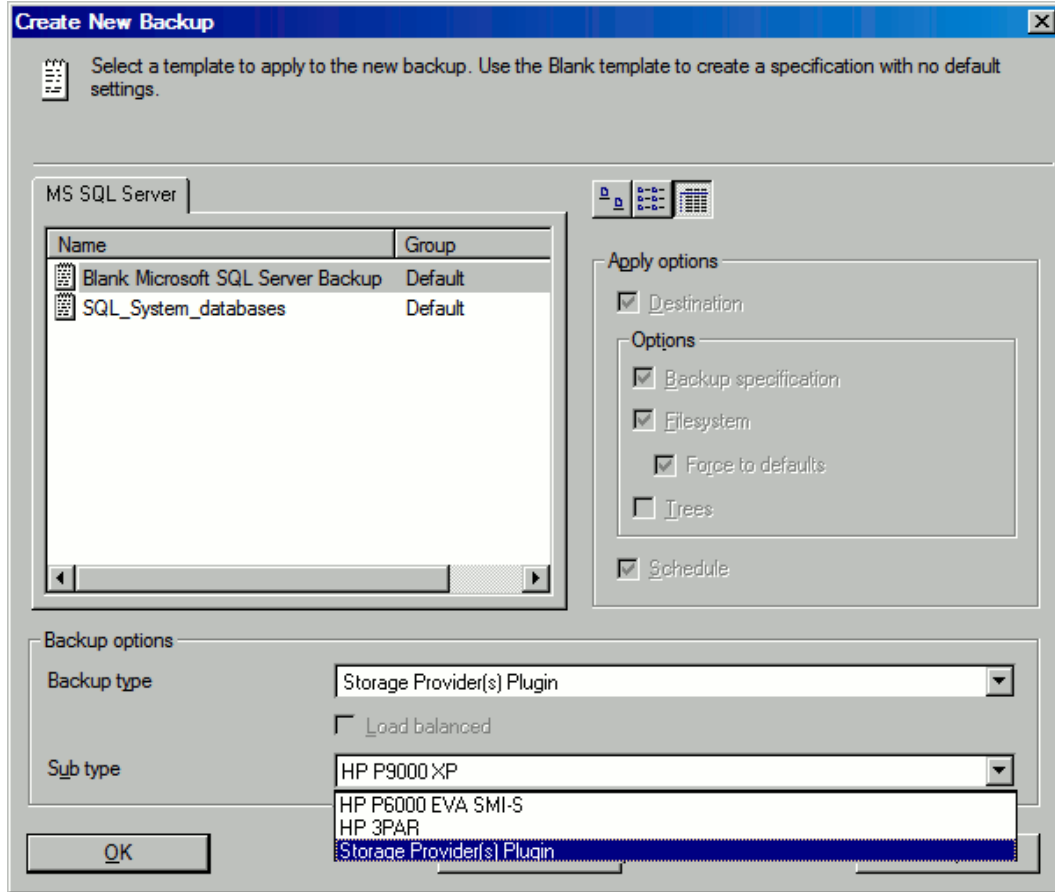
Creating ZDB specifications

Create a ZDB specification, using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the **Blank Microsoft SQL Server Backup** template.

From the **Backup type** drop-down list, select **Snapshot or split mirror backup**, and from the **Sub type** drop-down list, select the appropriate disk array agent. The agent must be installed on the application system and the backup system. See [Selecting a Microsoft SQL Server backup template and the snapshot or split mirror backup, on the next page](#).

Selecting a Microsoft SQL Server backup template and the snapshot or split mirror backup



Click **OK**.

4. Under **Client systems**, select the SQL Server system. In cluster environments, select the virtual server of the SQL Server resource group.

In the **Backup system** drop-down list, select the backup system.


Select other disk array-specific backup options (see for EMC, [P9000 XP Array backup options, on page 205](#) for P9000 XP Array, [P6000 EVA Array backup options, on page 206](#) for P6000 EVA Array, [NetApp Storage backup options, on page 208](#) for NetApp Storage, [EMC VNX Storage backup options, on page 209](#) for EMC VNX Storage, or [EMC VMAX Storage backup options, on page 210](#) for EMC VMAX Storage). For detailed information on the backup options, press **F1**.

EMC:

In the EMC GeoSpan for Microsoft Cluster Service environment, select the backup system for the active node and specify the TimeFinder configuration.

After a failover in EMC GeoSpan for MSCS, select the backup system for the currently active node and save the backup specification.

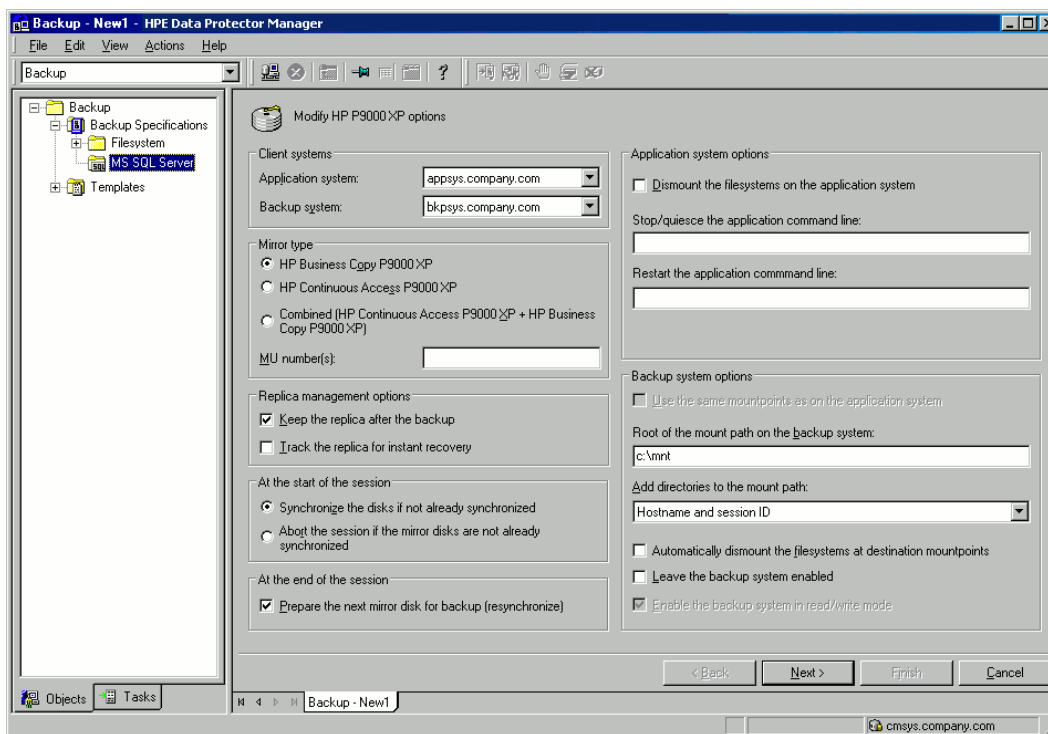
EMC backup options



P9000 XP Array:

To enable instant recovery, leave **Track the replica for instant recovery** selected.


P9000 XP Array backup options



P6000 EVA Array:


To enable instant recovery, select **Track the replica for instant recovery**.

P6000 EVA Array backup options

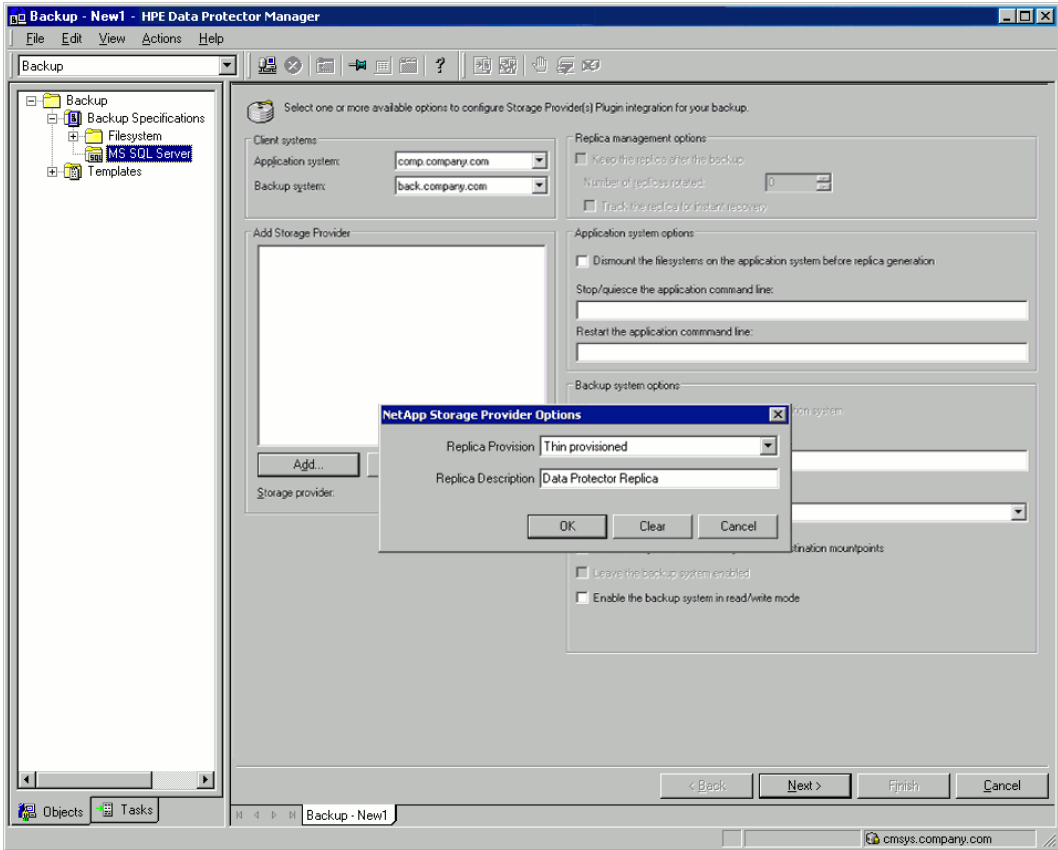


P10000 3PAR Array specifics:


P10000 3PAR Array backup options



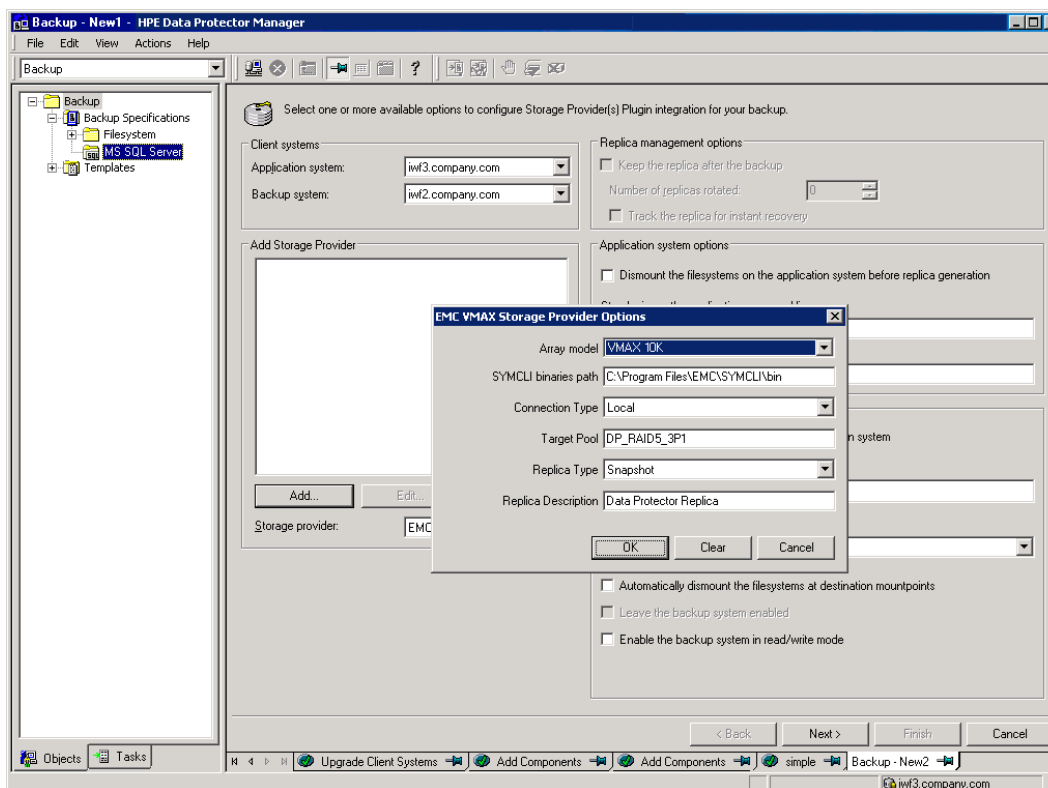
NetApp Storage:
NetApp Storage backup options



EMC VNX Storage: EMC VNX Storage backup options



EMC VMAX Storage: EMC VMAX Storage backup options



Click **Next**.

5. In **Application database**, specify the name of the SQL Server instance.

Windows Server 2008: If you intend to use **Integrated authentication** and you want that the backup session runs under the specified operating system user account, specify the **Specify OS user** option. For information on the **User and group/domain** options, press **F1**.

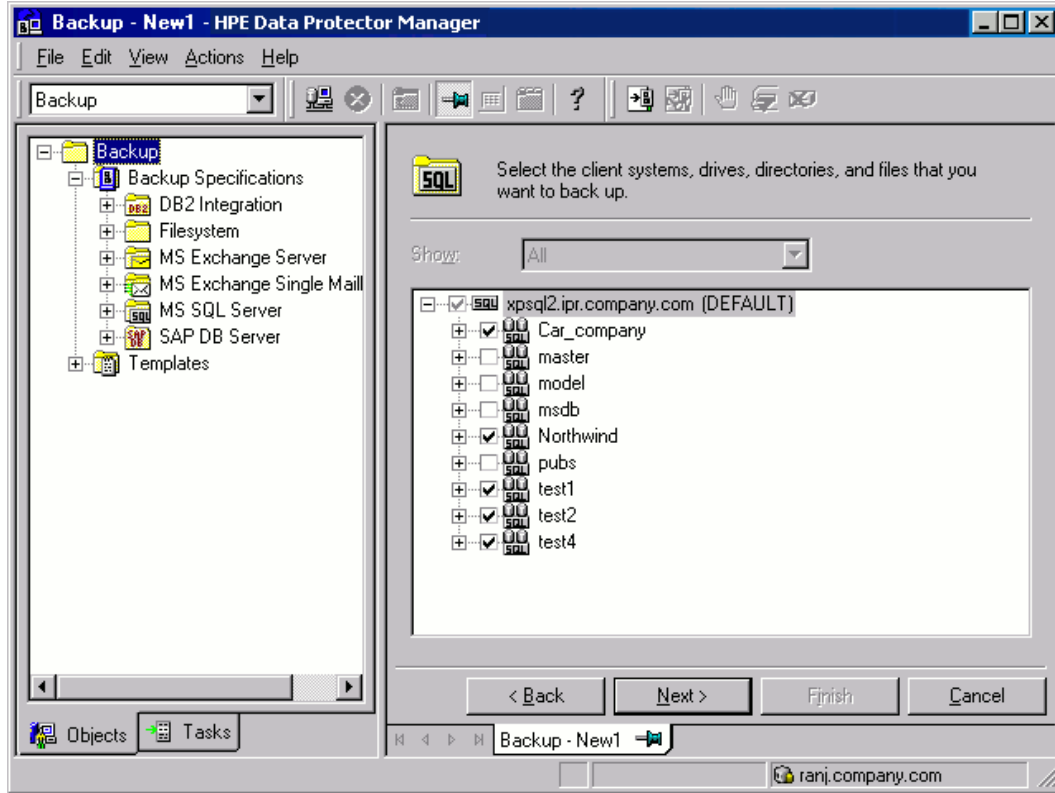
Click **Next**.

6. If the client is not configured, the **Configure MS SQL Server** dialog box appears. Configure it as described in [Configuring SQL Server instances, on page 198](#).
7. Select the databases to be backed up.

IMPORTANT:

To enable instant recovery, create different backup specifications for user and system databases.

Selecting user databases



Click **Next**.

8. Select the devices. Click **Properties** to set the media pool and preallocation policy. The device concurrency is set to 1 and cannot be changed. For more information on options, press **F1**.
To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror/Remove mirror**. Select separate devices for each mirror. The minimum number of devices for mirroring equals the number of devices used for backup.
For more information on object mirroring, see the *HPE Data Protector Help*.

NOTE:
Object mirroring is not supported for ZDB to disk.

Click **Next**.


9. Select backup options.
For information on **Backup Specification Options** and **Common Application Options**, see the *HPE Data Protector Help*.
For information on **Application Specific Option**, see [SQL Server-specific backup options, on the next page](#).
Click **Next**.
10. Click **Save As** to save the backup specification, specify a name and a backup specification group. Optionally, you can click **Save and Schedule** to save, and then schedule the backup specification. For more information on how to create and edit schedules, see Scheduler in Data

Protector in *HPE Data Protector Administrator's Guide*.

Note that only **Full** backup is performed.

You start the backup specification by clicking **Start Backup**.


Saving a backup specification



SQL Server-specific backup options

Specify SQL Server-specific backup options by clicking **Advanced** in the **Application Specific Options** group box page.

Application-specific options



SQL Server backup options

Pre-exec	Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server before backup. Resides in the default Data Protector commands directory. Only the filename must be provided in the backup specification.
-----------------	---

Post-exec	Specifies a command with arguments or a script started by <code>sql_bar.exe</code> on SQL Server after backup. Resides in the default Data Protector commands directory. Only the filename must be provided in the backup specification.	
Concurrent streams	Sets the number of concurrent streams used to back up SQL Server databases from the replica to tape. Applicable for ZDB-to-tape and ZDB-to-disk+tape sessions.	
Fast direct mode	Ignored for ZDB sessions.	
Check database integrity	Performs data integrity validation before backup. If the check fails, the session completes with warnings.	
SQL backup compression	Specify how Data Protector should handle the Microsoft SQL Server backup compression.	
	SQL Server settings (default)	Handles the backup compression according to the Microsoft SQL Server settings.
	Enable	Executes the backup compression regardless of the Microsoft SQL Server settings.
	Disable	Specifies that the backup compression should not be executed regardless of the Microsoft SQL Server settings.
Exclude from backup <i>(available for standalone instance backup only)</i>	Excludes specific databases from backup.	
	Availability Group Databases	Excludes databases belonging to any availability group from backup.
	Standalone databases	Excludes all standalone databases from backup.
	None (default)	Does not exclude any database from backup.
	Use SQL server settings (default)	Performs backup according to the Microsoft SQL Server settings.

NOTE:

Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

Scheduling backups

You can run unattended ZDB at specific times or periodically. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

NOTE:

You cannot run ZDB to disk or ZDB to disk+tape if **Track the replica for instant recovery** is not selected in the backup specification.

Starting backup sessions

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


Using the Data Protector GUI

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

For ZDB sessions, the backup type is set to **Full**.

For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

Starting interactive backups



Using the Data Protector CLI

To start ZDB to tape or ZDB to disk+tape, execute:

```
omnib -mssql_list ListName
```

To start ZDB to disk, execute:

```
omnib -mssql_list ListName -disk_only
```

where *ListName* is the name of the backup specification. For more information on omnib, see its man page.

Restore

Data Protector offers restore from backup media to the application system on LAN (standard restore), where you can select various restore options depending on your restore scenario, and instant recovery. For more information, see the following sections.

Before you begin

- Verify that the databases to be restored are not in use.
- In an availability group configuration, restore to a different client and instance is mandatory. User must select the restore options with appropriate values for the fields 'Restore to another Client' and 'Restore to another Instance'. Make sure that you do not select an availability group listener for the target client (as it is not supported), and that the selected SQL Server instance exists on the target client. Also make sure that the database which you selected for the restore does not belong to any availability group.

Standard restore

NOTE:

There is no need to create an empty database before restore, because the database and its files are generated automatically by SQL Server.


Proceed as follows using the Data Protector Manager:

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects, MS SQL Server**, and then select the client (backup system) from which you want to restore. A list of backed up objects is displayed in the Results Area.
3. Select the SQL Server objects that you want to restore. See [Selecting backup objects for restore, on the next page](#).

NOTE:


Selecting datafiles for restore means that a file restore will always be done even if all the datafiles under database were selected. To restore entire database, user should explicitly select the database level.

Selecting backup objects for restore



To select backup object-specific options, right-click the object and select **Properties**.

Selecting object-specific options



In the Version tab, select the backup version (backup date) which you want to use for restore.

Select other restore options as appropriate. Note that some options are not available for restore of data files. See [Restore options, on the next page](#) for details.

Click **OK**.

5. In the **Options** property page, specify new locations for the databases, if you want to restore your data to a different client or instance.

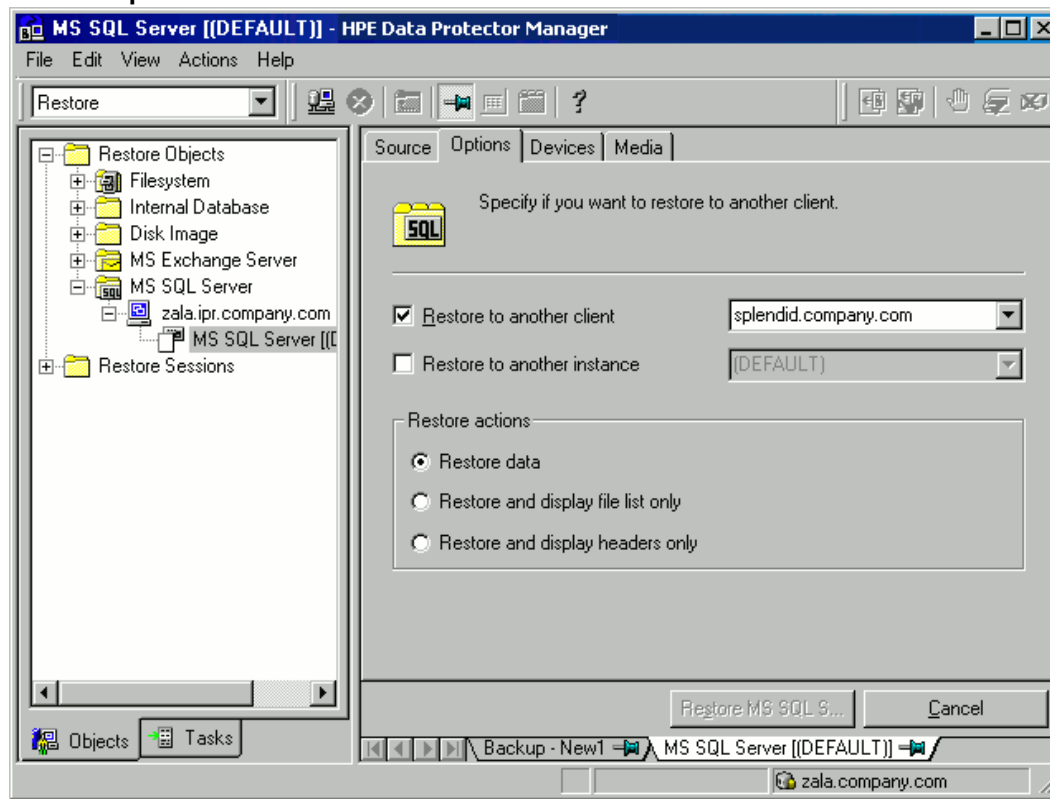
IMPORTANT:

- When you click **Options**, the cell is browsed for running SQL Server instances that can become target instances for restore. If no instances are found, **Restore to another instance** is disabled and the message **There are no instances on this client system** is displayed.
- Make sure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.

Select one of the following **Restore actions**:

- **Restore data** . Select to restore the whole database. This option is selected by default.
- **Restore and display file list only** . Select if you do not know the original filenames. In this case, the files backed up in a particular session are displayed.
- **Restore and display headers only** . Select if you need specific details about backup. SQL Server header information is displayed.

Restore options



6. In the **Devices** page, select the devices to be used for the restore.
 For more information on how to select devices for a restore, see the *HPE Data Protector Help* index: “restore, selecting devices for”.
7. Click **Restore MS SQL Server** and then **Next** to select **Report level** and **Network load**.

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

8. Click **Finish** to start the restore.
 The statistics of the restore session, along with the message `Session completed successfully` is displayed at the end of the session output.

Restore options

Microsoft SQL Server database restore options

Option	Description
Backup version	Specifies the backup session from which the selected objects will be restored.

Option	Description
<p>Point-in-time restore</p>	<p>This option is only available for database objects.</p> <p>Specifies a point in time to which the database state will be restored (you also need to select Backup version and set Stop at). After recovery, the database is in the state it was at the specified date and time.</p> <p>Only transaction logs written before the specified date and time are applied to the database.</p>
<p>Stop at</p>	<p>This option is only available for database objects.</p> <p>Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable database recovery to a particular point in time, backup you restore from must be a transaction log backup.</p> <p>You cannot use this option with NORECOVERY or STANDBY. If you specify Stop at time that is after the end of RESTORE LOG operation, the database is left in a non-recovered state (as if RESTORE LOG is run with NORECOVERY).</p>
<p>Restore only this backup</p>	<p>If you restored a database version and left it in a non-operational or standby state, you can subsequently restore differential or transaction log backups one by one, leaving each version non-operational to restore additional backups.</p>
<p>Full restore of the database</p>	<p>All necessary versions are restored, including the latest full backup, the differential backup (if one exists), and all transaction log backups from the differential up to the selected version.</p>
<p>Force restore over the existing database</p>	<p>Select this option if a database with the same name but a different internal structure already exists at the target Microsoft SQL Server instance.</p> <p>If this option is not selected, the Microsoft SQL Server does not let you overwrite the existing database - the restore will fail.</p> <p>If you are restoring a data file from the PRIMARY group to an existing database, you must specify the option at the data file level.</p> <p>When using this option, make sure that the most recent logs are backed up before the restore.</p>
<p>Put database in single user mode - log off all users</p>	<p>Disconnects all users that are connected to the target Microsoft SQL Server database and puts the database in the single user mode. Note that if the database is not in the simple recovery mode, the Force restore over the existing database option should also be selected.</p>
<p>Recovery completion state</p>	<p>Enables selecting the database state after recovery. You may select from:</p> <ul style="list-style-type: none"> Leaving the database operational. Once the last transaction log is

Option	Description
	<p>restored and the recovery completed, the database becomes operational.</p> <ul style="list-style-type: none">Leaving the database non-operational after the last transaction log is restored. You may restore additional transaction logs one by one.Leaving the database in read-only mode. You may restore additional transaction logs before the database is set to read-write mode. <p>This selection is only available for database objects.</p>
Restore database with a new name	<p>This option is only available for database objects.</p> <p>Restores the database under a different name. Specify the database logical filename and the destination filename (suboptions of Restore files to new locations).</p>
Restore files to new locations	<p>Restores files to a new location. Specify the database logical filename and a destination target filename for the specified logical filename. Use this option to restore data to a different client, a different instance, or to make a database copy on the same client.</p>

TIP:

To allow different restore scenarios, you can combine general restore options, such as **Restore database to another Microsoft SQL Server** and **Restore using a different device**, with object-specific restore options, such as **Point-in-time restore**, **Recovery completion state**, **Force restore over the existing database**.

Restoring to a different SQL Server instance or/and different SQL Server

Prerequisites

- Both SQL Servers must have the same local settings (code page and sort order). This information is displayed in the session monitor for each backup.
- The target SQL Server must be configured and reside in the same Data Protector cell as the original SQL Server.
For the configuration procedure, see [Creating ZDB specifications, on page 202](#).

Procedure

- Select the databases you want to restore and their versions.
- Select the following:
 - To restore to a different SQL Server client, select **Restore to another client** and the target client from the drop-down list.
 - To restore to a different SQL Server instance, select **Restore to another instance**. If there are no instances in the drop-down list, enter the instance name by yourself.

IMPORTANT:

Make sure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.

3. Specify new database locations.
4. Start restore. See [Restore, on page 215](#).

In case of the following backup sequence:

- **F(Italic) F T(Italic)(bold) T**: To restore the version marked T, all the backup versions in italic are restored. Chain was constructed by LSN.
- **F F(Italic) T(Italic) T(Italic)(bold)**: To restore the version marked T, all the backup versions in italic are restored. Chain is constructed by time.

When Data Protector is not able to construct chain by LSNs, the old formula by time still applies as it was before. So the latest full backup, the latest differential backup and all transaction log backups after the last full or differential backup must be restored.

Instant recovery

See the *HPE Data Protector Concepts Guide* and *HPE Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

Prerequisites

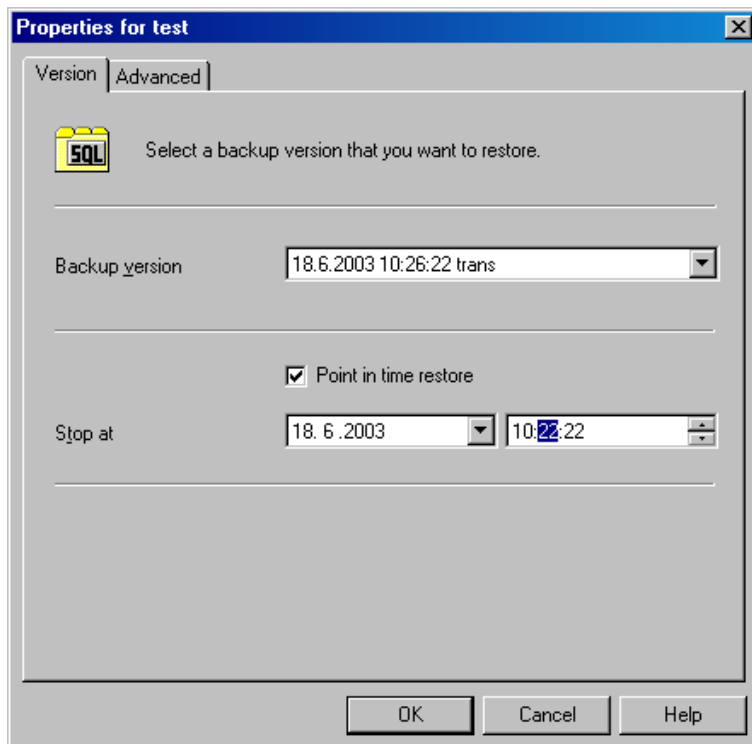
- If you restore user databases, put the databases offline:
 1. Start SQL Server Enterprise Manager.
 2. Selecting the database and click **Action**.
 3. Select all tasks and take them offline.
- If you restore system databases, put SQL Server offline by starting SQL Server Enterprise Manager, right-clicking SQL Server, and clicking **Stop**.

Perform instant recovery using the Data Protector Manager:

1. In the Context List, click **Instant Recovery**.
2. Expand **MS SQL Server** and select the backup session (replica) from which you want to restore. By default, the database will be recovered until the last backed up transaction.
3. To recover user databases to a specific point in time:
 - a. In the **Source** property page, under **Restore Objects**, right-click a database and click **Properties**.
 - b. In the **Backup version** drop-down list, select the required replica. The latest version is selected by default.

Select **Point in time restore**. From the **Stop at** drop-down list, select the point in time to which the transactions should be applied, and click **OK**. If no transaction logs are available, this option is disabled.

Point-in-Time restore

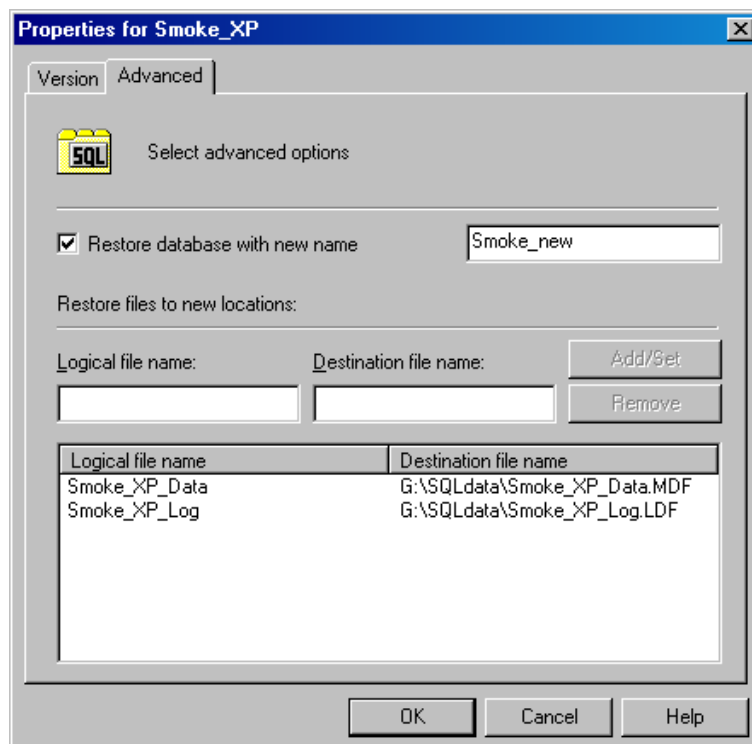


To restore the database under a different name, click **Advanced** and select **Restore database with new name**. See [Restoring database with a new name, below](#).

IMPORTANT:

If the logical filename and physical filename are not listed, add them to the list. Specify the same names as used for ZDB; otherwise, instant recovery fails.

Restoring database with a new name



4. Click **Restore MS SQL Server**.

If you restore system databases, SQL Server displays errors because its services are offline. Therefore, when restore completes, start SQL Server manually using SQL Server Enterprise Manager.

Monitoring sessions

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

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

For information on monitoring sessions, see the *HPE Data Protector Help* index: “viewing currently running sessions” and “viewing finished sessions”.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector SQL Server integration. Start at [Problems, on page 225](#). If you cannot find a solution there, perform general checks and verifications.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

For general ZDB, restore, and instant recovery related troubleshooting, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Before you begin

- Make sure that the latest official Data Protector patches are installed. For details on how to verify this, see the *HPE Data Protector Help* index: "patches".
- See the *HPE Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <https://softwaresupport.hpe.com/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Check that SQL Server services are running.
- Examine system errors reported in `debug.log` on the SQL Server client. Additionally, check `errorlog` and `VDI.log` files in the `MSSQL\log` directory.
- Make a test filesystem backup and restore of the problematic client. For information, see the *HPE Data Protector Help*.
- Check that every SQL Server used with Data Protector has the MS SQL Integration component installed.
- Connect to SQL Server via SQL Server Enterprise Manager using the same login ID as you specified in the Data Protector **Configuration** dialog box.
- Perform a database backup using SQL Server Enterprise Manager. If the backup fails, fix any SQL Server problems, and then perform a backup using Data Protector.

Additionally, if your backup failed:

- Verify the configuration file to check if the Cell Manager is correctly set on SQL Server.
- If you do not see the SQL Server instance as the application database when creating a backup specification, enter the instance name yourself. When "not-named instance" is not displayed, insert the `DEFAULT` string.
- If Data Protector reports that the integration is properly configured, verify that the SQL Server user has appropriate rights to access the required databases.

During master database restore, the following error occurs when executing an SQL statement:

```
Error has occurred while executing an SQL statement.  
Error message: 'SQLSTATE:[42000] CODE:(3108) MESSAGE:[Microsoft]  
[ODBC SQL Server Driver][SQL Server]To restore the master database,  
the server must be running in single user mode. For information on  
starting in single user mode, see "How to: Start an Instance of SQL  
Server (sqlservr.exe)" in Books Online.
```

Note that this behavior is expected when the master database is not restored in single user mode.

Problems

Problem

The integration is properly configured but the database backup fails after a timeout

- With an error similar to:

```
[Warning] From: OB2BAR@computer.company.com "SQLSRV"  
Time: 7/29/2011 8:19:22 PM  
Error has occurred while executing SQL statement.  
[Microsoft][ODBC SQL Server Driver][SQL Server]Backup or restore  
operation terminating abnormally.'  
[Critical] From: OB2BAR@computer.company.com "SQLSRV"  
Time: 7/29/11 8:19:24 PM  
Received ABORT request from SM => aborting
```

- SQL Server error log contains an entry similar to:

```
2011-07-29 20:19:21.62 kernel  
BackupVirtualDeviceSet::Initialize: Open failure on backup  
device 'Data_Protector_master'.  
Operating system error -2147024891(Access is denied.).
```

- SQL Server VDI.LOG file contains an entry similar to:

```
2011/07/30 13:19:31 pid(2112)  
Error at BuildSecurityAttributes: SetSecurityDescriptorDacl  
Status Code: 1338, x53A Explanation: The security descriptor  
structure is invalid.
```

SQL Server service and Data Protector Inet are running under different accounts. The integration cannot access SQL Server due to security problems.

Action

Restart the Data Protector Inet service under the same account as the SQL Server service is running.

Problem

Backup fails with "The object was not open"

When backing up Microsoft SQL Server databases, the session fails with an error similar to the following:

```
[Critical]From : OB2BAR_Main@wemaoldb2dr "Aolins" Time:11/12/2011 02:01:34 AM  
Microsoft SQL Server reported the following error during login : The object was not  
open
```

The error may appear if the SQL Server Browser service is not running.

Action

Proceed as follows:

1. Start the SQL Server Browser service.
2. Start a new backup session.

Problem

Backup fails if the appropriate drive letter on the backup system does not exist

Backup fails with an error, similar to:

```
[Major] From: SSEA@computer1.com "" Time: 02-Feb-11 14:07:54
      Filesystem \\.\Volume{ef58fe0e-b2b8-11db-aa08-000802804af6} could not be mounted
      to Q:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\.
```

```
([87] The parameter is incorrect. ).
```

Backup fails because SSE agent tries to mount the filesystem to the drive letter which does not exist on the backup system. The drive letter must be the same as on the application system. SSE or SMI-S agents always mount the filesystems to the same drive letters on the backup systems as if the ZDB_PRESERVE_MOUNTPOINTS option is set to 1.

Action

For source volumes to be successfully replicated, create the same drive letters on the backup system as they are used on the application system for mounting the source volumes.

Problem

Database is left in unrecovered state after “Invalid value specified for STOPAT parameter” is reported

The database remains in an unrecovered state as if the RESTORE LOG operation was run with **Leave the database non-operational**.

Action

Recover the database to the latest point in time using SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

Problem

Transaction logs cannot be restored from tape

The recovery completed successfully and the database was put in norecovery state, but transaction logs cannot be restored from tape.

Action

Recover the database to the state of ZDB to disk using the SQL Query Analyzer:

```
RESTORE DATABASE database name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

Problem

Instant recovery of SQL Server databases fails

If the SQL Server service is offline prior to the instant recovery, instant recovery of the SQL Server databases fails.

The following errors are displayed:

```
[Critical] From: computer@company.com "(DEFAULT)" Time: 4/9/2011 7:01:42 PM
Microsoft SQL Server reported the following error during login:
The object was not open.
[Warning] From: computer@company.com "(DEFAULT)" Time: 4/9/2011 7:01:42 PM
[152:9208] Data Protector is probably not configured for use with SQL Server on
this host.
```

Action

Perform one of the following:

- Set the Data Protector `omnirc` option `OB2_SQLRESTORE_STARTSRV`, which starts the SQL Server service prior to the recovery of SQL databases, to 1.

During the master database restore, the following error is displayed:

```
RESTORE master with SNAPSHOT is not supported.
```

Note that this behavior is expected. No further steps are needed after the instant recovery.

- Restart SQL Server instance services after instant recovery completes. If restarting services does not automatically start the recovery of all system databases, start the SQL Server instance in single user mode and manually start the recovery of the master database. Follow the same procedure for other system databases. At the end, restart SQL Server instance services.

Problem

Restore to another client in the Data Protector cell not configured for use with SQL Server fails

Action

Configure the SQL integration on this client (see [Configuring the integration, on page 196](#)).

Problem

Database is left in unrecovered state after restore completed successfully

If you set the time for `Stop` at beyond the end of the `RESTORE LOG` operation, the database remains in the unrecovered state as if the `RESTORE LOG` operation was run with `Leave the database non-operational`.

Action

Recover the database to the latest point in time by using the SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

Problem

Instant recovery of a Microsoft SQL Server database configured on a Microsoft Cluster Server system fails with "The physical filename may be incorrect"

Instant recovery of the Microsoft SQL Server data in an HPE Business Copy P9000 XP configuration on a Microsoft Cluster Server system fails with the following error:

```
[Microsoft][ODBC SQL Server Driver][SQL Server]Device activation error.
```

The physical file name '<Data/Log filename>' may be incorrect.

Action

Perform the following steps:

1. Using Microsoft SQL Server Enterprise Manager, detach the Microsoft SQL Server database that you want to recover.
2. Using Cluster Administrator, take the Microsoft SQL Server Disk resource offline.
3. On the application system, configure the ZDB_TAKE_CLUSRES_ONLINE option.
For details, see ???.
4. Start instant recovery.
5. When the message Please, take MS SQL cluster resources online appears in the Data Protector GUI, bring the Microsoft SQL Server cluster resources online using Cluster Administrator.

Problem

Database restore fails

The restore session aborts with a major error similar to:

```
Error has occurred while executing a SQL statement.Error message: 'SQLSTATE:[42000] CODE:(3159) MESSAGE:[Microsoft][ODBC SQL Server Driver][SQL Server]The tail of the log for the database "test2" has not been backed up. Use BACKUP LOG WITH NORECOVERY to backup the log if it contains work you do not want to lose. Use the WITH REPLACE or WITH STOPAT clause of the RESTORE statement to just overwrite the contents of the log. SQLSTATE:[42000] CODE:(3013) MESSAGE:[Microsoft][ODBC SQL Server Driver][SQL Server]RESTORE DATABASE is terminating abnormally.'
```

Action

To solve the problem perform a transaction log backup to obtain the most recent transaction logs.

Chapter 4: Data Protector Microsoft Exchange Server 2010+ ZDB integration

Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server 2010 ZDB integration, where Data Protector integrates with Microsoft Exchange Server 2010 or Microsoft Exchange Server 2013 (hereinafter, both Exchange Servers are called **Microsoft Exchange Server**, unless differences are pointed out). It describes concepts and methods you need to understand to back up and restore Microsoft Exchange Server 2010 mailbox and public folder databases or Microsoft Exchange Server 2013 mailbox databases (**databases**).

Both standalone environments and Database Availability Group (**DAG**) environments are supported.

The Data Protector Microsoft Exchange Server 2010 integration is based on the Volume Shadow Copy Service (**VSS**) technology. For details on VSS concepts, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Backup

During backup, databases can be used actively (**online backup**). In DAG environments, you can back up active and/or passive database copies.

As Microsoft Exchange Server, ZDB disk array, and VSS are involved, you can specify different kinds of backup types:

- Microsoft Exchange Server backup types
- VSS backup types
- ZDB backup types

You can select among the following Microsoft Exchange Server backup types:

- Full
- Copy
- Incremental
- Differential

You can select among the following ZDB backup types:

- ZDB-to-disk
- ZDB-to-disk+tape
- ZDB-to-tape

You can select among the following VSS backup types:

- Local or network backup
- VSS transportable

For details on the backup types, see [Backup types, on page 236](#).

Restore

You can restore Microsoft Exchange Server databases using standard restore or instant recovery.

During restore, each database can be restored using a different restore method. The following methods are available:

- Repair all passive copies with failed status
- Restore to the latest state
- Restore to a point in time
- Restore to a new mailbox database
- Restore files to a temporary location

This chapter provides information specific to the Microsoft Exchange Server 2010 integration. For additional limitations, see the *HPE Data Protector Product Announcements, Software Notes, and References*. For general Data Protector procedures and options, see the *HPE Data Protector Help*.

Integration concepts

Data Protector integrates with Microsoft Exchange Server through the Data Protector Microsoft Exchange Server integration agent, which channels communication between the Data Protector Session Manager and the clients in the Microsoft Exchange Server environment. The agent communicates with Microsoft Exchange Server through the Microsoft Exchange Management Shell and uses VSS to back up data.

Supported environments


Data Protector supports Microsoft Exchange Server Database Availability Group environments (**DAG environments**) as well as environments with standalone Microsoft Exchange Server systems (**standalone environments**).

Standalone environments

In a standalone Microsoft Exchange Server environment, each Microsoft Exchange Server system stands on its own.

In one session, you can back up databases from only one Microsoft Exchange Server system. Data Protector sends backup and restore requests directly to the Microsoft Exchange Server system.

Standalone environment (example)
Standalone environment

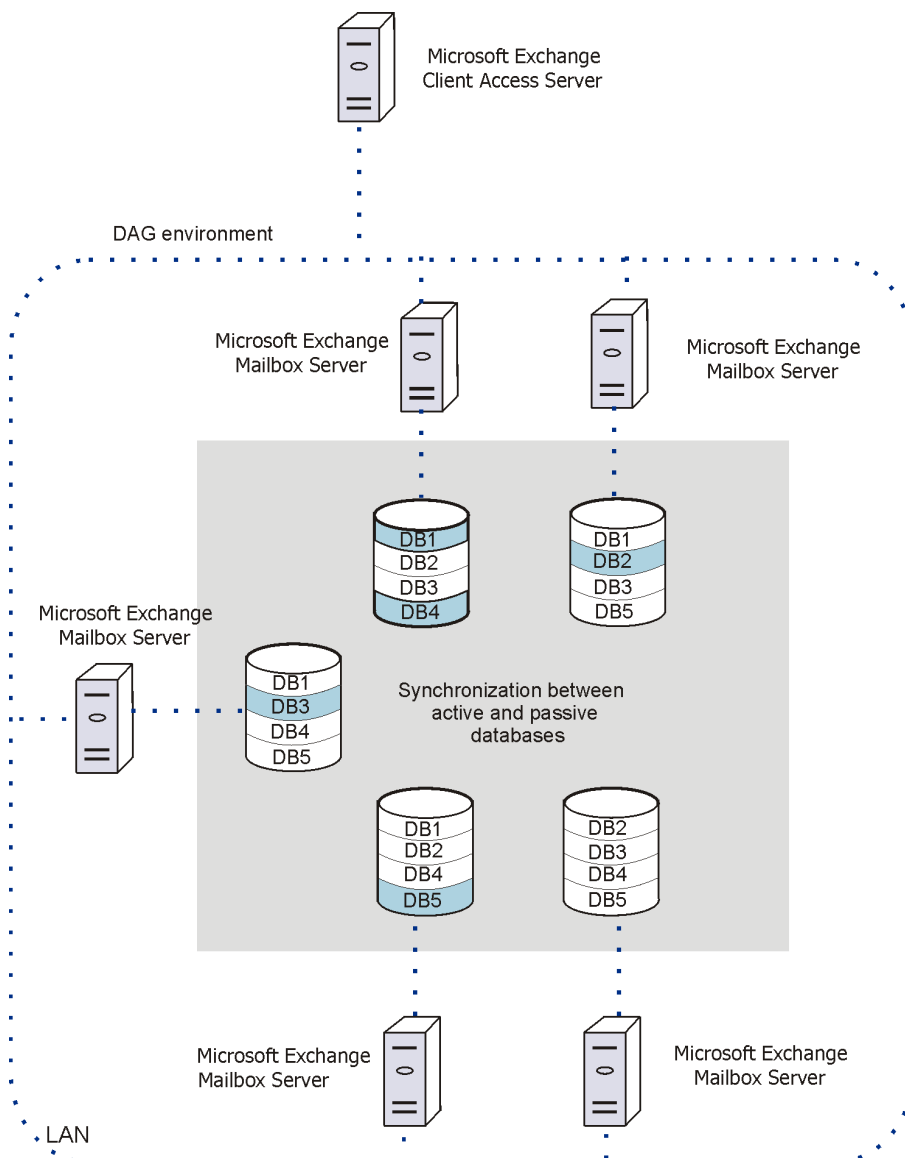


DAG environments

In a DAG environment, Data Protector communicates with the DAG using one of the Microsoft Exchange Server systems (the one that is currently active in the environment). All backup and restore requests are sent there.

In one session, you can back up active and/or passive database copies from different Microsoft Exchange Server systems that belong to the same DAG.

DAG environment (example)



In [DAG environment \(example\)](#), on the previous page, active databases are shaded in blue.

If a database has multiple passive copies, you can specify which particular passive copy you want to back up, using one of the following backup policies:

- minimize the number of hosts
- lowest activation preference
- highest activation preference
- shortest replay lag time
- longest replay lag time
- longest truncation lag time

You can also specify from which Microsoft Exchange Server systems database copies should not be backed up.

For a brief description of the activation preference number, replay lag time, and the truncation lag time, see [Microsoft Exchange Server parameters in DAG environments, below](#).

Microsoft Exchange Server parameters in DAG environments

Parameter	Description
Activation preference number	The activation preference number determines which passive copy is activated if multiple passive copies meet the same criteria; the copy assigned the lowest activation preference number is activated.
Replay lag time	The <code>ReplayLagTime</code> parameter plays a role when synchronizing a passive copy with the active copy. As soon as a log file at the active copy side is filled up, it is copied to the passive copy side. By default, the newly copied log is also applied to the passive copy database files. However, if the passive copy <code>ReplayLagTime</code> parameter is set to a value greater than 0, the log is applied with a lag, creating a lagged database copy. The maximum value is 14 days.
Truncation lag time	The <code>TruncationLagTime</code> parameter specifies how long the Microsoft Exchange Replication service waits before truncating log files that have already been applied to the database files. The maximum value is 14 days.

Configuring the integration

Prerequisites

- Ensure that you have correctly installed and configured the Microsoft Exchange Server environment.
 - For supported versions, platforms, devices, and other information, see the latest support matrices at <https://softwaresupport.hpe.com/manuals>.
 - On Microsoft Exchange Server, install .NET Framework 3.5.1.

NOTE:

On Windows Server 2012, the installation of .NET Framework 3.5.1 is done manually and not by default.

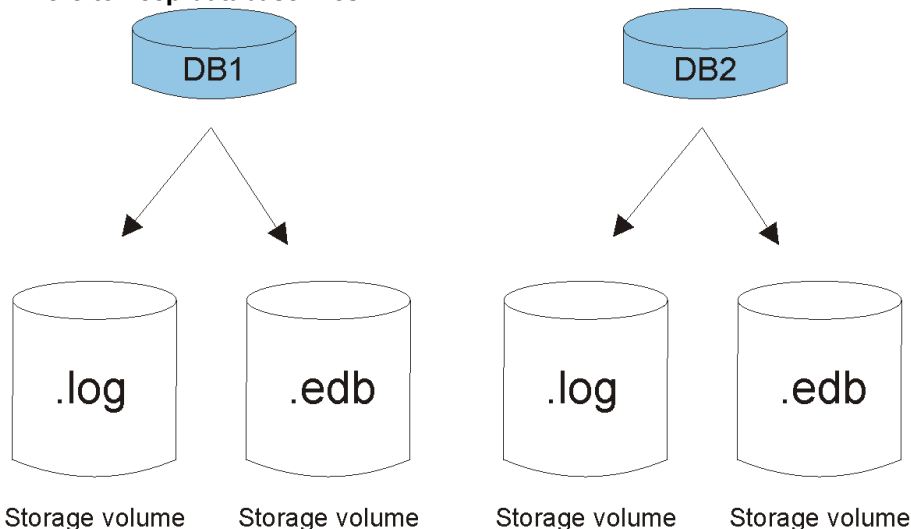
- For information on installing, configuring, and using Microsoft Exchange Server, see the Microsoft Exchange Server documentation.
- In a Microsoft Exchange Server 2010 environment, if you intend to use the restore method **Restore to a point in time**, make sure that you have Microsoft Exchange Server 2010 SP1 installed.

- If you intend to run Incremental and Differential backup sessions, make sure that circular logging is disabled.

If you plan to run instant recovery sessions, it is advisable to keep Microsoft Exchange Server databases on separate storage volumes. Also, keep a database's files (.edb and .log) on separate storage volumes. See [Where to keep database files, below](#). Such a configuration provides better restore granularity.

Ensure that storage volumes on different Microsoft Exchange Server systems are the same size. Otherwise, you may experience problems during copy-backup instant recovery sessions.

Where to keep database files



- Ensure that you have correctly installed Data Protector. On how to install Data Protector in various architectures, see the *HPE Data Protector Installation Guide*.

Ensure that the following Data Protector components are installed on all Microsoft Exchange Server systems:

- MS Exchange Server 2010+ Integration
- MS Volume Shadow Copy Integration
- The appropriate Data Protector disk array agent

NOTE:

For VSS transportable backup sessions, the MS Volume Shadow Copy Integration component and the appropriate Data Protector disk array agent must also be installed on the backup systems.

In DAG environments, the DAG virtual system (host) must also be imported to the Data Protector Cell. On how to import a client to a Data Protector Cell, see the *HPE Data Protector Help* index: "importing, client systems".

- For limitations, see "Limitations and recommendations" in the *HPE Data Protector Product Announcements, Software Notes, and References*.

Limitations

- Due to incompatibility between Microsoft Exchange Server versions, backup objects belonging to a particular Exchange Server version cannot be restored to Data Protector clients on which a different Exchange Server version is installed.

Before you begin

- Configure devices and media for use with Data Protector.
- To test whether a Microsoft Exchange Server system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore on every Microsoft Exchange Server client in your environment.

Configuring user accounts

Windows domain user account for backup and restore sessions

Backup and restore sessions are started by the Data Protector Inet service, which by default runs under the Windows local user account SYSTEM. Consequently, a backup or restore session is performed using the same user account.

However, you can specify that the Data Protector Inet service should use a different Windows domain user account to start a session:

- To perform a backup session under a different user account, specify the **Specify OS user** option (see [Specifying view type, on page 242](#)) when creating a backup specification.
- To perform a restore session under a different user account, specify the **User name** and **Group/Domain name** options in the Options page (when performing standard restore, see [Restore options, on page 265](#)) or Advanced page (when performing instant recovery, see [Instant recovery – advanced, on page 273](#)).

Before you specify a different Windows domain user account, configure the user account as follows:

1. Grant the user appropriate permissions to back up and restore Microsoft Exchange Server databases.
2. Add the user to the Data Protector admin or operator user group. For details on adding users, see the *HPE Data Protector Help* index: "adding users".
3. Save the user and its password to a Windows Registry on the Microsoft Exchange Server system on which you plan to start the integration agent (`e2010_bar.exe`). To save the user account, use the Data Protector `omniinetpasswd` or `omnicc` command.

NOTE:

The user account saved in the Windows Registry will be used by the Data Protector Inet service when needed.

For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: "Inet user impersonation".

Example

To save the user `jane` from the domain `HPE` and with the password `mysecret` to a Windows Registry, log on to the Microsoft Exchange Server system and execute the following command:

```
omniinetpasswd -add jane@HPE mysecret
```

User account for executing Exchange Management cmdlet operations

In the Microsoft Exchange Server 2013 environment, you need user credentials with specific Exchange Management Roles assigned to create a remote runspace for executing the Exchange Management cmdlet operations remotely. These operations are executed as part of Microsoft Exchange Server backup and restore operations.

Configure your user account with the following Exchange privileges:

- As a member of the Organization Management role group.
- As a member of the Discovery Management role group.
- As a member of the Administrators group of the Microsoft Exchange Server system on which the integration is installed.

Configure a valid Exchange domain user account, when creating a backup specification. The user account is saved in the user credentials specific configuration file located in the `Data_Protector_program_data\Config\Server\Integ\Config\E2010` directory and named by the domain name. The saved user credentials will be used by Data Protector when needed.

For details, see [Specifying view type, on page 242](#).

For information on the Exchange Management cmdlet operations, see the Microsoft Exchange Server documentation.

Backup

When you back up a Microsoft Exchange Server database, the following files are automatically backed up:

- database files (`.edb`)
- transaction logs (`.log`)
- checkpoint files (`.chk`)

However, depending on the Microsoft Exchange Server backup type you select, not all files are always backed up. For details, see [Microsoft Exchange Server backup types, on the next page](#).

Backup types

As Microsoft Exchange Server, ZDB disk array, and VSS are involved, you can specify different kinds of backup types:

- Microsoft Exchange Server backup types
- ZDB backup types
- VSS backup types

Microsoft Exchange Server backup types

You can select among the following Microsoft Exchange Server backup types:

Backup types

Full	<p>Backs up the database files (.edb), transaction logs (.log), and checkpoint files (.chk), and then truncates the transaction logs.</p> <p>DAG environments: If multiple copies of a database are selected for backup, Data Protector first performs a Full backup of the passive copy that has the fewest logs applied to the database file, and then performs a Copy backup of all the remaining copies, with the active copy being backed up last. The copies are backed up sequentially due to a Microsoft Exchange Server VSS writers limitation.</p>
Copy	<p>Backs up the database files (.edb), transaction logs (.log), and checkpoint files (.chk), without truncating the transaction logs.</p>
Incremental	<p>Backs up the transaction logs (.log) that have been created since the last Full or Incremental backup, and then truncates the transaction logs.</p> <p>DAG environments: If multiple copies of a database are selected for backup, Data Protector backs up the transaction logs of only one copy (one of the passive copies is selected).</p>
Differential	<p>Backs up the transaction logs (.log) that have been created since the last Full backup, without truncating the transaction logs.</p>

NOTE:

An incremental or differential backup of a database cannot be performed:

- If a full backup has not been performed.
- If an incremental backup is started just after a differential backup has been performed, or the other way around.
- If Microsoft Exchange Server circular logging is enabled.

ZDB backup types

You can select among the following ZDB backup types:

- ZDB-to-disk
- ZDB-to-disk+tape
- ZDB-to-tape

NOTE:

For Microsoft Exchange Server Incremental and Differential backup types, only the ZDB-to-tape backup type is available.

VSS backup types

You can select among the following VSS backup types:

- Local or network backup
- VSS transportable

For details, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Backup parallelism

- During a backup session, copies of different databases are backed up in parallel, however, copies of the same database are not, due to a Microsoft Exchange Server VSS writers limitation.
- If multiple backup sessions that intend to back up the same database are started in parallel, only the session that first locks the database can back up the database; the other sessions cannot. In DAG environments, this also applies if backup sessions intend to back up different copies of the same database; only the session that first locks the database (that is, all its copies) can back up the database copies; the other sessions cannot.

NOTE:

This behavior ensures that the construction of a restore chain is valid. For example, suppose that several full backup sessions that intend to back up the same database are started in parallel. If all the sessions backed up the database, it might happen that the session given the latest Session ID is not the one that backed up the database last. For details on restore chains, see [Restore chain, on page 256](#).

Replica rotation in DAG environments

Data Protector Microsoft Exchange Server 2010 integration enables you to perform ZDB sessions in multi-system environments (DAG environments). This brings some changes to the existing replica rotation functionality.

In standalone environments, the replica rotation functionality works as it used to; it limits the number of backups that are kept in the Data Protector IDB database for instant recovery purposes. For example, if the **Number of replicas rotated** option is set to 1, only one backup session is available for instant recovery at a time. If you start another backup session (with the same backup specification), backup storage volumes created in the previous backup session are deleted before new ones are created and the previous session is removed from the Instant Recovery context.

In DAG environments, a database can be backed up from different systems in different sessions. This introduces the changes. For example, you want to back up the database DB1, which is active on `node1.company.com`, and the database DB2, which is active on `node2.company.com`. Suppose the **Number of replicas rotated** option is set to 1 and your backup policy is such that the active copy is always backed up. After the backup, the Data Protector VSSDB database contains the following entries:

Backup 1 (2013/10/05-1):

- 2013/10/05-1:node1.company.com (containing DB1)
- 2013/10/05-1:node2.company.com (containing DB2)

Suppose that a failover occurs and the database DB1 becomes active on node2.company.com. You start another backup session. Now both databases are backed up from node2.company.com. As a result, the VSSDB database contains the following entries:

Backup 2 (2013/10/05-2):

- 2013/10/05-1:node1.company.com (containing DB1)
- 2013/10/05-2:node2.company.com (containing DB1 and DB2)

The entry 2013/10/05-1:node2.company.com is no longer in the VSSDB database as it has been rotated out (the corresponding backup storage volumes have been deleted) due to the replica rotation functionality.

NOTE:

Entries are rotated per system and not per session. As a result, entries that are created in the same session can be rotated out at different points in time (that is, in different sessions).

Now another failover occurs and both databases become active on node1.company.com. You start another session and both databases are backed up from node1.company.com. As a result, the VSSDB database contains the following entries:

Backup 3 (2013/10/05-3):

- 2013/10/05-3:node1.company.com (containing DB1 and DB2)
- 2013/10/05-2:node2.company.com (containing DB1 and DB2)

Note that the entry 2013/10/05-1:node1.company.com has been rotated out as well. Since both parts created in the session 2013/10/05-1 have been rotated out, the session 2013/10/05-1 can no longer be used for instant recovery (it is removed from the Instant Recovery context).

If you go to the Instant Recovery context and select the session created in Backup 1 after you have performed Backup 2, the source page incorrectly shows that both databases can be restored. The entry 2013/10/05-1:node2.company.com (containing DB2) has been rotated out in Backup 2. If you select DB2 for restore and start instant recovery, the session fails. Therefore, make sure to restore databases from sessions for which the necessary entries in the VSSDB database still exist.

Backup considerations

- *Backup strategy:*

Choose one of the following strategies to back up your data:

- Full
- Full, Incremental, Incremental, ...
- Full, Differential, Differential, ...
- Full, Copy, Incremental, ..., Copy, Incremental, ...

IMPORTANT:

An Incremental backup session cannot be followed by a Differential backup session, nor the other way around. You must first run a Full backup session.

- *Active copies as opposed to passive copies:*

There is no difference between the active and passive copy, except in the currently active log file (at the active copy side), which is not copied to the passive copy side until the file is filled up (that is, reaches 1 MB). Consequently, if you back up a passive copy, the transactions in the currently active log file are not included.

- *Lagged database copies:*

Backing up a lagged database copy is equivalent to backing up a non-lagged database copy. If you restore from the backup of a lagged database copy, files are not only restored, but logs are also applied to the database file, returning the database to its most recent state. However, restoring the logs and applying them to the database file is time-consuming and, therefore prolongs the restore session. Also note that you need enough disk space to restore all the necessary logs.

On the other hand, restoring from the backup of a lagged database copy enables you to restore the database to a point in time before the backup was taken. Restore the database without performing database recovery and mounting. Then remove unwanted logs, and finally recover and mount the database.

- *Public folders :*

In the Microsoft Exchange Server 2010 environment, backup of Microsoft Exchange Server public folders with activated replication is not supported.

- *Concurrent backup sessions:*

Backup sessions that back up the same database cannot run in parallel.

Object operations considerations

- Object copy and object verification

When copying or verifying Microsoft Exchange Server objects you need to select all Data Protector backup objects created in the same session. To make sure that you do not select only a few objects from the session, the Data Protector GUI does not list Microsoft Exchange Server backup objects for interactive object copy or object verification sessions in the Objects scope of the Object Operations context.

Use the Session or the Media scope instead.

Creating backup specifications

Create a backup specification using the Data Protector GUI (**Data Protector Manager**).

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS Exchange 2010+ Server**, and click **Add Backup**.
3. In the Create New Backup dialog box, specify **Backup type** (VSS backup type). For details, press **F1**. Click **OK**.

4. In **Application system**, select the Microsoft Exchange Server system that you want to back up. In a DAG environment, select the DAG virtual system or a Microsoft Exchange Server system.

NOTE:

The **Application system** drop-down list contains all clients that have the Data Protector MS Exchange Server 2010+ Integration component installed. In a DAG environment, the list contains also the DAG virtual system (host).

The backup session (that is, the integration agent `e2010_bar.exe`) will be started on the client that you specify here. If you select a DAG virtual system, the integration agent is started on the currently active Microsoft Exchange Server node. To find out which node is currently active, see [Restore, on page 252](#).

NOTE:

In the Microsoft Exchange Server 2010 environment, to back up public folders residing on a Microsoft Exchange Server system that is a part of a DAG environment, select the Microsoft Exchange Server system and not the DAG virtual system (host). If you select the DAG virtual system, you can back up only databases that belong to the DAG. The Microsoft Exchange Server public folders database is not the part of it.

Depending on the VSS backup type you selected, specify the following:

- If you selected **Local or network backup**, in **Provider**, select **Use hardware provider**.
- If you selected **VSS transportable backup**, specify **Backup system**.

Specify ZDB-specific options.

For details, press **F1** or see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

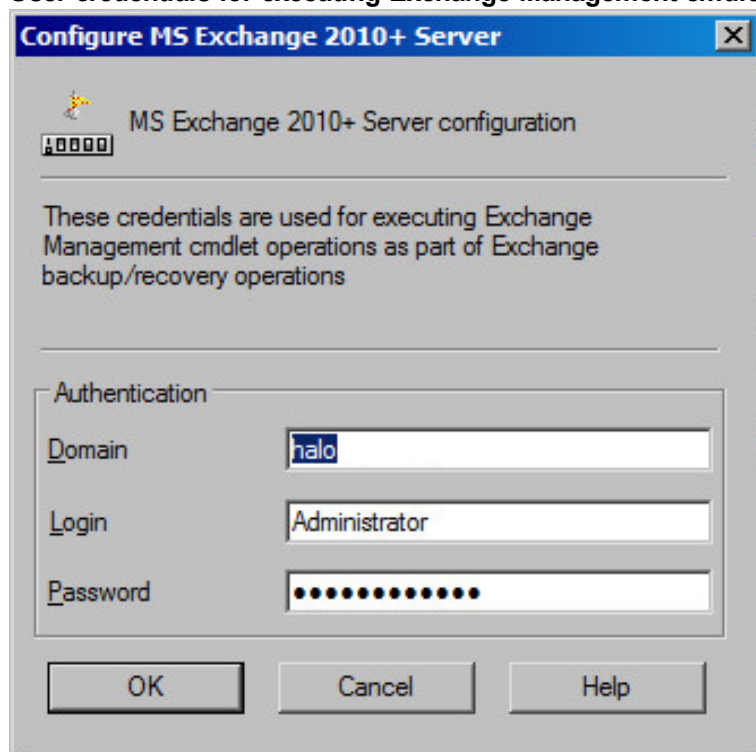
NOTE:

In an HPE P9000 XP Disk Array Family resync mode environment, the maximum number of replica storage volumes (S-VOL) that can be created for a given storage volume (P-VOL) is limited by the hardware provider configuration — MU range (maximum is 3). To enable performing of incremental and differential sessions, the **Number of replicas rotated** option in the backup specification must be set to one less than the MU range. Thus, one replica storage volume is kept free for incremental and differential backup sessions.

Click **Next**.

5. In the **Configure MS Exchange 2010+ Server** dialog box, provide the domain, username, and password to browse or backup or recover the Exchange server.

User credentials for executing Exchange Management cmdlet operations




Click **OK**.

6. The MS Exchange 2010+ Server is configured. Exit the GUI or proceed with creating the backup specifications at [Step 7](#).
7. If you selected the DAG virtual system (host), specify **View Type** to define how Microsoft Exchange Server databases should be organized in the next page (source page):

By Role	All databases in the DAG are displayed.
By Client	All clients in the DAG are displayed, together with all the databases (active or passive) residing on them. Active databases have the label (active) appended at the end. Passive databases have no label.

Specifying view type



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

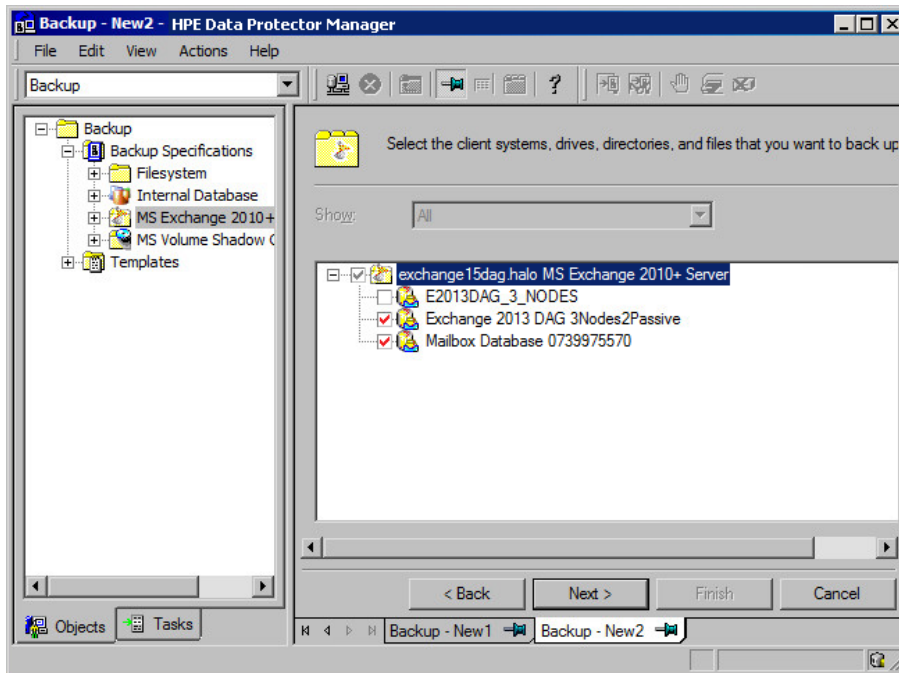
NOTE:
If no valid user credentials for executing the Exchange Management cmdlet operations remotely are specified, the Microsoft Exchange Server configuration dialog box is displayed.

8. Select which Microsoft Exchange Server databases to back up.


NOTE:
DAG environments:
In a single session, you can back up one of the following:

- multiple databases, but only one copy of each
- a single database, but multiple copies of it


Selecting databases (DAG environment – by role)



Selecting databases (DAG environment – by client)




Selecting databases (standalone environment)



9. The following applies in DAG environments if you selected the **By Role** view type. Specify the backup policy options.

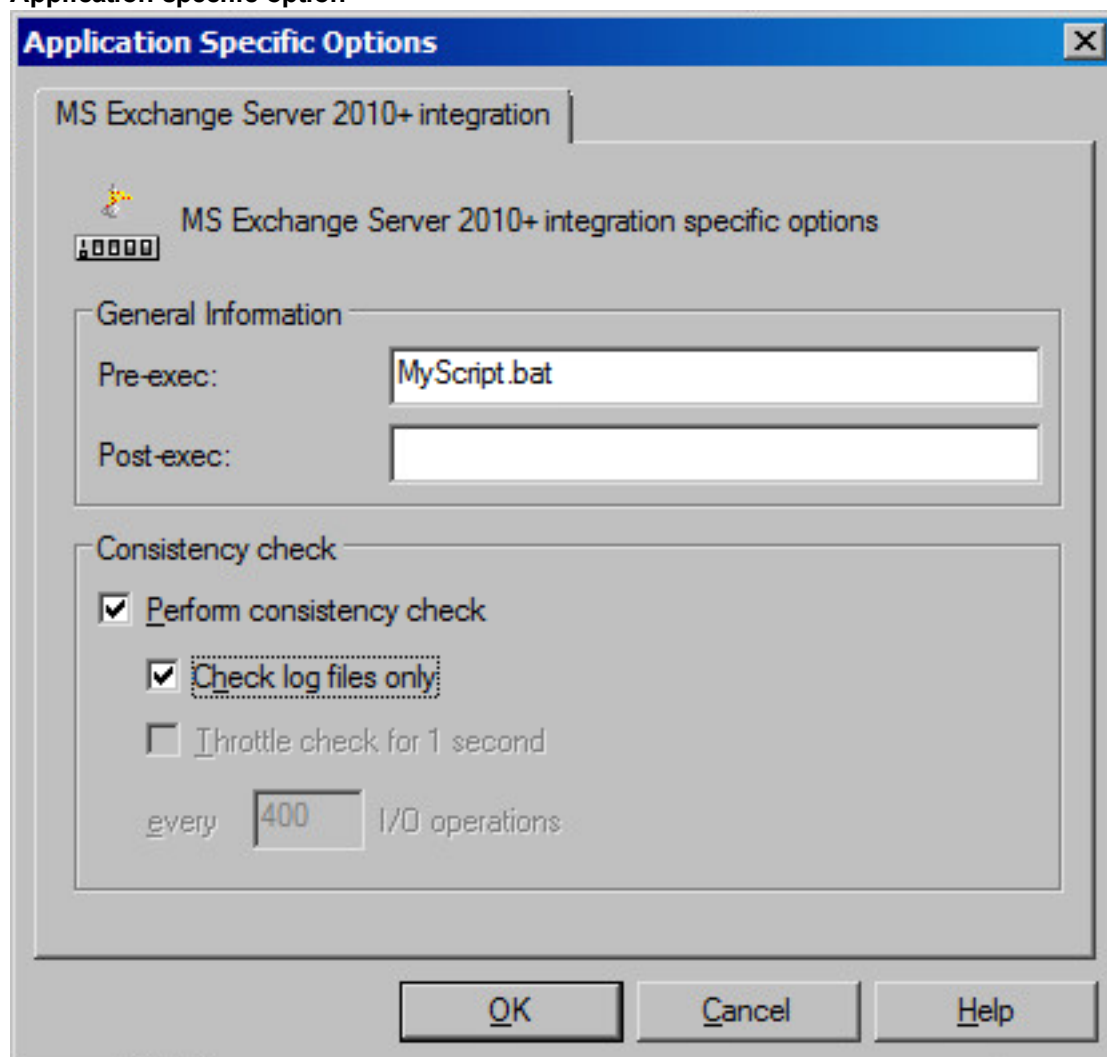
Backup policy options



For details, see [Backup policy options, on the next page.](#)

10. Select the devices to use for the backup.
To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and the media pool to use.
Click **Next**.
11. Set backup options.

Application-specific option



For information on application-specific backup options, see [Application-specific backup options, on page 248.](#)

Click **Next**.

12. Click **Save As** to save the backup specification, specify a name and a backup specification group. Optionally, you can click **Save and Schedule** to save, and then schedule the backup specification. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

TIP:
 Preview your backup specification before using it for real. See [Previewing backup sessions, on page 250](#).

Backup policy options

Options	Description	
Back up active database	If this option is selected, the active copy is backed up.	
Back up passive copy	If this option is selected, a passive copy is backed up. If a database has multiple passive copies, specify which particular copy you want to back up, using one of the following policies:	
	minimize the number of hosts (default)	If this option is selected, the minimum number of clients is involved in the backup. For example, if databases to be backed up have each a passive copy residing on the same client, they are all backed up from this client (and not one database from one client and another database from another client).
	lowest/highest activation preference	If this option is selected, the database copy with the lowest/highest activation preference number is backed up.
	shortest/longest replay lag time	If this option is selected, the database copy with the shortest/longest replay lag time is backed up.
	longest truncation lag time	If this option is selected, the database copy with the longest truncation lag time is backed up.
	For a brief description of the activation preference number, replay lag time and transaction lag time parameters, see Microsoft Exchange Server parameters in DAG environments, on page 233 . For details, see the Microsoft Exchange Server documentation.	
	Use active copy for backup if no passive copies are available	Available if Back up passive copy is selected. If this option is selected, the active copy is backed up when no passive copy is available.
Back up all copies	Available if only one database is selected for backup. If this option is selected, all copies (active and passive) are backed up. This is useful when you create ZDB-to-disk or ZDB-to-disk+tape backups (that is, backups that can be used for instant recovery). If multiple copies are backed up, during instant recovery, multiple copies can be restored, as each copy has its own	

Options	Description
	<p>replica storage volumes to be restored from. For details, see Instant recovery in DAG environments, on page 256.</p> <p>When you create a ZDB-to-tape backup, it is enough that a single copy is backed up; you can restore different copies of a database from the ZDB-to-tapebackup of a single copy.</p>
Exclude clients from backup	Creates a list of clients. The database copies that reside on these clients are not backed up.

Application-specific backup options

Options	Description
Pre-exec , Post-exec	<p>Specifies which command line to run on a Microsoft Exchange Server system before (<i>pre-exec</i>) or after (<i>post-exec</i>) the backup.</p> <p>The command line is executed only on the Microsoft Exchange Server system on which the backup session is started (that is the system on which the Data Protector Microsoft Exchange Server integration agent <code>e2010_bar.exe</code> is started).</p> <p>Type only the name of the command and ensure that the command is located in the default Data Protector commands directory on the same system. Do not use double quotes.</p> <p>DAG environments: If you selected the DAG virtual system (host) in the Application system option, ensure that the command is located on the currently active node. To find out which Microsoft Exchange Server node is currently active, see Restore, on page 252.</p>
<p>Perform consistency check</p> <p><code>[-exch_check</code> <code>[-exch_throttle <i>Value</i>]</code> <code>-exch_checklogs]</code></p>	<p>If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed.</p> <p>The check is performed on the replica storage volumes after the backup data is created. If the data is found corrupt, the replica storage volumes are discarded and the database backup fails.</p> <p>Default: selected</p>

Options	Description
	<p>If the Check log files only option is selected, only the backup data of the log files is checked, which is enough for Microsoft Exchange Server to guarantee data consistency.</p> <p>Default: selected</p> <p>By default, the consistency check is I/O intensive, which can negatively affect disk performance. The Throttle check for 1 second option throttles down the consistency check of the database file .edb to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second.</p> <p>This option is not available if only the log files are checked.</p> <p>Default: not selected</p>

Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context.

In the Microsoft Exchange Server 2013 environment, in the Source page, you can change the Exchange domain user credentials for executing the Exchange Management cmdlet operations remotely by right-clicking the selected backup object and clicking **Configure**. You can also validate your configuration by clicking **Check configuration**.

Click other desired tabs, and apply the changes.

NOTE:

To see all databases in the source page, select **All** in the **Show** option. In a DAG environment, this not only shows all databases, but also updates the current status of databases (active or passive).

To see only the databases you selected for backup or excluded from backup, select **Selected** in the **Show** option. If any databases are not displayed, it means that they are not excluded from the backup specification and will be backed up. When a new database is added on the client selected for backup, it will be automatically included to the backup specification.

Scheduling backup sessions

You can schedule a backup session to start automatically at specific times or periodically. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

Previewing backup sessions

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

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS Exchange 2010+ Server**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

1. Log in to the Cell Manager or to any client with the Data Protector User Interface component installed, under a user account that is configured as described in [Configuring user accounts, on page 235](#).
2. Execute the following command:

```
omnib -e2010_list BackupSpecificationName -test_bar
```

What happens during the preview?

The following are tested:

- Communication between the Microsoft Exchange Server system on which the backup session is started and the Cell Manager
- If each selected database has at least one copy available for backup after the **Backup policy** options and **Client filtering** options have been applied (this applies to backup specifications that contain backup policy options)
- If the selected databases are ready to be backed up (that is, they should not be dismounted, suspended, or in a failed state)

Starting backup sessions

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

To start a backup, use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS Exchange 2010+ Server**. Right-click the backup specification you want to use and click **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

1. Log in to the Cell Manager or to any client that has the Data Protector User Interface component installed under a user account that is configured as described in [Configuring user accounts, on page 235](#).

2. Execute the following command:

```
omnib -e2010_list BackupSpecificationName [-barmode E2010Mode] [LIST_OPTIONS]
```

where *E2010Mode* is one of the following:

```
full|copy|incr|diff
```

The default is `full`.

For *ListOptions*, see the `omnib` man page or the *HPE Data Protector Command Line Interface Reference*.

Examples

To start a Full backup using the backup specification `MyDatabases`, execute:

```
omnib -e2010_list MyDatabases -barmode full
```

To start a Differential backup using the same backup specification, execute:

```
omnib -e2010_list MyDatabases -barmode diff
```

Backup objects

For each database (copy), Data Protector creates the following backup objects:

- *Database file object*
 - *ClientName* :/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/DBID/File [MSVSSW-APP]
(standalone database or active copy)
 - *ClientName* :/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/DBID/File [MSVSSW-APP]
(passive copy)
- *Log file object*
 - *ClientName* :/Microsoft Exchange Writer (Exchange Information Store)/Microsoft Information Store/DBID/Logs [MSVSSW-APP]
(standalone database or active copy)
 - *ClientName* :/Microsoft Exchange Writer (Exchange Replication Service)/Microsoft Information Store/DBID/Logs [MSVSSW-APP]
(passive copy)
- *Database object*

```
ClientName :/DBID/DBName [E2010]
```

The database object contains information needed to construct the restore chain. For details on restore chains, see [Restore chain, on page 256](#).

- *VSS metadata object*

```
/BackupSession/Metadata [MSVSSW-APP]
```

Information on whether the objects were successfully backed up or not is saved in the Data Protector IDB. On how to retrieve the information from the IDB, see [Finding information for restore, on page 257](#).

Restore

You can restore Microsoft Exchange Server data by performing a standard restore or instant recovery session.

For details, see [Standard restore, on page 258](#) and [Instant recovery, on page 268](#).

IMPORTANT:

After you have restored a database, start a Full backup session for the database. Otherwise, the subsequent Incremental and Differential backup sessions will fail.

Considerations

- A Microsoft Exchange Server database that was backed up using the Data Protector Microsoft Volume Shadow Copy Service integration cannot be restored using the Data Protector Microsoft Exchange Server 2010 integration, nor the other way round.

Restore methods

There are various reasons for restoring a Microsoft Exchange Server database. Here are some examples:

- The database has become corrupt.
- The synchronization between an active and passive database copy is broken, but you want to avoid reseeding the passive copy, or simply the resume operation does not work.
- The database needs to be restored to a different point in time.
- The database's backup data needs to be restored for investigation purposes.
- The database's backup data needs to be restored to a recovery database in order to extract individual mailboxes or mailbox files.
- The database's backup data needs to be restored to a dial tone database.

To suit your needs, Data Protector offers different restore methods. You can choose among the following:

- **Repair all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**

- **Restore to a new mailbox database**
- **Restore to a temporary location**

You can specify different restore methods for different databases in the same session.

NOTE:

The first three methods restore backup data to the original database and are therefore only available if the original database still exists. The last two methods restore backup data to a new location.

Repair all passive copies with failed status

This method is available only for databases that are part of a DAG. It is useful if some of a database's passive copies become corrupt, acquiring the status `Failed` or `FailedAndSuspended`. The method automatically restores all the corrupt passive copies from the backup created in the last backup session (and the corresponding restore chain). After the data is restored, the copies are synchronized with the active copy, provided that the **Resume database replication** option is selected.

Restore to the latest state

This method is used to restore a corrupt database to the latest possible point in time. Data Protector restores the database from the backup created in the last backup session (and the corresponding restore chain). For details, see [Restore chain, on page 256](#).

Once the files are restored, all the logs (not only those restored from the backup, but also any existing logs) are replayed to the database file.

NOTE:

DAG environments:

When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `ReplayLagTime` parameter setting.

Restore to a point in time

This method is used to restore a database to a specific point in time.

NOTE:

Standard restore:

When you restore a standalone database or active copy, the existing `.log` and `.chk` files are renamed (a `.keep` extension is added to their names). This feature is useful when you restore files without performing database recovery. It enables you to apply additional logs to the database file; just delete the `.keep` extension of the log files that you also want to be applied and start a database recovery manually. In this way, you can fine-tune the point in time the database is restored to.

When you restore a passive copy, the existing files are deleted.

Once the files are restored, the logs are replayed to the database file (`.edb`) if the **Perform database recovery** option is selected.

NOTE:***DAG environments:***

- When a passive copy is restored, Microsoft Exchange Server ensures that the logs are replayed to the database file in accordance with the `RepLayLagTime` parameter setting.
- For passive copies that are not restored, a full reseed is required once the restore session completes.

Restore to a new mailbox database

This method is used to restore data to a different database, either because the original database no longer exists or in order to move the data elsewhere.

Using it, you can restore data also to a Microsoft Exchange Server recovery database.

NOTE:***Instant recovery:***

This option is not available for replica types whose data can only be restored to the original storage volumes.

Restore files to a temporary location

Using this method, you can restore database files to a location of your choice.

- When you restore from a Differential or Incremental backup session, you can restore the complete restore chain or only the files (.log) backed up in the selected session.
- When you restore data from a Full backup session, you have an option to restore only the database file (.edb).

NOTE:***Instant recovery:***

This option is not available for replica types whose data can only be restored to the original storage volumes.

Restore destination

Backup data can be restored:

- to an existing database (standalone database, active copy, passive copy),
- to a new database,
- to a temporary location.

Restoring to a standalone database

Restore to the original standalone database (standalone environment) progresses as follows:

1. The database is dismounted.
2. Backup data is restored.

3. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file .edb and the database is mounted.

To restore to the original standalone database, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

Restoring to an active copy

Restore to the active copy (DAG environment) progresses as follows:

1. The database is dismounted.
2. All replications are suspended.
3. Backup data is restored.
4. Optionally, the newly-restored logs (and pre-existing ones if you are performing the **Restore to the latest state** method) are replayed to the database file .edb and the database is mounted.

To restore to the active copy, use one of the following restore methods:

- **Restore to the latest state**
- **Restore to a point in time**

Restoring to a passive copy

Restore to a passive copy (DAG environment) progresses as follows:

1. The replication is suspended.
2. Backup data is restored.
3. Optionally, the replication with the active copy is resumed.

To restore to a passive copy, use one of the following restore methods:

- **Restore all passive copies with failed status**
- **Restore to the latest state**
- **Restore to a point in time**

Restoring data to a new database

Restore to a new database progresses as follows:

1. A new mailbox database is created.
2. Backup data is restored to the new database.

NOTE:

If you restore to a recovery database, first the backup data is restored and then a recovery database is created.

To restore data to a new mailbox database or recovery database, use the **Restore to a new mailbox database** restore method.

Restoring data to a temporary location

You can restore the database file (.edb and/or .log and/or .chk) to a client and directory of your choice. Select the **Restore files to a temporary location** restore method.

Instant recovery in DAG environments

When you back up a database in a DAG environment, you can decide whether to back up all its copies or only a single copy. If all copies are backed up, during instant recovery, all copies can be restored, as each copy has its own replica storage volumes to be restored from. If only a single copy is backed up, note the following:

- In most cases, only one database copy can be restored from the backup of a single database copy. Some replica types are directly connected with the source (*dependent* replica types) while others are *independent*, allowing the data to be restored to a different location. With the latter, you can restore either the original or a different database copy.

NOTE:

The following replica types are independent:

- HPE P9000 XP Disk Array Family (split mirror replica type in the VSS compliant mode)
- HPE P6000 EVA Disk Array Family (snapclone replica type)

For dependent replica types, Data Protector automatically grays out those clients in the **Target Nodes** option whose database copies were not backed up, because these copies cannot be restored.

- The only replica type that enables you to restore multiple database copies from the backup of a single database copy is snapclone (only with the HPE P6000 EVA Disk Array Family). However, you must also ensure that both the **Restore using HPEP6000 EVA SMI-S** and **Copy replica data to the source volumes** instant recovery options are selected, in which case data from the replica storage volumes is sequentially copied to multiple locations, restoring one database copy after the other.

Restore chain

By default, when you select a Differential or Incremental backup session for restore, Data Protector restores not only the logs (.log) backed up in the selected session but also files backed up in preceding sessions (**restore chain**):

- If a Differential backup session is selected, Data Protector restores:
 1. The .edb file and .log files backed up in the Full or Copy backup session.
 2. The .log files backed up in the selected Differential backup session.
- If an Incremental backup session is selected, Data Protector restores:
 1. The .edb file and .log files backed up in the Full or Copy backup session.
 2. The .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session.

- If a Full or Copy backup session is selected, Data Protector restores the .edb file and .log files backed up in the selected session.

NOTE:

- If the **Restore to the latest state** method is used, .log files from the Full or Copy backup session are not restored.
- The only method that enables you to restore only .log files backed up in the selected Incremental or Differential session is **Restore to a temporary location**.

Restore chain during instant recovery

During an instant recovery session, you first select which:

- Full (ZDB-to-disk or ZDB-to-disk+tape) session or
- Copy (ZDB-to-disk or ZDB-to-disk+tape) session

to use for instant recovery. From database specific options you then specify whether additional logs should also be restored, by selecting a subsequent Incremental or Differential session in the **Restore additional logs until** option.

In an instant recovery session Data Protector restores:

1. The .edb file and .log files backed up in the selected Full or Copy backup session.
2. The .log files backed up in the selected Differential backup session or in all subsequent Incremental backup sessions, up to the selected Incremental backup session.

Restore parallelism

If device concurrency allows, database copies are restored in parallel, except in the following cases:

- If database copies were backed up from the same client, but are now restored to different clients.
- If backup data of the same database copy is used as a restore source for multiple database copies.

Finding information for restore

You can retrieve information about backup sessions (such as information on the backup type and media used, and the messages reported during the backup) from the Data Protector IDB.

To retrieve information, use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Internal Database**.
2. In the Scoping pane, expand **Objects** or **Sessions**.

If you expand **Objects**, backup objects are sorted according to the Microsoft Exchange Server databases for which they were created.

NOTE:

The backup object name contains the database GUID. To find out which GUID belongs to

which database, see the *database object/DB_GUID/DB_Name*.

For example, the *database object* for the database DB1 with the GUID 08bca794-c544-4e27-87e8-533fb81fd517 is:

```
/08bca794-c544-4e27-87e8-533fb81fd517/DB1
```

If you expand **Sessions**, backup objects are sorted according to the sessions in which they were created. For example, backup objects created in the session 2013/02/7-7 are listed under 2013/02/7-7.

To view details on a backup object, right-click the backup object and click **Properties**.

TIP:

To view the messages reported during the session, click the **Messages** tab.

Using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the Data Protector User Interface component installed under a user account that is configured as described in [Configuring user accounts, on page 235](#).

2. Get a list of Microsoft Exchange Server backup objects created in a backup session:

```
omnidb -session SessionID
```

3. Get details on a backup object:

```
omnidb -e2010 BackupObjectName -session SessionID -catalog
```

Here is one example of a backup object name:

```
devy.company.com:/08bca794-c544-4e27-87e8-533fb81fd517/DB1
```

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

Standard restore

Standard restore is restore from backup data residing on Data Protector media (for example, a tape). Such data is created in ZDB-to-disk+tape and ZDB-to-tape sessions.

You can restore multiple Microsoft Exchange Server databases in the same standard restore session, specifying a different restore method for each database. For details, see [Restore methods, on page 252](#).

To perform a standard restore, use the Data Protector GUI or CLI.


Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Exchange 2010+ Server**, expand the DAG virtual system or standalone Microsoft Exchange Server system and click **MS Exchange 2010+ Server**.
3. In the Source page, Data Protector displays all Microsoft Exchange Server databases backed up from the selected DAG or standalone environment.

Select the Microsoft Exchange Server databases to restore.


When you select a database, the Properties for Database dialog box is displayed automatically. Specify a restore method and click **OK**. For databases that are part of a DAG, the default restore method is **Repair all passive copies with failed status**. For standalone databases, the default is **Restore to the latest state**.

Selecting databases for restore



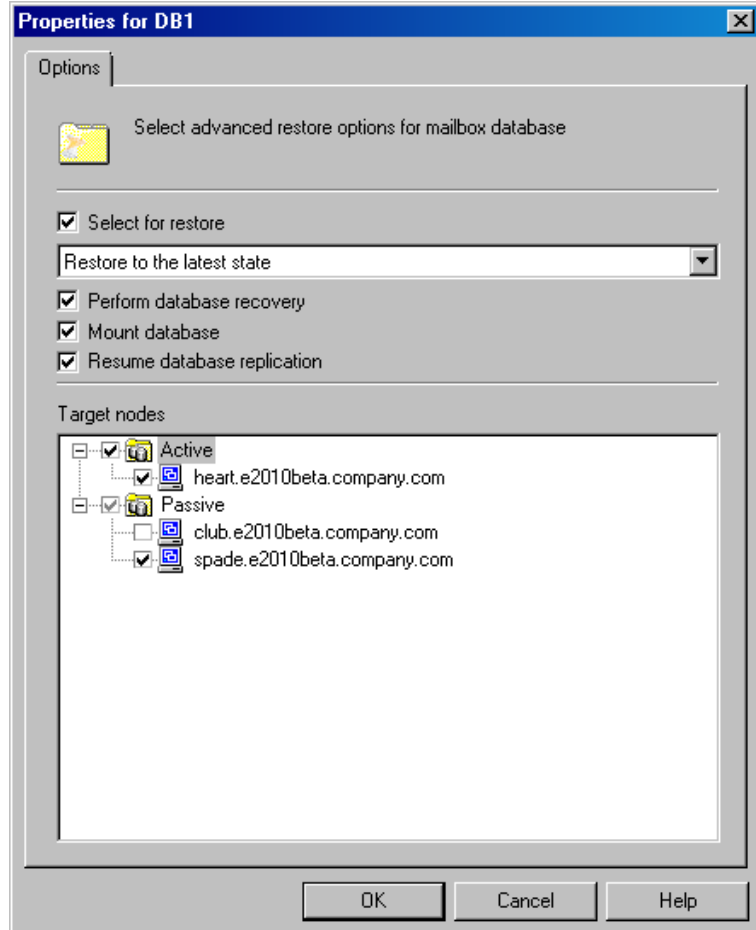
To change the restore method, right-click the database and click **Properties**.

Repair all passive copies with failed status



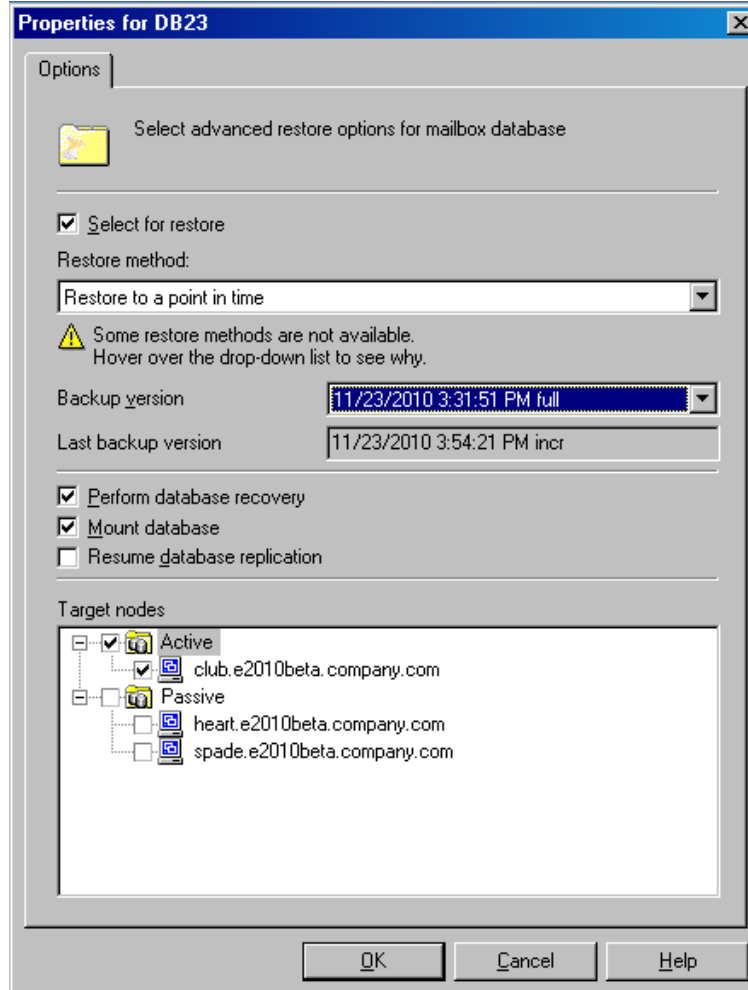
For details, see [Repair all passive copies with failed status](#), on page 276.

Restore to the latest state




For details, see [Restore to the latest state, on page 253](#).

Restore to a point in time




For details, see [Restore to a point in time, on page 253](#).



For details, see [Restore to a new mailbox database, on page 254](#).


Restore files to a temporary location



For details, see [Restore files to a temporary location, on page 254](#).

4. In the **Options** page, specify the Data Protector Microsoft Exchange Server 2010 integration restore options. For details, see [General restore options, on page 280](#).

Restore options



5. In the **Devices** page, select the devices to use for restore.
For details on how to select devices to be used for restore, see the *HPE Data Protector Help* index: “restore, selecting devices for”.
6. Click **Restore**.
7. In the **Start Restore Session** dialog box, click **Next**.
8. Specify **Report level** and **Network load**.

NOTE:

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

9. Click **Finish** to start the restore.
The statistics of the restore session, along with the message `Session completed successfully` is displayed at the end of the session output.

Restoring using the Data Protector CLI

1. Log in to the Cell Manager or to any Microsoft Exchange Server client with the User Interface component installed under a user account that is configured as described in [Configuring user accounts, on page 235](#).
2. Execute the following:

```
omnir -e2010  
-barhost ClientName  
[VSS_EXCHANGE_SPECIFIC_OPTIONS]  
Database [Database ...]  
[-user User:Domain]
```

```
[GENERAL_OPTIONS]
```

Database

```
{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID }
[-source SourceClientName]
{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS E2010_REPAIR_
METHOD_OPTIONS
[-no_resume_replication]
```

E2010_LATEST_METHOD_OPTIONS

```
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]
```

E2010_PIT_METHOD_OPTIONS

```
-session BackupID
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]
```

E2010_NEW_METHOD_OPTIONS

```
-session BackupID
-client TargetClientName
-location TargetDatabasePath
-name TargetDatabaseName
[-recoverydb]
[-no_recover]
[-no_mount]
```

E2010_TEMP_METHOD_OPTIONS

```
-session BackupID
-client TargetClientName
-location TargetDatabasePath
[-no_chain]
[-edb_only]
[-no_recover]
```

For a brief description of the options, see [Restore options, on page 276](#). For details, see the `omnir` man page or the *HPE Data Protector Command Line Interface Reference*.

NOTE:

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy

session still exists. To restore the objects, you must specify the session ID of the original backup session (that is, the backup ID) and not the session ID of the object copy session.

The omnir syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

Example (Restore method – repair)

DAG environment

To restore all corrupt passive copies of the database DB1, which was backed up from a DAG whose virtual system name was `dag0.company.com`, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, execute:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source dag0.company.com -repair
```

Example (Restore method – latest)

Standalone environment

To restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`, to the latest possible point in time, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange2.company.com`, execute:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source exchange1.company.com -latest
```

DAG environment

Suppose you want to restore the active copy of the database DB1, which resides on the client `exchange1.company.com`, and the passive copies of the database that reside on the clients `exchange2.company.com` and `exchange3.company.com`. Suppose the database DB1 is part of a DAG whose virtual system name is `dag0.company.com`, and that you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange2.company.com`. Execute the following command:

```
omnir -e2010 -barhost exchange2.company.com -db_name DB1 -source dag0.company.com -latest -node exchange1.company.com -node exchange2.company.com -node exchange3.company.com
```

Example (Restore method – pit)

Standalone environment

Suppose you want to restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`, using the backup data created in the session `2013/5/14-1`. Suppose you want the integration agent (`e2010_bar.exe`) to be started on the client `exchange1.company.com`. Execute the following command:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -pit -session 2013/5/14-1
```

NOTE:

The `-source` option is not specified, in which case Data Protector assumes that the database was backed up from the client specified with the `-barhost` option.

Example (Restore method – new)

DAG environment

Suppose you want to restore the backup of the database DB1 to a recovery database that should be created on the client `exchange2.company.com` and named `Recovery1`, with the files in the `C:\Recovery1Folder` directory. Suppose the database DB1 was backed up in the session `2013/5/14-1` from a DAG whose virtual system name was `dag0.company.com`. To also ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, execute the following command:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -source dag0.company.com -
new -session 2013/5/14-1 -client exchange2.company.com -location C:\Recovery1Folder
-name Recovery1 -recoverydb
```

Example (Restore method – temp)

Standalone environment

Suppose you want to restore the transaction logs of the database DB1, which resides on the client `exchange2.company.com`. The logs were backed up in the Incremental backup session `2013/5/14-1`. To restore the logs to the client `exchange2.company.com` to the directory `C:\DB1TransactionLogFolder` without performing database recovery, and to ensure that the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, execute the following command:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -source
exchange2.company.com -temp -session 2013/5/14-1 -client exchange2.company.com -
location C:\DB1TransactionLogFolder -no_chain -no_recover
```

Restoring using another device

You can restore using a device other than that used for a backup. For details, see the *HPE Data Protector Help* index: “restore, selecting devices for”.

Instant recovery

To be able to perform an instant recovery session, you need backup data that is stored on replica storage volumes. Such backup data is created in ZDB-to-disk and ZDB-to-disk+tape sessions.

You can restore multiple Microsoft Exchange Server databases in the same instant recovery session, specifying a different restore method for each database. For details, see [Restore methods, on page 252](#).

To start an instant recovery session, use the Data Protector GUI or CLI.

Performing instant recovery using the Data Protector GUI

To perform an instant recovery session:

1. In the Context List, click **Instant Recovery**.
2. Expand **Exchange 2010+ Server** and select which ZDB-to-disk or ZDB-to-disk+tape session to

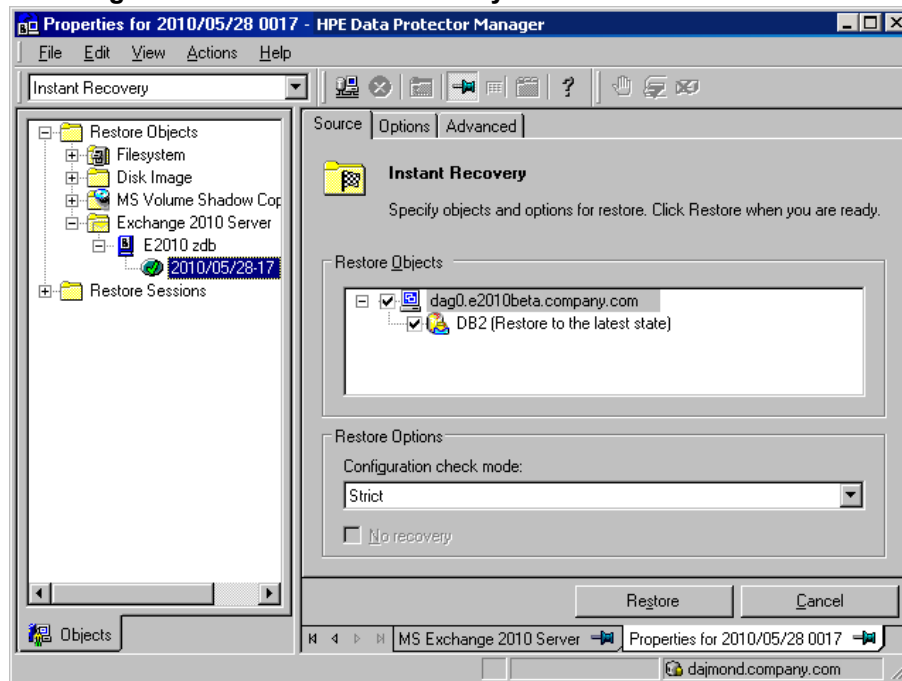
use for instant recovery. The sessions are sorted according to the backup specifications used.

3. In the **Source** page, select which Microsoft Exchange Server databases to restore.


When you select a database, the Properties for Database dialog box is displayed automatically. Specify a restore method and click **OK**. For databases that are part of a DAG, the default restore method is **Repair all passive copies with failed status**. For standalone databases, the default is **Restore to the latest state**. To change the restore method, right-click the database and click **Properties**.

For details on **Configuration check mode**, press **F1**.

Selecting databases for instant recovery

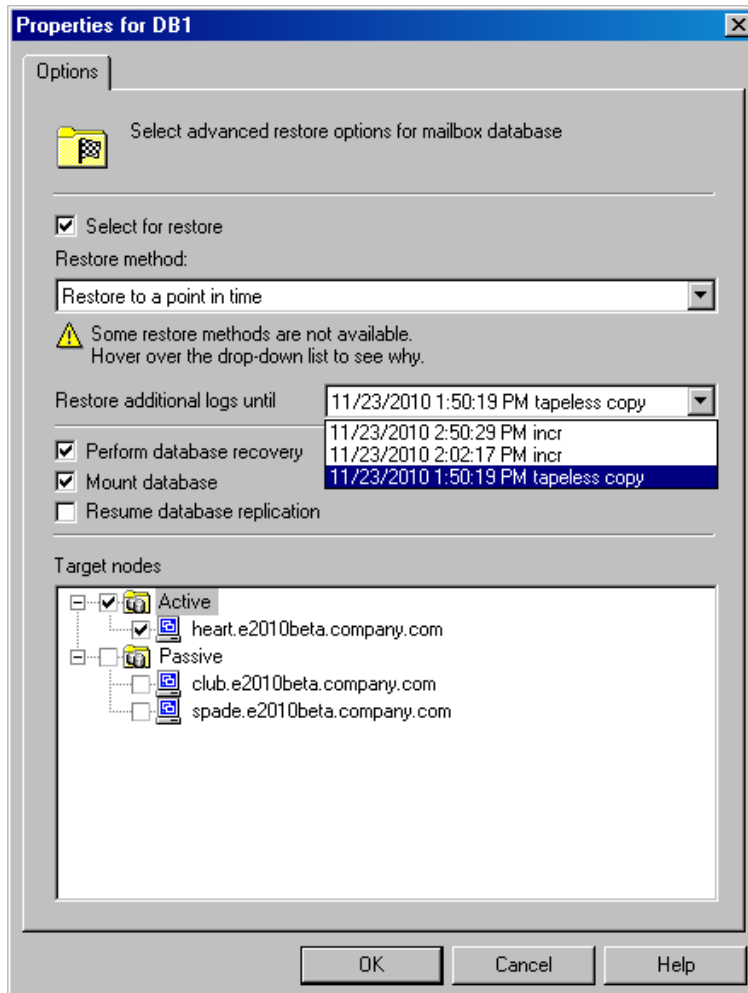


Restore to the latest state



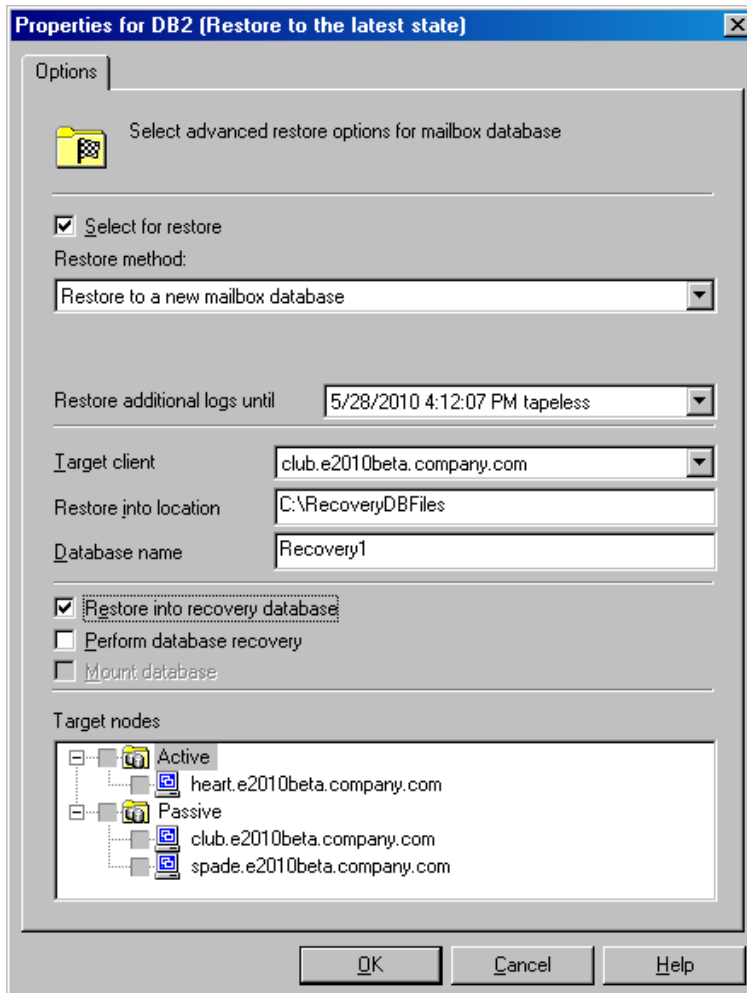
For details, see [Restore to a point in time, on page 277](#).

Restore to a point in time




For details, see [Restore to a point in time](#), on page 277.

Restore to a new mailbox database



For details, see [Restore to a new mailbox database, on page 254](#).

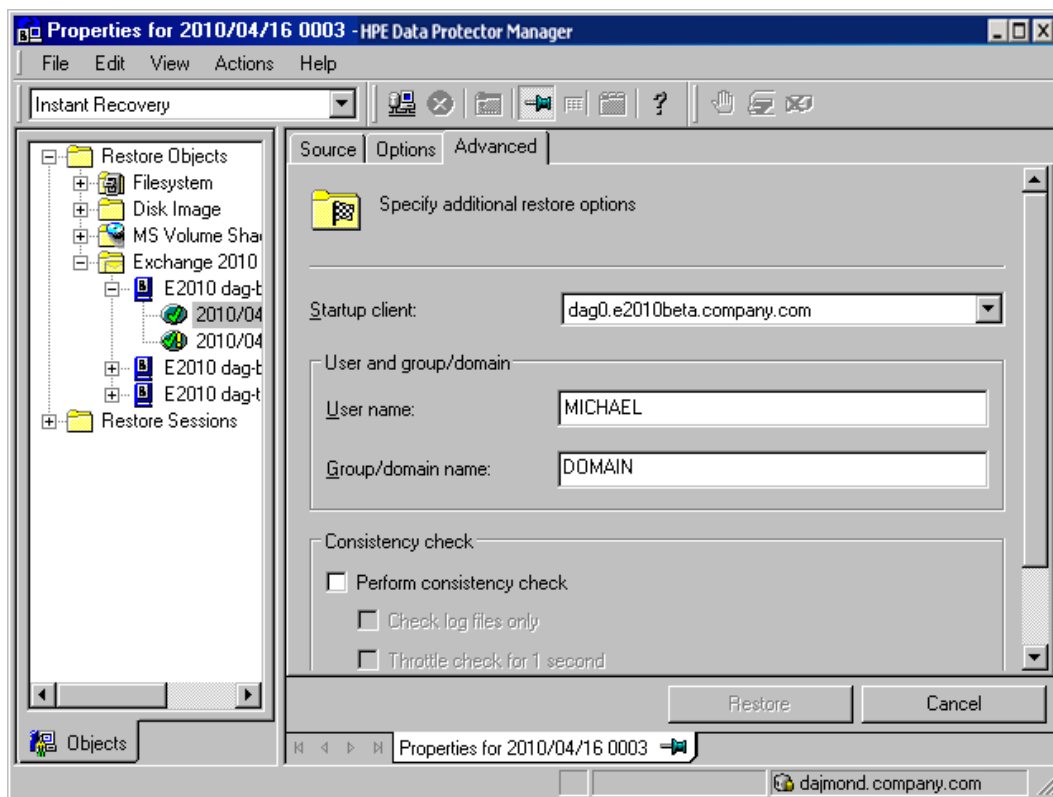
Restore files to a temporary location



For details, see [Restore files to a temporary location, on page 254](#).

4. In the **Options** page, specify ZDB-specific options. For details, press **F1**.
5. In the **Advanced** page, specify the Data Protector Microsoft Exchange Server 2010 integration instant recovery options. For details, see [General restore options, on page 280](#).

Instant recovery – advanced



6. Click **Restore**.

Performing instant recovery using the Data Protector CLI

1. Log in to the Data Protector Cell Manager or to any Microsoft Exchange Server client under a user account as described in [Configuring user accounts, on page 235](#).
2. Execute the following:

```
omnir -e2010
      -barhost ClientName
      -instant_restore
          [VSS_INSTANT_RECOVERY_OPTIONS]
      [VSS_EXCHANGE_SPECIFIC_OPTIONS]
          Database [Database ...]
      [-user User:Domain]
      [GENERAL_OPTIONS]
```

Database

```
{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID}
[-source SourceClientName]
{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS E2010_REPAIR_
METHOD_OPTIONS
[-no_resume_replication]
```

E2010_LATEST_METHOD_OPTIONS

```
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]
[E2010_IR_SPECIFIC_OPTIONS] E2010_PIT_METHOD_OPTIONS
-session SessionID
[-node TargetNode ... | -all]
[-no_resume_replication]
[-no_recover]
[-no_mount]
[E2010_IR_SPECIFIC_OPTIONS] E2010_NEW_METHOD_OPTIONS
-session SessionID
-client TargetClientName
-location TargetDatabasePath
-name TargetDatabaseName
[-recoverydb]
[-no_recover]
[-no_mount]
[E2010_IR_SPECIFIC_OPTIONS] E2010_TEMP_METHOD_OPTIONS
-session SessionID
-client TargetClientName
-location TargetDatabasePath
[-no_chain]
[-edb_only]
[-no_recover]
[E2010_IR_SPECIFIC_OPTIONS] E2010_IR_SPECIFIC_OPTIONS
[-from_session SessionID]
```

For brief description of the options, see the section [Restore options, on the next page](#).

For details on the options, see the *HPE Data Protector Command Line Interface Reference* or the `omnir` man page.

Example (Restore method – latest)

Standalone environment

Suppose you want to restore the corrupt standalone database DB1, which resides on the client `exchange1.company.com`. To ensure the integration agent (`e2010_bar.exe`) is started on the client `exchange1.company.com`, and that the database is restored to the latest state, execute:

```
omnir -e2010 -barhost exchange1.company.com -instant_restore -copy_back -db_name DB1 -latest
```

Example (Restore method – temp)

DAG environment

Suppose you want to restore the database DB1, which was part of a DAG whose DAG virtual system (host) name was `dag0.company.com`. The database was backed up in the session `2013/5/14-1`. To

restore the database to a temporary location on the client exchange1.company.com to the directory C:\BackupDatabase, and to ensure that the integration agent (e2010_bar.exe) is started on the client exchange1.company.com, execute:

```
omnir -e2010 -barhost exchange1.company.com -instant_restore -copy_back -db_name DB1 -source dag0.company.com -temp -session 2013/5/14-1 -client exchange1.company.com -location C:\BackupDatabase
```

Restore options

Repair all passive copies with failed status

Option in the GUI / CLI	Description
Resume database replication / -no_resume_replication	Available in DAG environments. Resumes the replication between the active and passive copies after the copies are restored. Note that the CLI option -no_resume_replication has the opposite meaning. If it is specified, the replication is not resumed.
Restore additional logs until -session	This is an instant recovery-specific option. Not available.
Target nodes	Not available. The clients (that is, copies) that have the status Failed or FailedAndSuspended are automatically selected.

Restore to the latest state

Option in the GUI / CLI	Description
Select for restore	Specifies whether the database should be restored.
Restore additional logs until -session	This is an instant recovery-specific option. Not available.
Perform database recovery / -no_recover	Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Applies the logs to the database file (.edb) after the restore completes. Note that the CLI option -no_recover has the opposite meaning. If it is specified, the database recovery is not performed.
Mount database /	Available when restoring a standalone database (standalone environment) or an active copy (DAG environment). Mounts the

Option in the GUI / CLI	Description
-no_mount	<p>database after the database recovery completes. This option is available only if Perform database recovery is selected.</p> <p>Note that the CLI option -no_mount has the opposite meaning. If it is specified, the database is not mounted.</p>
Resume database replication / -no_resume_replication	<p>Available when restoring passive copies (DAG environment). Resumes the replication between the active and passive copies after the copies are restored.</p> <p>Note that the CLI option -no_resume_replication has the opposite meaning. If it is specified, the replication is not resumed.</p>
Target nodes -node -all	<p>Available only in DAG environments. Specifies which clients (that is, database copies) to restore.</p>

Restore to a point in time

Option in the GUI / CLI	Description
Select for restore	<p>See the description in Restore to the latest state, on the previous page.</p>
Backup version / -session	<p>This is a standard restore-specific option.</p> <p>It specifies from which backup data to restore. Select a backup ID.</p> <p>If a Differential backup session is selected, the .log files backed up in the selected Differential backup session are restored.</p> <p>If an Incremental backup session is selected, the .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session, are restored.</p>
Last backup version	<p>This is a standard restore-specific option.</p> <p>It shows the session in which the database was last backed up.</p>
Restore additional logs until -session	<p>This is an instant recovery-specific option.</p> <p>If a Differential backup session is selected, the .log files backed up in the selected Differential backup session are restored.</p> <p>If an Incremental backup session is selected, the .log files backed up in all subsequent Incremental backup sessions, up to the selected Incremental backup session, are restored.</p>

Option in the GUI / CLI	Description
Perform database recovery / -no_recover	See the description in Restore to the latest state, on page 276 .
Mount database / -no_mount	
Resume database replication / -no_resume_replication	
Target nodes / -node -all	See the description in Restore to the latest state, on page 276 . The node (client) hosting the active copy is automatically selected for restore.

Restore to a new mailbox database

Option in the GUI / CLI	Description
Select for restore	See the description in Restore to the latest state, on page 276 .
Restore additional logs until -session	This is an instant recovery-specific option. See the description in Restore to a point in time, on the previous page .
Target client / -client	Specifies the client to restore to.
Restore into location / -location	Specifies the directory to restore to (standard restore) or the directory to mount the replica storage volumes to (instant recovery).
Database name / -name	Specifies the name to be used for the new database. If another database with the same name already exists, the restore fails.
Restore into Recovery database / -recoverydb	Restores the data to a Microsoft Exchange Server recovery database. Although multiple recovery databases can exist in parallel, only one recovery database can be mounted to the Microsoft Exchange Server at a time.
Backup version / -session	See the description in Restore to a point in time, on the previous page .

Option in the GUI / CLI	Description
Last backup version	
Perform database recovery / -no_recover	See the description in Restore to the latest state, on page 276 .
Mount database / -no_mount	
Target nodes	Not available.

Restore files to a temporary location

Option in the GUI / CLI	Description
Select for restore	See the description in Restore to the latest state, on page 276 .
Restore additional logs until -session	This is an instant recovery-specific option. See the description in Restore to a point in time, on page 277 .
Restore chain	If this option is set to Restore only this backup , only files backed up in the selected session are restored. If this option is set to Full restore (full, incr, diff backups) , the complete chain is restored.
Target client / -client	See the description in Restore to a new mailbox database, on the previous page .
Restore into location / -location	
Backup version / -session	See the description in Restore to a point in time, on page 277 .
Last backup version	
Restore database files only / -edb_only	Restores only the database files (.edb). The logs (.log) and checkpoint files (.chk) are not restored.
Perform database recovery / -no_recover	See the description in Restore to the latest state, on page 276 .
Target nodes	Not available.

General restore options

Option in the GUI / CLI	Description
<p>Startup client / -barhost</p>	<p>Specifies the client on which the integration agent (e2010_bar.exe) should be started. If the DAG virtual client (host) is selected, the integration agent is started on the currently active node. To find out which node is currently active, see Restore, on page 252.</p> <p>Default: The same client that was specified for the backup session. If the DAG virtual client was specified, this client is now selected. However, note that the integration agent may not be started on the same physical node as during the backup session; it depends which node is currently active.</p>
<p>Username Group/Domain name / -user</p>	<p>Specifies which Windows domain user account to use for the restore session. Ensure that the user is configured as described in Configuring user accounts, on page 235.</p> <p>If these options are not specified, the restore session is started under the user account under which the Data Protector Inet service is running.</p>
<p>Perform consistency check / [-exch_check [-exch_throttle <i>Value</i>] -exch_checklogs]</p>	<p>If this option is selected, Microsoft Exchange Server checks the consistency of a database's backup data. If this option is not selected, the session finishes earlier, but the backup data consistency is not guaranteed.</p> <p>The check is performed at the target location on the source storage volumes after the backup data is restored. You do not need to perform the consistency check if it was already performed at the time of backup.</p> <p>Default: not selected</p> <p>If the Check log files only option is selected, only the log file backup data is checked, which is enough for Microsoft Exchange Server to guarantee data consistency.</p> <p>Default: not selected</p> <p>By default, the consistency check is I/O intensive, which can negatively affect disk performance. The Throttle check for 1 second option throttles down the consistency check of the database file .edb to lessen impact on the disk performance. Specify after how many input/output operations the check should stop for one second.</p> <p>This option is not available if only the log files are checked.</p> <p>Default: not selected</p>

TIP:
To find out which Microsoft Exchange Server node is currently active, connect to one of the nodes and execute:

```
cluster group
```

Example

```
C:\Users\administrator.E2010BETA>cluster group  
Listing status for all available resource groups:
```

Group	Node	Status
Available Storage	spade	Offline
Cluster Group	club	Online

The currently active node has the status `Online`. In the example, this is `club`.

Monitoring sessions

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

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

To monitor a session, see the *HPE Data Protector Help* index: “viewing currently running sessions”.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft Exchange Server 2010 integration.

Because the Data Protector Microsoft Exchange Server 2010 integration is based on the Data Protector Microsoft Volume Shadow Copy Service integration, also see troubleshooting information in the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the *HPE Data Protector Help* index: “patches”.
- See the *HPE Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <https://softwaresupport.hpe.com/manuals> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your browsing, backup, or restore failed:

- Examine system errors reported in the `debug.log` file.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HPE Data Protector Help*.

Problems

Problem

It takes a long time to display Microsoft Exchange Server topology in the Data Protector GUI

When you open the Data Protector GUI and try to display the source page, either in the Backup or Restore context, you must wait a long time.

This may happen if there is an unresponsive system in the same domain (for example, a system that is shut down). The problem occurs even if the unresponsive system is not part of your backup environment. This is due to Microsoft Exchange Server problems with execution of Microsoft Exchange Server Shell commands.

Action

Remove the system from the domain or fix the problem.

Problem

A database backup cannot be performed

When you start a backup session for a database, the database is not backed up, appearing to be locked by other session, though there are no other backup sessions currently running. A message similar to the following is displayed:

```
[Minor] From: OB2BAR_E2010_BAR@exch03.e2010.company.com "MS Exchange Server" Time: 1/17/2013 3:07:13 PM  
[170:313] One or more copies of database DEMAR are already being backed up in a different session.
```

This may happen if the integration agent (`e2010_bar.exe`) was terminated by force while a previous backup session was in progress, either because the Microsoft Exchange Server system was restarted or for some other reason, so the lock remains.

Action

Execute the following command:

```
omnidbutil -free_cell_resources
```

NOTE:

This command line removes all existing locks, so ensure that none of the existing locks is still needed.

Problem

Restore fails

When you try to restore a database, the session fails.

This may happen if a database has been restored before (probably unsuccessfully), and during that previous restore session, the Microsoft Exchange Server created an `.env` file in the database directory. This file now prevents the database from being restored again.

Action

Delete the `.env` file and start a new restore session.

Problem

Restore from an object copy fails in a DAG environment

When restoring a database from a media set created in an object copy session, as the media set created in the original backup session no longer exists, the session fails with an error similar to the following:

```
[Critical] From: OB2BAR_E2010_BAR@computer1.company.com "MS Exchange 2010 Server"  
Time: 28/02/2013 16:08:12 No mailbox database copy can be selected for  
restore/instant recovery.
```

Action

1. Verify that the media set created in the object copy session still exists.
2. On all Microsoft Exchange Server system nodes, set the environment variable `OB2BARHOSTNAME` to the name of the DAG virtual system and restart the Data Protector Inet service.
3. Start a new restore session.

Problem

Restore to the latest state fails

When you try to restore a database all of whose log files were lost, using the restore method **Restore to the latest state** with the **Perform database recovery** option selected, the database recovery fails.

This may happen if a database is restored from a Full backup (that is, the restore chain consists of only the Full backup session). Since in the **Restore to the latest state** session, only the `.edb` file is restored from the Full backup (see [Restore chain, on page 256](#)), when the database recovery is started, there are no logs to be applied to the database file, and the database recovery fails.

Action

Restore the database using the restore method **Restore to a point in time**. For details, see [Restore to a point in time, on page 253](#).

Problem

After instant recovery to a point in time, passive copies remain in the Failed state

When you start an instant recovery session in a DAG environment to restore active and passive copies of the same database, using the restore method **Restore to a point in time**, the data is successfully

restored, but the synchronization between the active copy and passive copies fails, leaving the passive copies in the Failed state.

This problem occurs if, after the data is restored, the passive copy has extra log files, which are not present at the active copy side, and so synchronization cannot be established. This can happen if, during a Full backup session, multiple copies of a database are selected for backup. Data Protector first performs a Full backup of the passive copy that has the fewest logs applied to the database file, and then performs a Copy backup of all the remaining copies, with the active copy being backed up last. While the backup session is in progress, a new log may be created at the active copy side, so when the active copy is backed up, the newly created log is also backed up. If, further on in time, a failover occurs (one of the passive copies becomes the active copy) and you perform a **Restore to a point in time** instant recovery, each copy is restored from its own replica storage volumes. This results in the active copy (which was passive at the time of backup) having fewer logs than the passive copy (which was active at the time of backup). Consequently, synchronization cannot be established.

Action

Perform a full reseed for all Failed passive copies.

Problem

In a DAG, a copy-back instant recovery fails when restoring a non-original database copy

Suppose you backed up a database copy by creating a snapclone replica (HPE P6000 EVA Disk Array Family). Using the copy-back method, this replica can be used to restore the original database copy and/or the related database copies that reside on different Microsoft Exchange Server systems in the DAG. If the size of the source storage volumes on those Microsoft Exchange Server systems differs from the size of the source storage volumes that were backed up, the instant recovery session fails.

If the **Retain source for forensics** option is selected, a message similar to the following is displayed:

```
[Warning] From: SMISA@dizzy.e2008.company.com "SMISA" Time:
 1/17/2013 2:51:08 PM
[236:8001] This pair of source and target storage volumes are
not the same size.
Source storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\dizzy\dizzy-DB1-data\ACTIVE
Source size : 4 GB
Target storage volume : 50014380025B4860\\Virtual Disks\VSS\
FizzyDizzy\DAG\fizzy\hpVSS-LUN-06Jul10 02.23.27\ACTIVE
Target size : 3 GB
```

```
[Major] From: OB2BAR_VSSBAR@dizzy.e2008.company.com "MS Exchange
Server" Time: 1/17/2013 3:11:16 PM
The system failed to refresh symbolic links in kernel object namespace.
```

If the **Retain source for forensics** option is not selected, a message similar to the following is displayed:

```
[Major] From: SMISA@dizzy.e2008.company.com "SMISA" Time:
 1/17/2013 2:51:08 PM
[236:8001] This pair of source and target storage volumes are not
the same size.
Source storage volume : 50014380025B4860\\Virtual Disks\VSS\
```


FizzyDizzy\DAG\dizzy\dizzy-DB1-data\ACTIVE

Source size : 4 GB

Target storage volume : 50014380025B4860\\Virtual Disks\VSS\

FizzyDizzy\DAG\fizzy\hpVSS-LUN-06Jul10 02.23.27\ACTIVE

Target size : 3 GB

Action

Ensure that storage volumes on different Microsoft Exchange Server systems are the same size.

Chapter 5: Data Protector Microsoft SharePoint Server VSS based solution

Introduction

This chapter explains how to configure and use the Data Protector Microsoft SharePoint Server VSS based solution (**VSS based solution**). In reality, the solution is based on the Data Protector Microsoft Volume Shadow Copy Service integration (**VSS integration**). For details on the VSS integration, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

The chapter describes concepts and methods you need to understand to back up and restore Microsoft SharePoint Server 2010 and Microsoft SharePoint Server 2013 data that is stored in Microsoft SQL Server databases. For example:

- The configuration database (SharePoint_Config)
- Content databases (SharePoint_AdminContent_Label, WSS_Content_Label,...)
- SharePoint Service Applications databases (SSA_DB) (Microsoft SharePoint Server 2010/2013)

In addition, you can also back up and restore Microsoft SharePoint Server search index files.

From now on, both Microsoft SharePoint Server versions are called **Microsoft SharePoint Server**, unless the differences are pointed out.

Backup

Microsoft SharePoint Server data that is stored in Microsoft SQL Server databases is backed up using one of the following Microsoft SQL Server VSS writers:

- MSDE writer (for Microsoft SQL Server 2000 databases)
- SqlServerWriter (for Microsoft SQL Server 2008 databases)

Microsoft SharePoint Server 2010 search index files are backed up using the following VSS writers:

- OSearch14 VSS writer
- SPSearch4 VSS writer

Microsoft SharePoint Server 2013 search index files are backed up using following VSS writers:

- OSearch15 VSS writer

Microsoft FAST Search Server 2010 search index files are backed up:

- using the Data Protector Disk Agent (in case of the standard filesystem backup with VSS enabled)
- using the Data Protector VSS integration (in case of the ZDB filesystem backup)

You can create and run backup specifications using the Data Protector PowerShell command which is described in [Backup, on page 290](#).

Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- **Microsoft SharePoint Server 2010:** With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only Full backup type is supported.

Restore

Restore can be started using the Data Protector GUI or CLI as described in [Restore, on page 301](#).

Installation and configuration

ZDB prerequisites

If you plan to run ZDB and instant recovery (IR) sessions, ensure that the SPSearch and OSearch index files of each SSP or SSA, and the FAST Search index files reside on a disk array.

Microsoft Office SharePoint Server 2007

The default location for the SPSearch index files is:

```
C:\Program Files\Microsoft Office Servers\12.0\Data\Applications
```

The default location for the OSearch index files is:

```
C:\Program Files\Microsoft Office Servers\12.0\Data\Office Server\Applications
```

To move the index files to the disk array:

1. Open the Command Prompt and change the directory to:

```
C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\12\BIN>
```
2. To move the SPSearch index files, execute:

```
stsadm -o spsearch -indexlocation PathToNewLocation
```
3. To move the OSearch index files, execute:

```
stsadm -o editssp -title SSPname -indexlocation PathToNewLocation
```

Microsoft SharePoint Server 2010

The default location for the SPSearch index files is:

```
C:\Program Files\Microsoft Office Servers\14.0\Data\Applications
```

The default location for the OSearch index files is:

```
C:\Program Files\Microsoft Office Servers\14.0\Data\Office Server\Applications
```

To move the index files to the disk array:

1. Open the Command Prompt and change the directory to:
`C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\14\BIN>`
2. To move the SPSearch index files, execute:
`stsadm -o spsearch -indexlocation PathToNewLocation`
3. To move the OSearch index files use the Central Administration (modify farm topology).

The FASTSearch home folder must be installed to the disk array during the FAST Search Server 2010 system installation. In case of a multiple FAST Search Server system farm, ensure that the FASTSearch home folders on all the systems have the same path (drive and path name).

Licensing

The Data Protector VSS based solution requires one online-extension license per each Microsoft SharePoint Server client participating in the backup and restore process. This means one online-extension license for each system on which the Data Protector MS Volume Shadow Copy Integration component is installed.

Installing the integration

For details on how to install a Data Protector cell, see the *HPE Data Protector Installation Guide*.

To be able to back up Microsoft SharePoint Server objects, install the following installation packages and Data Protector components:

- Windows PowerShell 2.0 or later and the Data Protector User Interface component on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands and on which you install the Data Protector MS Volume Shadow Copy Integration component. See the next bullet.

If not already available on your Windows system, you can download Windows PowerShell from <http://www.microsoft.com/windowsserver2003/technologies/management/powershell/default.aspx>.

- The Data Protector MS Volume Shadow Copy Integration component on the Microsoft SQL Server system and the Microsoft SharePoint Server systems that have at least one of the following services enabled:

Microsoft Office SharePoint Server 2007:

- Windows SharePoint Services Database
- Windows SharePoint Services Help Search
- Office SharePoint Server Search

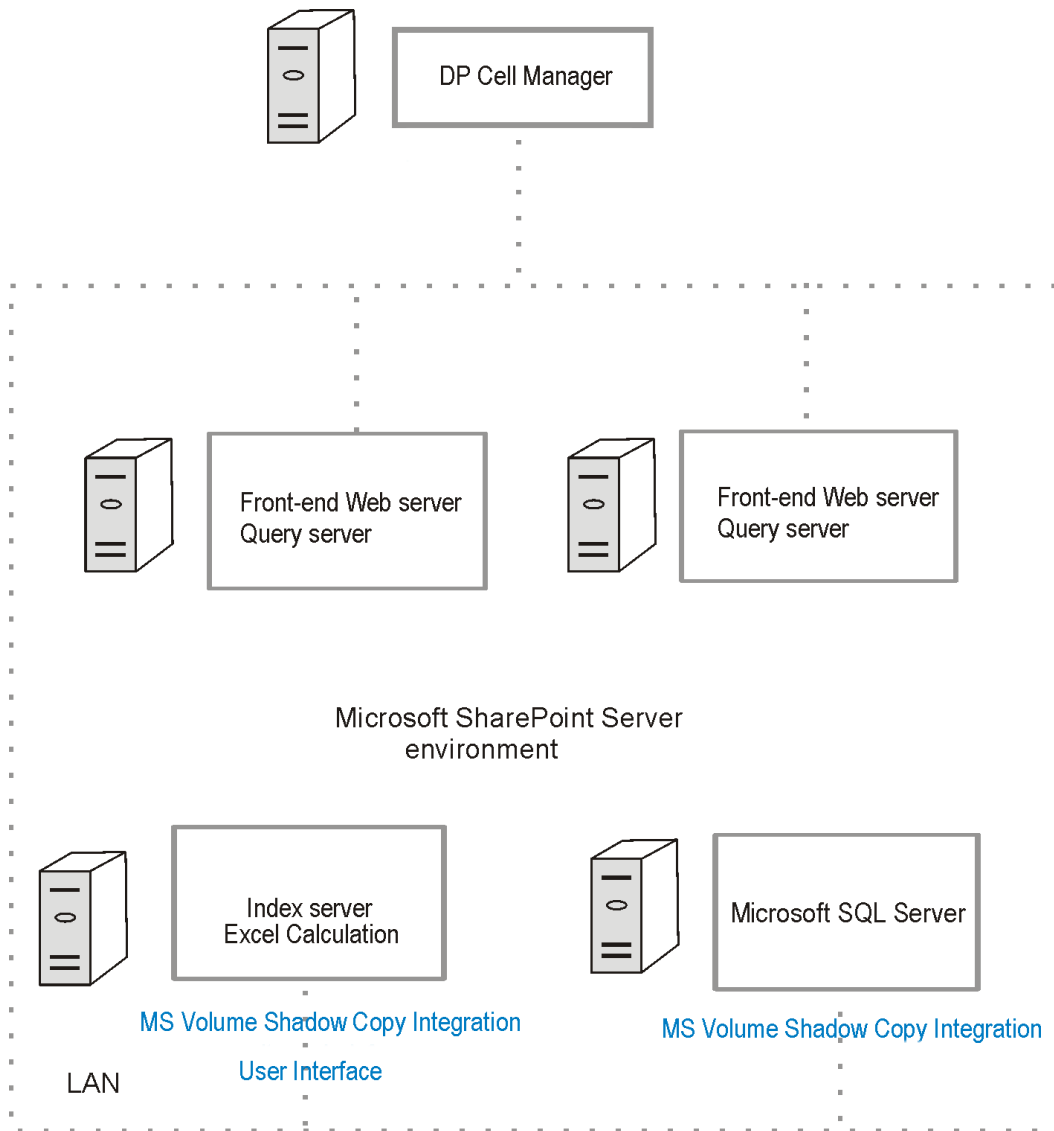
Microsoft SharePoint Server 2010/2013:

- SharePoint Foundation Database
- SharePoint Foundation Help Search (Microsoft SharePoint Server 2010 only)

- SharePoint Server Search
- FAST Search Server 2010 for SharePoint
- The Data Protector Disk Agent component on each Microsoft FAST Search Server 2010 system for SharePoint (Microsoft SharePoint Server 2010, in case of the mixed filesystem + ZDB environment)

Ensure that the Volume Shadow Copy service is started on all these clients.

Installing a medium farm (example)



In **Installing a medium farm (example)**, above, the Data Protector components that you need to install are colored blue.

Configuring the integration

For details on how to configure the Data Protector VSS integration, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Configuring user accounts

Create or identify a Windows domain user account that has Windows administrative rights on the Microsoft SharePoint Server system on which you plan to execute the Data Protector commands. This user must also be granted Microsoft SharePoint Server administrative rights and must be added to the Data Protector `admin` user group.

Backup

To back up Microsoft SharePoint Server data, create backup specifications and start backup sessions using the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`.

Prerequisites

- The Windows Remote Management service (which is used for starting and stopping Windows services remotely, and suspending and resuming FAST for Microsoft SharePoint Server 2010) must be configured on all systems.

To configure and analyze the WinRM service, execute the `winrm quickconfig` command.

For more information, see the Windows Remote Management service documentation.

- In the case of Microsoft SharePoint Server 2010/2013 which uses Microsoft SQL Server 2008/2012 for storing data, and Remote BLOB Storage (RBS) is used with the FILESTREAM provider, ensure that FILESTREAM access level is set to `Full access enabled` or `Transact-SQL access enabled`.

For details of how to configure RBS and FILESTREAM, see the Microsoft SQL Server 2008 documentation.

Limitations

- The only supported way to start backup sessions is using the Data Protector PowerShell command. Starting the backup sessions using the Data Protector GUI or CLI is not supported.
- **Microsoft SharePoint Server 2010:** With VSS based solution, the FAST Search index files can also be backed up incrementally when using the Data Protector Disk Agent. For all other Microsoft SharePoint Server data only `Full` backup type is supported.

Recommendations

- Use the Data Protector PowerShell command to create backup specifications and not the Data Protector GUI.

- Use the Data Protector GUI to modify backup specifications (for example, to add backup devices).
- Use the simple mode for the SQL Server databases. In case you want to use the full mode anyway, ensure that you truncate the transaction logs. Otherwise, you may run out of disk space.
- Whenever you change the farm configuration, perform a new backup.
- In case you want to back up the Single Sign-On database, do not forget to back up the encryption key as described in:
<http://technet.microsoft.com/en-us/library/cc262932.aspx#Section32>.
Otherwise, you will not be able to restore the database.

How the command works

When you execute the Data Protector PowerShell command `SharePoint_VSS_backup.ps1`, Data Protector first queries for information about the Microsoft SharePoint Server environment. Then it creates backup specifications.

The newly created backup specifications are named `SharePoint_VSS_backup_ClientName` and have the same backup device specified for use (the one that you specified at command runtime).

Once the backup specifications are created, the command starts backup sessions (one session for each backup specification).

Microsoft SharePoint Server 2010

In a Microsoft SharePoint Server 2010 environment, the command creates a separate backup specification for each Microsoft SharePoint Server system that has at least one of the following services enabled:

- SharePoint Foundation Database
- SharePoint Foundation Help Search
- SharePoint Server Search 14
- FAST Search Server 2010 for SharePoint (FAST Search)

For a system with the SharePoint Foundation Database service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server 2008) object selected.

For a system with the SharePoint Foundation Help Search and SharePoint Server Search services enabled, the command creates a backup specification that has the `SPSearch4 VSS Writer` and `OSearch14 VSS Writer` objects selected.

For a system with the FAST Search Server 2010 service enabled, the command with the `-hardware` option specified creates a VSS backup specification that has the complete `FASTSearch` home folder (including `bin` and `lib`) selected.

Microsoft SharePoint Server 2013

In a Microsoft SharePoint Server 2013 environment, the command creates a separate backup specification for each Microsoft SharePoint Server system that has at least one of the following services enabled:

- SharePoint Foundation Database
- SharePoint Server Search 15

For a system with the SharePoint Foundation Database service enabled, the command creates a backup specification that has the `SqlServerWriter` (Microsoft SQL Server 2008/2012) object selected.

For a system with the SharePoint Server Search services enabled, the command creates a backup specification that has the `OSearch15 VSS Writer` objects selected.

Considerations

- The command options enable you to split the process into two parts: first you create the backup specifications and then you start backup sessions. In this way, you can manually modify the newly-created backup specifications in the Data Protector GUI before the backup is actually started.
- If Microsoft SQL Server instances are used not only by Microsoft SharePoint Server but also by other database applications, modify the backup specifications so that only the databases that belong to Microsoft SharePoint Server are selected for backup. See the section [Modifying backup specifications, on page 299](#).
- If you have Microsoft SQL Server database mirroring enabled, a failover can occur and so a different Microsoft SQL Server system becomes active. Since the command creates backup specifications only for the currently active Microsoft SQL Server systems, it is advisable to update (recreate) the backup specifications before the backup is started.

The command syntax

```
SharePoint_VSS_backup.ps1 -help | -version  
SharePoint_VSS_backup.ps1 -createonly CreateOptions  
SharePoint_VSS_backup.ps1 -backuponly BackupOptions  
SharePoint_VSS_backup.ps1 -resume farm [-preview] | -resume cert
```

CreateOptions

```
{-device DevName | -hardware {no_keep|keep|ir} [-device DevName]}  
[-overwrite]  
[-prefix PrefixName]  
[-excludeindex]
```

BackupOptions

```
[-outfile PathToFile]  
[-prefix PrefixName]  
[-preview]  
[-snapshot {diskonly | disktape | tapeonly}]  
[-reduce]  
[-mode {full | incremental | incremental1 ... | incremental9}]  
[-timeout Timeout]
```

IMPORTANT:

- The command must be executed from the `Data_Protector_home\bin` directory on the front-

end Web Server system. Ensure that you are logged in under a user account that is configured as described in [Configuring user accounts, on page 290](#) and that you open the command prompt with administrative rights.

- Do not close the PowerShell console while the backup session is in progress. If you close the console during the backup, some actions are not performed: the backup sessions started do finish, but the farm does not resume the original state. To resume the farm, first execute the command with the `-resumeFarm` option and then unquiesce the farm manually using the Microsoft SharePoint Server Central Administration or `stsadm`.

Option description

<code>-help</code>	Displays the <code>SharePoint_VSS_backup.ps1</code> command usage.
<code>-version</code>	Displays the <code>SharePoint_VSS_backup.ps1</code> version.
<code>-createonly</code>	If this option is specified, Data Protector only creates backup specifications. Backup is not started.
<code>-backuponly</code>	If this option is specified, Data Protector only starts backup sessions using the existing backup specifications. The <code>-device</code> option is not required.
<code>-deviceDevName</code>	<p>Specifies which Data Protector device to use for backup. You can specify only one device.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>IMPORTANT: If only one device is used to back up a multi-system farm, the corresponding backup sessions cannot run in parallel. This prolongs the time during which the farm is in read-only mode. Specifically, the farm is in read-only mode from the moment when the backup sessions are started up until all VSS snapshots are created.</p> <p>To enable backup sessions to run in parallel, select different or additional devices in each backup specification before the backup is started. See the section Modifying</p> </div>

	<p>backup specifications, on page 299.</p>
<p>-hardware {no_ keep keep ir}</p>	<p>Specifies that the hardware provider should be used (instead of the software provider with -device option specified) and, consequently, ZDB options set. The default values for ZDB options are as follows:</p> <ul style="list-style-type: none"> • Keep the replica for instant recovery: selected if <i>ir</i> is specified. • Keep the replica after the backup: selected if <i>ir</i> or <i>keep</i> is specified. • Configuration check mode: Strict • Replica type: Mirror/Clone (Plex) • Numbers of replica rotated: 3 <p>The default ZDB backup types are as follows (provided a device is also specified):</p> <ul style="list-style-type: none"> • <i>no_keep</i>: ZDB-to-tape • <i>keep</i>: ZDB-to-disk+tape • <i>ir</i>: ZDB-to-disk+tape
<p>-overwrite</p>	<p>By default, Data Protector does not create backup specifications if they already exist. If this option is specified, Data Protector overwrites the existing backup specifications with the newly-created ones. Not applicable if -<i>backuponly</i> is specified.</p>
<p>-prefix<i>PrefixName</i></p>	<p>With this option specified, the backup specifications are created under a different name: <i>SharePoint_VSS_backup_PrefixName_ClientName</i>.</p> <p>In case of backup, this option specifies which backup specifications to use: those which name contains <i>PrefixName</i>.</p> <p>Non-ASCII characters in <i>PrefixName</i> are not supported.</p>
<p>-outfile<i>PathToFile</i></p>	<p>If this option is specified, backup specification names, errors, sessions outputs, and <i>omnir</i> restore commands are written to the specified file.</p>

<p>-preview</p>	<p>If this option is specified, Data Protector displays information about the Microsoft SharePoint Server environment and describes the related actions without actually performing them.</p>
<p>-snapshot { diskonly disktape tapeonly}</p>	<p>Applicable when starting ZDB backup sessions (that is, sessions that use backup specifications in which a hardware provider is specified for use). Performs a ZDB-to-disk (diskonly), ZDB-to-tape (tapeonly) or ZDB-to-disk+tape (disktape) session.</p>
<p>-reduce</p>	<p>Microsoft SharePoint Server 2010: If this option is specified, the command excludes mirrored query components from backup to reduce the backup size.</p> <p>Microsoft SharePoint Server 2013: If this option is selected, the command selects primary index replicas of each index partition to reduce the backup size.</p>
<p>-excludeindex</p>	<p>Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010/2013). If this option is specified, Data Protector excludes data_index folder contained in the FASTSearch home folder from backup specification. This way, the backup is faster, but the restore is more time consuming. The option enables balancing between a backup size and a time to recovery.</p>
<p>-mode { full incremental incremental1... incremental9}</p>	<p>Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010/2013). With this option specified, either a Full or Incremental or leveled incremental backup can be started. By default, the Full backup is performed.</p> <p>When the incremental option is specified and the Full backup does not exist, the option is ignored and the Full</p>

	filesystem backup of the FAST Search index files is started.
-resumecert	<p>Applicable only to Microsoft FAST Search Server 2010/2013. If this option is specified, the FAST Search certificates for the content and the query connectors are reinstalled.</p> <p>IMPORTANT: The SharePoint_VSS_backup.ps1 -resumecert command must be started on the Microsoft SharePoint Server system where the SharePoint Server Search 14 service is enabled.</p>
-resumefarm	<p>To be used after restore. This option returns the farm to a working state by resuming all background activities and crawling, unlocking sites, and starting Microsoft SharePoint Server services.</p> <p>IMPORTANT: The command with the -resumefarm option specified uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. To ensure its proper operation, an exception must be added to the Windows Default Firewall for Remote administration, which adds the WMI ports, or for the WMI directly. For details, see: http://support.microsoft.com/kb/154596.</p>
-timeout <i>Timeout</i>	This option sets the timeout in minutes after which the crawl of the FAST Search index files is aborted and the farm is resumed. If not specified, the default timeout is 15 minutes.

Starting Windows PowerShell

1. Log in to the Microsoft SharePoint Server system where Windows PowerShell and User Interface component are installed, under a user account that is configured as described in

[Configuring user accounts, on page 290.](#)

2. Open the Windows PowerShell CLI. For example:
Start > Programs > Accessories > Windows PowerShell > Windows PowerShell
3. In case you have Windows User Account Control (UAC) enabled, ensure that you open the CLI with administrative rights. Otherwise, you will not be able to run the Data Protector PowerShell command.
4. Ensure that the Windows PowerShell execution policy is set to RemoteSigned or Unrestricted.

[Displaying the Data Protector PowerShell command syntax, below](#) shows how the Windows PowerShell execution policy is set to Unrestricted and how the Data Protector PowerShell command syntax is displayed.

Displaying the Data Protector PowerShell command syntax

```
Administrator:SharePoint 2010 Management Shell
PS C:\Program Files\OmniBack\bin> .\SharePoint_USS_backup.ps1 -help

Usage synopsis:

SharePoint_USS_backup.ps1 -version | -help
SharePoint_USS_backup.ps1 -createonly CreateOptions
SharePoint_USS_backup.ps1 -backuponly [BackupOptions]
SharePoint_USS_backup.ps1 -resumefarm [-preview] | -resumecert

CreateOptions
<-device DeviceName | -hardware {no_keep | keep | ir} [-device DeviceName]>
[-overwrite]
[-prefix PrefixName]
[-excludeindex]

BackupOptions
[-outfile PathToFile]
[-prefix PrefixName]
[-preview]
[-snapshot {diskonly | disktape | tapeonly}]
[-reduce]
[-mode {full | incremental | incremental1 .. | incremental9}]

-version
Shows the version of script.
-help
Displays this help information.
-preview
Shows all the farm information and actions to be taken. Does not actually
perform any action and does not start the backup(s).
-createonly
Only creates backup specifications.
-overwrite
Overwrite the backup specifications during their creation. Not applicable
for -backuponly.
-backuponly
Performs backup only. Backup specification are not created, -device option
not required with -backuponly
-device <DP device name>
Device name to be used in created backup specifications.
For backup specification creation either '-device' or '-hardware' option
has to be present.
If more than one host is backed up, the backups will not run in parallel
with one device. In the destination page of the backup specification, you
can select different or additional devices. For more details about
modifying backup specification, see documentation.
-hardware {no_keep | keep | ir}
With this option created backup specification uses USS hardware providers.
Specify no_keep, keep or ir to specify whether to keep created disk copy and
tracks it for instant recovery.
For backup specification creation either '-device' or '-hardware' option has
to be present.
-prefix <prefix>
Additional prefix for backup specifications names.
-reduce
Script will exclude mirrored query components from backup to reduce the size
of backup. Applicable only for SharePoint 2010.
-excludeindex
Exclude FASTSearch index data from datalist. Applicable only for FASTSearch DA datalis
-mode {full | incremental | incremental1 .. | incremental9}
This option is used for starting full or incremental or leveled incremental backup.
If you don't use this option or if you use this option on wrong way by default
backup mode will be full. Applicable only for FASTSearch DA datalist.
-resumecert
Reinstall FASTSearch certificates for content and query connectors.
-snapshot {diskonly | disktape | tapeonly}
This option is used for starting backup session to disk or to tape or disk+tape .
Must be in use for backup specification that use hardware provider.
-outfile <filename>
Writes backup specifications names/restore and/or recovery commands/session
output to file specified.
-resumefarm
Resumes all farm(s) activities.

PS C:\Program Files\OmniBack\bin> _
```

Creating backup specifications (examples)

1. To create backup specifications in which the backup device filelib_writer1 is specified for use, execute:
SharePoint_VSS_backup.ps1 -createonly -device filelib_writer1
2. To create backup specifications with the label weekly in their names and in which the backup

device dev1 is specified for use, execute:

```
SharePoint_VSS_backup.ps1 -createonly -device dev1 -prefix weekly
```

3. To create ZDB backup specifications in which the backup device dev1 and the hardware provider (ZDB disk array) are specified for use, and in which the ZDB option Keep the replica for instant recovery is enabled, execute:

```
SharePoint_VSS_backup.ps1 -createonly -hardware ir -device dev1
```

4. Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010).

To create filesystem backup specifications in which the backup device dev1 is specified for use and with the data_index folder, contained in the FASTSearch home folder, excluded from the backup of the FAST Search index files, execute:

```
SharePoint_VSS_backup.ps1 -createonly -device dev1 -excludeindex
```

Modifying backup specifications

To modify a backup specification, open the Data Protector GUI. In the Context list, select **Backup** and, under **MS Volume Shadow Copy Writers** or under **Filesystem** (if performing a standard filesystem backup of the FAST Search index files), click the name of the backup specification that you want to modify.

Source page

To modify the Source page of the backup specification (for example, you want to back up individual Microsoft SharePoint Server databases), consider the following:

- The configuration database and the Central Administration content database must both be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure data consistency.
- **Microsoft SharePoint Server 2010:** The Help Search database and the associated index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure data consistency.
- **Microsoft SharePoint Server 2010:** The FAST search index files and the FAST Content SSA crawl components must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure data consistency.
- **Microsoft SharePoint Server 2010/2013:** The SharePoint Service Applications, Search database (SSA_Search_DB), and the associated search index files must all be backed up during the same time period which starts with the suspend of the Microsoft SharePoint Server (SharePoint farm) and ends with the resume of the SharePoint farm to ensure data consistency.

Otherwise, after restore, the Microsoft SharePoint Server data may not be consistent.

Destination page

In the Destination page of the backup specification, you can select different or additional devices and set the device and media options.

Options page

In the Options page of the backup specification, you can modify backup options. For the standard filesystem backup of the FAST search index files leave the **Use Shadow Copy** option specified to enable the use of the VSS. To modify ZDB options, click **Advanced** in the Backup Specifications field and then, in the Backup Options dialog box, click the **Advanced backup options** tab.

Starting backup sessions (examples)


1. To preview the actions that are performed when a backup session is started, execute:
`SharePoint_VSS_backup.ps1 -backuponly -prefix dev -preview`
2. To start backup sessions using the existing backup specifications that have no prefix in their names, execute:
`SharePoint_VSS_backup.ps1 -backuponly`
3. To start backup sessions using the existing backup specifications that have the prefix weekly in their names, execute:
`SharePoint_VSS_backup.ps1 -backuponly -prefix weekly`
4. To start backup sessions using the existing backup specifications that have no prefix in their names and to save the output of the sessions and the associated restore commands to the file `c:\logs\shp.log`, execute:
`SharePoint_VSS_backup.ps1 -backuponly -outfile C:\logs\shp.log`
5. To start ZDB-to-disk backup sessions using the existing ZDB backup specifications that have no prefix in their names, execute:
`SharePoint_VSS_backup.ps1 -backuponly -snapshot diskonly`
6. To start incremental filesystem backup sessions of the FAST Search index files (Microsoft SharePoint Server 2010), execute:
`SharePoint_VSS_backup.ps1 -backuponly -mode incremental`

Scheduling backup sessions

You can schedule backup sessions using the Windows system scheduler.

1. On the front-end Web server system, create a Windows PowerShell scheduled task. Go to:
Start > Settings > Control Panel > Scheduled Tasks > Add Scheduled Task
2. Open advanced properties for the task.

Scheduling a backup session using the Windows scheduler



In the **Run** text box, enter:

```
Windows_PowerShell_home \powershell.exe SharePoint_VSS_backup.ps1[Options]
```

For details on Options, see [The command syntax, on page 292](#).

In the **Start in** text box, enter:

```
Data_Protector_home \bin
```

In the **Run as** text box, enter a Windows domain user account `DOMAIN\UserName` that is configured as described in [Configuring user accounts, on page 290](#).

Restore

To restore Microsoft SharePoint Server data:

- Stop Microsoft SharePoint Server services
- Restore the data.
- Return the farm to a working state.

For details, see the following sections.

Before you begin

- Stop and disable the following services:
 - SharePoint Server Search 14 (Microsoft SharePoint Server 2010)
 - SharePoint Server Search 15 (Microsoft SharePoint Server 2013)

In addition, stop the following services:

Microsoft Office SharePoint Server 2007:

- Windows SharePoint Services Administration
- Windows SharePoint Services Search
- Windows SharePoint Services Timer

Microsoft SharePoint Server 2010:

- SharePoint 2010 Administration
- SharePoint Foundation Search V4
- SharePoint 2010 Timer
- SharePoint 2010 Tracing
- FAST Search for SharePoint
- FAST Search for SharePoint Monitoring

Microsoft SharePoint Server 2013:

- SharePoint Administration
- SharePoint Search Host Controller
- SharePoint Timer Service
- SharePoint Tracing Service

- Put the Microsoft SQL Server instance offline if you plan to restore one of the following Microsoft SQL Server databases:
 - master
 - model
 - msdb
 - a database for which Microsoft SQL Server mirroring is enabled

NOTE:

- If you use `SqlServerWriter`, you can restore the `model` and `msdb` databases also when the Microsoft SQL Server instance is online. This is one advantage over `MSDE writer`.
- *Microsoft SQL Server mirroring*: If the original and mirror database reside in separate Microsoft SQL Server instances, put offline both Microsoft SQL Server instances.

Restoring data

You can restore Microsoft SharePoint Server data using the Data Protector GUI or CLI.

Considerations

- The configuration database and the Central Administration content database must both be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure the data consistency. Since the configuration database and the Central Administration content database contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.
- **Microsoft SharePoint Server 2007/2010:** The Help Search database and the associated index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure data consistency.
- **Microsoft SharePoint Server 2010:**
 - Since the FAST configuration database and the FAST Search home folder contain system-specific information, you can restore them only to the original environment or to an environment that has precisely the same configuration, software updates, server names, and number of servers.
 - The FAST Search index files and the FAST Content SSA crawl components must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure data consistency.
- **Microsoft SharePoint Server 2010/2013:** The SharePoint Service Applications, Search database (`SSA_Search_DB`) and the associated index files must all be restored using backups from the same point in time (the backups performed in the same period in which the Microsoft SharePoint Server (SharePoint farm) was in suspended mode) to ensure data consistency.
- The following table shows which VSS restore modes are supported for which writers:

VSS supported restore modes and writers

Writers	VSS restore modes	
	Restore to another client	Restore files to temporary location

MSDE writer SqlServerWriter	No	Yes (manual attach needed)
OSearch VSS writer OSearch14 VSS writer/OSearch15 VSS writer	Yes	No
SPSearch VSS writer SPSearch4 VSS writer	Yes	No

Prerequisites

- Applicable only to a Data Protector filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010). Before restoring the FAST Search index files, the **Overwrite** option must remain selected to ensure the data consistency. It is selected by default.

Restoring using the Data Protector GUI

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **MS Volume Shadow Copy Writers**, expand the client which data you want to restore, and then click **MS Volume Shadow Copy Writers**.
 If performing a filesystem restore of the FAST Search index files (Microsoft SharePoint Server 2010), expand **Filesystem**, expand the client which data you want to restore, and then click the filesystem object.
3. In the Source page, select the data that you want to restore.
4. In the Options page, specify the restore options.
5. In the Devices page, select devices to use for restore.
6. Click **Restore**, review your selection, and click **Finish**.

Restoring using the Data Protector CLI

You can restore Microsoft SharePoint Server data using the Data Protector `omnir` command. For details, see the `omnir` man page or the *HPE Data Protector Command Line Interface Reference*.

If you specified the `-outfile` option when you ran backup sessions, you can find the necessary `omnir` commands in the specified file. The following is an example of the `omnir` command from such a file.

```
omnir -vss -barhost SHP-APP
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/master"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server2005:SQLWriter)/SHP-APP/model"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/msdb"
```

```
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/
WSS_Content_SSPAdminAccounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/SHP-APP/SSP_Accounting"
-session 2011/09/25-13
-tree "/SqlServerWriter(SQL Server 2005:SQLWriter)/
SHP-APP/SSP_Accounting_Search"
```

Limitations

The `omnir` command syntax should not contain more than 8191 characters. If you have so many `-tree` objects that the syntax exceeds 8191 characters, split the objects and run two separate sessions.

After the restore

After the restore:

1. Enable the service Office SharePoint Server Search, SharePoint Server Search 14, or SharePoint Server Search 15.
2. Bring the Microsoft SQL Server instances online (if offline).
3. Return the farm to a working state (that is, resume background activities and crawling, unlock sites, and start the Microsoft SharePoint Server services) by executing:

```
SharePoint_VSS_backup.ps1 -resumefarm
```

NOTE:

- The command uses the WMI (Windows Management Instrumentation) to remotely start any stopped SharePoint services. Ensure its proper operation by adding an exception to the Windows Default Firewall for Remote administration, which adds the WMI ports, or for the WMI directly. For details, see: <http://support.microsoft.com/kb/154596>.
- If the FAST Search certificates for the content and query connectors are out of sync, you can reinstall them by executing:

```
SharePoint_VSS_backup.ps1 -resumecert
```

Start the command on the Microsoft SharePoint Server system where the SharePoint Server Search 14 service is enabled.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Microsoft SharePoint Server VSS based solution.

For Microsoft Volume Shadow Copy troubleshooting information, see the troubleshooting chapter in the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. On how to verify this, see the *HPE Data Protector Help* index: “patches”.
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HPE Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <https://softwaresupport.hpe.com/manuals>.

Checks and verifications

If your browsing, backup, or restore failed:

- Examine system errors reported in the `debug.log` file.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HPE Data Protector Help*.

After restore, you cannot connect to the Central Administration webpage

Problem

After restore, when you try to connect to the Microsoft SharePoint Central Administration webpage, an error similar to the following is displayed in your web browser:

Windows Internet Explorer:

Retrieving the COM class factory for component with CLSID (BDEADEE2-C265-11D0-BCED-00A0C90AB50F) failed due to the following error 800703fa.

Mozilla Firefox:

An unexpected error has occurred.

Action

1. Restart Microsoft SharePoint Server services on all clients in the farm.
2. Open the Internet Information Services (IIS) Manager and restart all application pools.
3. In case an application pool fails to be restarted with the following error:
Cannot Restore Application Pool. There was an error while performing this operation.
wait for a few seconds and then restart the operation.
4. Delete browsing history in your web browser.
5. Log in to the Central Administration webpage.

Backup fails with the error Failed to resume Service Windows SharePoint Services Help Search

Problem

When you start backup sessions, an error similar to the following is displayed:

```
Service Windows SharePoint Services Help Search on host
MOSS07-INDEX
-> Resuming background activity ...
ERROR: Failed to resume Service Windows SharePoint Services Help Search
on host MOSS07-INDEX
Web site URL: http://moss07-web:2001
Root title: as
-> Resuming background activity
```

Action

Execute:

```
SharePoint_VSS_backup.ps1-resumefarm
```

After restore, a quiesce operation fails

Problem

After you have restored the configuration database and executed `SharePoint_VSS_backup.ps1-resumefarm`, the data in the Microsoft SharePoint Server file system caches on front-end Web Server systems is not consistent with the data in the newly-restored configuration database. When you try to quiesce the farm, the operation fails with the following error:

```
An unhandled exception occurred in the user interface. Exception
Information: An update conflict has occurred, and you must re-try
this action. The object SessionStateService Parent=SPFarm Name=<
farm_config_database_name > is being updated by < domain\username
>, in the w3wp process, on machine < servername >. View the tracing
log for more information about the conflict.
```

Action

Clear the Microsoft Office SharePoint Server file system cache on all server systems in the farm.

For details, see: <http://support.microsoft.com/kb/939308>.

After restore, you cannot connect to the FAST Search Server

Problem

After restore, when you try to connect to the Microsoft FAST Search Server 2010 system for SharePoint, the operation fails.

FAST Query SSA search operations display an error similar to the following:

The search request was unable to connect to the Search Service.

Action

Execute:

```
SharePoint_VSS_backup.ps1 -resume-cert
```

NOTE:

The VSS based solution copies the FAST Search certificate FASTSearchCert.pfx from the FAST Admin Server system to the SharePoint Server system and installs it. Also, the SharePoint certificate is copied and installed to each FAST Search Server system. For details, see: <http://technet.microsoft.com/en-us/library/ff381244.aspx>.

The SharePoint_VSS_backup.ps1 script stops responding and the farm stays in read only mode

Problem

When starting a backup, the SharePoint_VSS_backup.ps1 script stops responding when a crawl of the Microsoft SharePoint Server is being performed. The issue can appear due to external conditions such as a corrupted SSA index, the need to reissue the certificate manually and so on.

As a result, the farm stays in read-only mode.

Action

Normally, the crawl should be aborted automatically after 15 minutes. If this does not happen:

1. Abort the script by pressing **Ctrl+C**.
2. Manually resume the farm.

You can specify a different timeout after which the crawl is aborted and the farm is resumed by using the `-timeout` option.

SharePoint Search service application not operational after restore

Problem

After restoring the SharePoint Search service application (SSA) and resuming the SharePoint Server 2013 farm, the SSA status reads `Paused for:External` request instead of `Running`, indicating that the SSA is not operational.

Action

Perform the following steps:

1. Using the SharePoint Online Server Management Shell, export the SharePoint Search service application (SSA) topology:

```
$ssa = Get-SPEnterpriseSearchServiceApplication -Identity "NameOfSSA"  
Export-SPEnterpriseSearchTopology -SearchApplication $ssa -Filename  
"TopologyFilename.xml"
```
2. Record the SSA application pool identify if it exists:

```
$ssaAppPool = $ssa.ApplicationPool
```

If it does not exist, create it by executing:

```
$ssaAppPool = New-SPServiceApplicationPool -name "ApplicationPoolName" -account  
"Domain\Username"
```
3. Delete the SSA by executing:

```
$ssa = Get-SPEnterpriseSearchServiceApplication -Identity "NameOfSSA"  
Remove-SPEnterpriseSearchServiceApplication -Identity $ssa -RemoveData  
Remove-SPEnterpriseSearchServiceApplicationProxy -Identity "NameOfSSAProxy"
```
4. Using the Data Protector Microsoft Volume Shadow Copy Service integration, restore the SSA databases with the Microsoft SQL Server VSS Writer. Ensure the database names start with the name of SSA.
5. Using the SharePoint Online Server Management Shell, restore the SSA itself by executing:

```
Restore-SPEnterpriseSearchServiceApplication -Name "NameOfSSA" -ApplicationPool  
$ssaAppPool -TopologyFile "TopologyFilename.xml" -KeepId
```
6. Create the SSA proxy by executing:

```
$ssa = Get-SPEnterpriseSearchServiceApplication -Identity "NameOfSSA"  
New-SPEnterpriseSearchServiceApplicationProxy -Name "NameOfSSAProxy" -  
SearchApplication $ssa
```
7. Stop the SharePoint Search Host Controller service on all systems where the SSA indexing components are installed. Execute the command:

```
Stop-Service SPSearchHostController
```
8. Using the Data Protector Microsoft Volume Shadow Copy Service integration, restore the SSA index files with the OSearch VSS Writer.
9. Using the SharePoint Online Server Management Shell, start the SharePoint Search Host

Controller service on all systems where the SSA indexing components are installed:

```
Start-Service SPSearchHostController
```

10. Resume the SSA by executing:

```
$ssa = Get-SPEnterpriseSearchServiceApplication -Identity "NameOfSSA"
```

```
Resume-SPEnterpriseSearchServiceApplication $ssa
```

Chapter 6: Data Protector Virtual Environment ZDB integration for VMware

Introduction

This chapter explains how to configure and use the Data Protector Virtual Environment ZDB integration for VMware. It describes concepts and methods you need to understand to back up and restore the VMware vSphere virtual environment that is using NetApp or HPE 3PAR storage.

The Data Protector zero downtime backup (ZDB) functionality offers online backup capabilities with minimal degradation of the application system performance. The load on the application system is significantly reduced because backup is nondisruptive and does not require downtime for virtual machines. The tape backup of the copied data is performed using a separate backup system.

Data Protector Virtual Environment ZDB integration for VMware supports environments where ESX(i) Server systems are set up with NetApp and HPE 3PAR storage system, and managed through a vCenter Server (vCenter environments).

Data Protector supports:

- **NetApp storage:** ZDB to tape
- **3PAR storage:** ZDB to disk, ZDB to tape, and ZDB to disk+tape

During the creation of a replica, the application system is in full operation and VMware virtual environment is actively used. The streaming of the data to tape media is subsequently performed on the backup system. The backed up data can be restored using standard Data Protector restore from tape.

Data Protector offers interactive and scheduled ZDB to tape and standard restore methods.

Supported disk arrays and disk array configurations

Supported array	Supported configurations
NetApp storage systems	Systems running Data ONTAP 8.2 for 7-Mode Systems running Data ONTAP 8.2 for C-Mode
HPE 3PAR storage systems	Systems running InForm OS version 3.1.2, 3.1.3, and 3.2.1

This chapter provides information specific to the Data Protector Virtual Environment ZDB integration for VMware. For general Data Protector procedures and options, see the *HPE Data Protector Help*. For details on ZDB terminology, ZDB types, advantages of offline and online backups, see the *HPE Data Protector Concepts Guide*. For more information on Data Protector virtual environment integration, see the *HPE Data Protector Integration Guide*. For more information on non-HPE Storage Arrays and HPE 3PAR, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*. For other limitations and recommendations, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

Recommendations

- HPE recommends not to install any Data Protector components on VMware vCenter Server system or VMware ESX(i) Server system.

Integration concepts

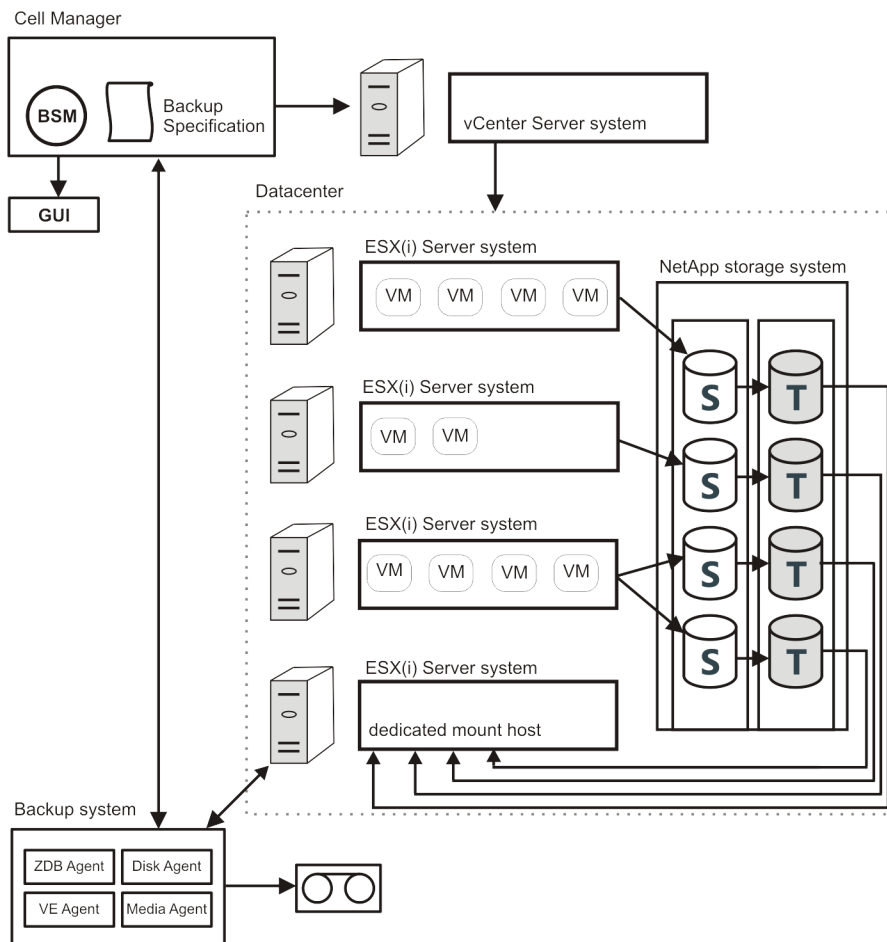
Supported environments

Data Protector supports environments where ESX and/or ESXi Server systems (ESX(i) Server systems) are managed through a vCenter Server (**vCenter environments**). Environments with standalone ESX(i) Server systems and mixed environments, in which some of the ESX(i) Server systems are managed through a vCenter Server system and some are standalone, are also supported. You can even have multiple vCenter Server systems in your environment, each managing its own set of ESX(i) Server systems.

However you set up your virtual environment, for the Data Protector Virtual Environment ZDB integration for VMware it is required that the ESX(i) Server systems are connected to the same storage system with Fibre Channel. Data Protector Virtual Environment ZDB integration for VMware is only supported in virtual environments running on NetApp, and HPE 3PAR storage system.

VMware ZDB integration architecture below shows the architecture of the Data Protector Virtual Environment ZDB integration for VMware.

VMware ZDB integration architecture



Virtual environment:

In this figure, the virtual environment is a vCenter system managing four ESX(i) Server systems. The ESX(i) servers are connected to the storage system through Fibre Channel SAN. The first three ESX(i) servers are active application systems, and the fourth one is a mount host, a system dedicated to mounting replicas, used for ZDB backup purposes only.

NOTE:

HPE recommends to use a dedicated ESX(i) Server for ZDB backup purposes (mount host).

The NetApp and HPE 3PAR storage system enables snapshot replication of the disk volumes used by virtual machines. The volumes containing the source or original data objects (S) are replicated to an equivalent number of target volumes (T). When the replication process is complete, the data in the target volumes constitutes the snapshot replica. For more information on disk replication, see the *HPE Data Protector Concepts Guide*.

The vCenter system is imported in the Data Protector Cell as **VMware vCenter** client.

Backup system:

The backup system is used to start the backup process and perform a backup of the resolved replicas to tape. It is recommended to use a dedicated physical backup system.

Data Protector components:

The following Data Protector components should be installed to enable the Data Protector Virtual Environment ZDB integration for VMware environment:

- ZDB Integration agent (NetApp Storage Provider or HPE P6000 / HPE 3PAR SMI-S Agent)

NOTE:
 The storage providers plug-ins enable ZDB integration within the Data Protector ZDB SMI-S Agent.

- Data Protector Virtual Environment Integration (VEAgent)
- Data Protector Disk Agent
- Data Protector General Media Agent

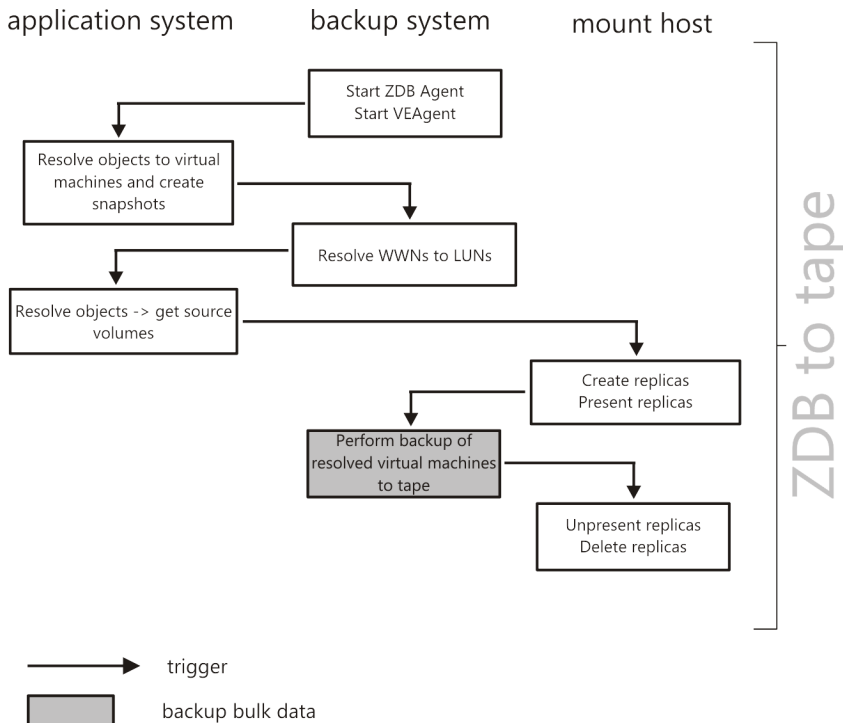
The ZDB Integration agent, the Data Protector Virtual Environment Integration component, and the Data Protector Disk Agent must be installed together on at least one Data Protector client in the cell. This client is called the backup system.

The Data Protector General Media Agent must be installed on clients that will transfer data to backup devices. It can be installed on the backup system.

Restoring virtual machines using Data Protector Virtual Environment ZDB integration for VMware is the same as restoring from a "local or network" backup type.

Backup process

VMware ZDB backup flow



See the *HPE Data Protector Concepts Guide* for a general description of ZDB concepts.

See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for a general description of the ZDB-to-tape session flow and for the explanation of actions triggered by ZDB options.

This section provides only the information relevant to the Data Protector Virtual Environment ZDB integration for VMware.

Operations on a replica described below are dependent on or triggered by ZDB options. See the *HPE Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

1. When you start zero downtime backup of your virtual environment, Data Protector starts the VEAagent and ZDB Agent on the backup system. VEAagent on the application system resolves the selected objects for backup to virtual machines and creates virtual machine snapshots.
2. VEAagent resolves the selected objects for backup to datastore WWNs and ZDB Agent resolves WWNs to LUNs.
3. ZDB Agent creates replicas of the resolved LUNs.
4. The replicas are presented to the ESX(i) mount host and backup system backs up all resolved virtual machines to tape.
5. When backup is completed, the replicas are unrepresented and deleted.

Backup concepts

What is backed up?

Using the Data Protector Virtual Environment ZDB integration for VMware, you can back up the following VMware objects:

VMware vSphere:

- Virtual machines
- Virtual machine disks
- Virtual machine templates

Virtual machines

When you back up a virtual machine, you actually back up virtual machine files of the following types:

- .vmx
- .vmdk

Virtual machine disks

Data Protector supports backup of individual virtual machine disks when using a vStorage Image backup method. In this case, all virtual machine files are backed up, except for virtual machine disks that are not specified. You can run full, incremental, and differential backups.

From Data Protector 9.05 onwards, the virtual machine disks are backed up in parallel and not serially.

To achieve the virtual machine disk parallelism, the virtual machine disks are considered as objects. When the object operations (such as object copy and object verification) are specified, the disk objects are not shown. The disks are considered when the virtual machine object is selected for the object operations.

NOTE:

After you add a new disk to a virtual machine, make sure you run a full backup session for the updated virtual machine.

Virtual machine templates

You can also back up virtual machine templates when using a vStorage Image backup method. When you create a backup specification, expand the **vm** folder and select the desired virtual machine templates.

vStorage Image backup method

The vStorage Image backup method provided by the Data Protector Virtual Environment ZDB integration for VMware is based on the VMware vStorage technology. For this method, a single central **backup host** is used to back up all virtual machines hosted by ESX(i) Server systems in a Data Protector cell.

During a ZDB backup, VEAgent performs vStorage Image backup. VEAgent first establishes connections between the backup system and the virtualization host (an ESX(i) Server system). This connection can either be through a vCenter Server system (in the case of a vCenter environment), or direct (in the case of a standalone ESX(i) Server environment). It then requests a snapshot of the virtual machine that is to be backed up, via the vStorage API for Data Protection (VADP). This ensures that the virtual machine is in a consistent state. ZDB Agent then creates a replica of the virtual machine disks. The replica is presented to the ESX(i) mount host, Media Agent client is initialized and the data for backup is streamed to tape.

Snapshot management

The vStorage Image backup method relies on being able to create virtual machine snapshots. A virtual machine snapshot is an operation that puts the virtual machine into a consistent state. All subsequent changes made to the virtual machine disks are recorded to the created snapshot.

NOTE:


The snapshot operation is not supported by all virtual machine disks. For example, snapshots of independent disks are not supported; consequently, this type of virtual disks cannot be backed up using the Data Protector Virtual Environment ZDB integration for VMware. For details, see the VMware documentation.

During a vStorage Image backup, Data Protector creates a snapshot, replicates virtual machine disks and copies the data from the consistent state to Data Protector media. Then Data Protector deletes the replica and the snapshot. Note that snapshots created by Data Protector (**DP snapshots**) are distinguished from other snapshots by the label `_DP_VEPA_SNAP_` containing the product name, a description, and a timestamp.

NOTE:

If a virtual machine has a user-created snapshot during CBT enabled backup, the user-created snapshot is backed up along with the other VM disk blocks. If the Data Protector detects a user-created snapshot at the time of restore in a virtual machine, the VM is not restored. To restore such a VM, you need to manually delete all the existing user-created snapshots.

Virtual machine snapshot tree



Existing virtual machine snapshots reduce the overall performance of a virtual machine. For this reason, Data Protector automatically removes DP snapshots once they are no longer needed.

IMPORTANT:

Do not use the label `_DP_VEPA_SNAP_` for snapshots that you create for other purposes, otherwise they will be deleted by Data Protector.

The replicas created by the ZDB Agent are always deleted after ZDB backup is completed.

Backup types

The type of backup to be performed is specified at backup specification level, either in the Scheduler page, or in the Start Backup dialog box for an interactive backup.

Using the vStorage Image backup method, you can perform the following backup types:

Backup types

Backup Types	Description
Full	Backs up the complete virtual machine (disk).
Incremental	Backs up the changes made to a virtual machine (disk) since the last full, incremental, or differential backup.
Differential	Backs up the changes made to a virtual machine (disk) since the last full backup.

For incremental and differential backup sessions, you must also specify how Data Protector should identify changes at disk block level.

To identify changes at disk block level, Data Protector uses the VMware changed block tracking functionality. For details, see [Changed block tracking, on the next page](#).

NOTE:

The number of snapshots remaining after backup is always 0. Only the Mixed snapshot handling mode is supported.

Mixed snapshot handling mode supports full, differential, and incremental backups in all possible backup chain forms.

For detailed description of snapshot handling modes, see *HPE Data Protector Integration Guide*.

When performing snapshot operations on a backed up virtual machine, you must be careful not to break your backup chains.

IMPORTANT:

A VMware object's snapshot which was not created with Data Protector cannot be used to set up a Data Protector backup chain (restore chain) for that object.

A backup chain gets broken if you perform any of the following operations:

- Delete a snapshot
- Revert to a snapshot
- Create a snapshot without involving Data Protector
- Change snapshot handling mode
- Add a new virtual machine disk or rename an existing one
- Restore a virtual machine
- Enable changed block tracking

After completing any of the above operations, you must first run a full backup to start a new backup chain. If you run a session for incremental or differential backup instead, the Data Protector switches the VEAgentdisk objects to have the effective backup type as full, whereas for VEAgent object, the backup type still remains as incremental or differential. This may create a backup chain with multiple sessions during restore; which impacts performance. So, it is recommended to execute a full backup.

Changed block tracking

Changed block tracking (CBT) is a feature of later versions of VMware that can be used to improve the efficiency and speed of your backups.

For CBT, change IDs are used. A change ID is an identifier for the state of a virtual disk at a specific point in time. It is saved by the virtual disk logic whenever a snapshot is taken of the disk.

The main advantage of using changed block tracking is most noticeable on incremental or differential backups, because:

- It is not necessary to keep virtual machine snapshots on the system until the next backup, greatly reducing the system overhead.
- The changes to be backed up are calculated more easily, by obtaining change information from the kernel, rather than calculating it from snapshots.

During full backups, only active blocks on the disk are backed up, and unallocated blocks are ignored. This makes the backups space-efficient and faster.

When changed block tracking is enabled, a virtual machine's performance is slightly impacted, but this is small relative to what you gain. You can enable it using the Data Protector GUI, if required.

NOTE:

The CBT functionality is supported for **vStorage Image + OpenStack** backup method.

When changed block tracking is used, Data Protector snapshots are still used to put a virtual machine in a consistent state. However, when the backup completes, they are deleted. Only changed block tracking log files change IDs are kept.

NOTE:

When working with changed block tracking, keep a note of the following points:

- Ensure that the prerequisites from VMware are met while using CBT backup. For more information, go to <https://kb.vmware.com/kb/1020128>
- Not all types of virtual disk support changed block tracking. In case of an unsupported disk, a virtual machine backup fails.
- When changed block tracking is first enabled, the next backup for a virtual machine will always be a full backup to provide a reference point for the tracking. In other words, a new backup chain is started.
- A CBT backup chain (Full, Differential, Incr,...) gets broken when you perform a restore session. When the restore session completes, run a full backup again to start a new backup chain, otherwise subsequent incremental and differential backup sessions will fail.

Backup flow

1. Data Protector triggers a snapshot.
2. Replicas of source volumes are created.
3. A change ID is recorded for the current backup.

If this is the first snapshot taken after changed block tracking was enabled, all active blocks are identified and change ID 0 is recorded.

In case of a full backup, this change ID becomes the start reference point for a new backup chain.


4. This step depends on the backup type selected:
 - *Full*: The blocks changed since change ID 0 are identified.
 - *Incremental*: The blocks changed since the change ID for the previous backup (full, incremental, or differential) are identified.
 - *Differential*: The blocks changed since the change ID for the previous full backup are identified.
5. The identified blocks are backed up.
6. The replica and the snapshot are deleted.

Example of a backup chain with changed block tracking

Snapshot	Change ID	Blocks identified	Blocks backed up
1st after CBT enabled	ID 0	All active blocks	—

Snapshot	Change ID	Blocks identified	Blocks backed up
Full backup	ID n	Changed since ID 0	All active blocks from ID 0 + changed blocks since ID 0
Incremental backup	ID $n+m$	Changed since ID n	Changed block since ID n
Incremental backup	ID $n+p$	Changed since ID $n+m$	Changed blocks since ID $n+m$
Differential backup	ID $n+q$	Changed since ID n	Changed blocks since ID n
Full backup	ID r	Changed since ID 0	All active blocks from ID 0 + Changed blocks since ID 0

Non-Changed Block Tracking (Non-CBT) backup



Non Changed Block Tracking (Non-CBT) backup is a feature which does not depend on block level changes for backup.

- In this feature, all the blocks of the virtual machine disks are backed up.
- This feature does not use the VMware CBT functionality to identify the modified blocks for backup.
- The size of the backed up image increases, as all the blocks of disk are backed up.

Allow fallback to non-CBT backups option is enabled when change block tracking backup fails.

The non-CBT backup can be used in the following scenarios:

- When the hardware version of virtual machine is lesser than 7.
- When backing up virtual machines without the older version of operating system installed (example, Windows 2003).
- When snapshots are available on virtual machine and CBT is not enabled.

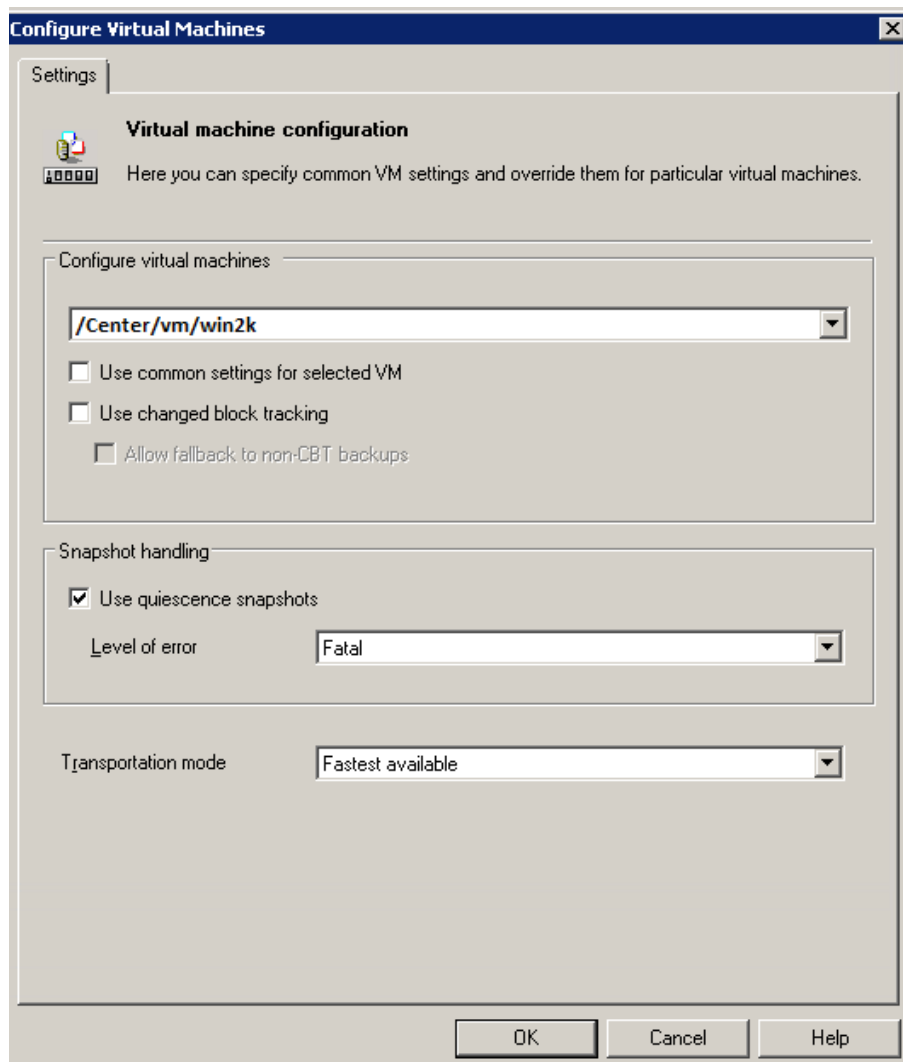
Limitations

- Only full backups are available for non-CBT backups.
- Incremental or differential backups are not supported.

Quiescence

If quiescence is selected, the snapshot process quiesces all the system writers, and the registered application writers. In Windows guest operating systems, the VSS framework freezes or quiesces the applications running in a virtual machine before a snapshot is created. Data Protector performs application-consistent quiescence every time the **Use Quiescence** option is selected for the backup.

The illustration below shows the options available for the Quiescence functionality to quiesce applications with VSS writers.



Select the **Use quiescence snapshots** check box. You can select the level of error to be reported. The following levels of errors can be selected:

- **Fatal:** If the quiescence snapshot fails, the session will fail .
- **Warning:** If the quiescence snapshot fails, a warning message is displayed and the backup process is continued.

For CLI options, see the **CLI Reference Guide**. Navigate to *Section 1M: Administrative commands > vepa_util.exe(1M)*.

NOTE: When you enable the Quiescence option, backup of virtual machines are application consistent for MS SQL, MS Sharepoint, MS Exchange and Oracle. The consistency works best when the applications are configured as recommended by their vendors. Data Protector recommends installing a respective integration agent within a virtual machine to protect these applications.

Prerequisites

- For VEPA-based quiescence backups of the Share Point server, you must register the VSS Writer before starting the backup of virtual machines. For more information, see the *Integration guide for Microsoft Volume Shadow Copy Service, Microsoft SharePoint Services writer specifics* section.
- For VEPA-based quiescence backups of the Oracle 11g Release 2 database, the Oracle VSS writer service should be in the active/running state, before starting the backup of virtual machines. This can be done by executing the command `oravssw /q /start`. For other applications such as SQL and SharePoint, the VSS writers are registered and active by default.

Limitations

- The Quiescence feature can slow down the speed of backup sessions considerably.
- The Microsoft SQL database on the Availability Group cluster application in the Windows OS cluster is not supported for VEPA quiescence.
- The cluster applications using SCSI controllers for disks in shared mode, are not supported for VEPA backups.
- vRDM disks can be used only for full backups.
- Quiescence backups for virtual machines in the Power Off state are not valid.

Considerations for Quiescence Operations

VMware considerations

- Do not disable the UUID attribute for virtual machines.
- The virtual machine must use only SCSI disks. Application-consistent quiescing is not supported for virtual machines with IDE disks. There must be as many free SCSI slots in the virtual machine as the number of disks.
- Physical RDM's don't support snapshots, therefore they cannot be used for quiescence.
- Quiescence is not supported for backups of virtual machines on the Microsoft cluster. For more information, see [Backup on virtual machines configured with bus-sharing](#).

Virtual machine considerations

- The virtual machine must not use dynamic disks.
- Ensure that the latest version of VMware Tools are installed in the virtual machine. For more information, see [Verifying a VMware Tools build version](#).
- VSS components must be explicitly specified during the VMware Tools upgrade process. VSS will not be installed in a non-interactive mode. For more information, see [Installing the Volume Shadow Copy Service with VMware Tools](#).
- Ensure that you are using Windows Server 2008 or higher. Previous versions of Windows, such as Windows XP and Windows 2000, do not include VSS and rely on the SYNC driver.
- The Distributed Transaction Coordinator service must be running during the installation of VMware Tools. Otherwise, VSS fails to quiesce Windows 2008 R2/ Windows 2012.
- Ensure that all the appropriate VSS application services are running, and the startup types are listed correctly.

For more information on VSS quiesce related issues, see [Troubleshooting Volume Shadow Copy \(VSS\) quiesce related issues](#).

Disk space requirements

A virtual machine backup requires sufficient disk space for snapshots and replicas on a storage system where the virtual machine disks reside.

Two replica provision types are available for managing storage resources:

- thin provisioned
With thin provisioning you can conserve storage resources. The disk space allocation is dynamically adapted to accommodate the demand for storage. It allows free space sharing between LUNs and enables LUNs to consume only the space they actually use.
The required free space is based on the delta file created by a Data Protector snapshot. Only changes made since the last backup session are written into the replica
- fully allocated
With this provision type space-reserved LUNs and snapshot copies have pre-allocated space that can be continually overwritten. This guaranteed space is not available to any other LUNs or snapshot copies within the volume.
The required free space is calculated based on the size of virtual machine disks just before the snapshot is created. The overall disk space required is double the space required for snapshots. A replica requires as much space as the source LUN. If the target is not space-reserved, ensure that the volume has enough free space to accommodate the replica.

Free space required option

You can use the Data Protector **Free space required %** option to make sure that a virtual machine is backed up only if there is enough free space.

The required free space is calculated based on the size of virtual machine disks just before the snapshot is created. Data Protector checks all datastores where the virtual machine disks reside. If one of the datastores does not meet the specified percentage of free space, no snapshot is created and the backup of the virtual machine fails with an error.

When backing up more than one virtual machine, the check is applied to each virtual machine separately. The virtual machines which pass the check are backed up and the ones which do not pass it are not.

If you specify 0%, the check is omitted.

Examples

The following examples illustrate how the **Free space required %** option works:

- Backup of a single virtual machine "test1" with the disk "disk1" residing on the datastore "datastore1":
 If you specify 30% in the **Free space required %** option and the datastore has a size of 100 GB, the backup succeeds if there is at least 30 GB of free space on the datastore.
- Backup of a single virtual machine "test1" with two disks, "disk1" and "disk2" residing on two datastores, "datastore1" and "datastore2":
 If you specify that 30% of free space is required, the backup succeeds if there is at least 30 GB of free space on each datastore.
- Backup of two virtual machines, the virtual machine "test1" with the disk "disk1" on the datastore "datastore1" and the virtual machine "test2" with the disk "disk2" on the datastore "datastore2":
 If you require 30% of free space, the backup of both virtual machines succeeds if there is at least 30 GB of free space available on each datastore. If for example, the datastore "datastore1" has less than 30% of free space and the datastore "datastore2" has at least 30% of free space, the backup of the virtual machine "test1" fails and the backup of the virtual machine "test2" succeeds. If both datastores have less than 30% of free space, the backup of both virtual machines fails.

Disk space requirements

Backup methods	Required disk space on datastores	Explanation
vStorage Image	The sum of the sizes of all the virtual machine disks, plus: <ul style="list-style-type: none"> The size of any quiescence zip files, if quiescence is specified. 	When a virtual machine snapshot is taken, changes made to the virtual machine disks are recorded to separate files (one delta file is created for each virtual machine disk). A delta file can grow up to the original virtual disk size.

Backup disk buffer

You can specify a disk buffer for your backup using the omnirc option `OB2_VEAGENT_BACKUP_DISK_BUFFER_SIZE`.

The SAN and the HotAdd backups support disk buffer sizes from 1 MB to 256 MB. By default, their disk buffer size is 8 MB. However, network backups, such as NBD and NBD (SSL), are always performed using the default disk buffer size of 1 MB.

NOTE:

- If there is not enough memory available for the specified disk buffer size, a fallback to 1 MB disk buffer size will happen to keep the backup running, and a warning message will be

displayed.

- Using bigger disk buffer sizes improves the backup performance, but it also increases the memory consumption. At the certain level the backup performance does not improve anymore due to the limits of your backup host.

Backup parallelism

By default, virtual machines are backed up in parallel. In rare cases this may lead to problems. For example, backup sessions may end unexpectedly. In such cases, you may set the Data Protector `OB2_VEAGENT_THREADED_BACKUP` omnirc option on the backup host to `0` to disable parallel backups.

NOTE:

Virtual machine disks are backed up sequentially in both cases.

For details on how to use Data Protector omnirc options, see the *HPE Data Protector Help* index: “omnirc options”.

Backup considerations

- **Change Block Tracking (CBT) and Mixed snapshot handling mode**

CBT backup method and mixed snapshot handling mode are supported.

- **Concurrent backup sessions**

Backup sessions that use the same devices cannot run in parallel.

You cannot back up virtual machines in parallel with an ESX(i) Server system or a datacenter where the virtual machines reside.

Backup sessions that backup virtual machines in the same data store cannot be run in parallel either in the same cell manager or different cell manager.

- **Transportation modes**

You can use various transportation modes for backup. For details, see [Configuring the integration, on page 335](#).

It is recommended to use CBT for making incremental/differential backups because it is faster and uses less space on the backup device.

The transportation mode can either be SAN or NBD, and this can be selected from the Data Protector GUI (This determines how the array will be accessed). The transportation mode can also be configured in the virtual machine options. The transportation mode configured in the virtual machine options is executed on priority, and it follows the order

`SAN:HOTADD:NBDSSL:NBD:FILESYSTEM`.

For example,

- If you have selected NBD in the virtual machine options from the Data Protector GUI
- If SAN is not available
- If the HOTADD solution is not possible for the zero downtime backup

Then the backups will go through NBDSSL. However, if you want the backups to go through NBD transport mode, then you must configure the NBD transport mode in the virtual machine options.

A template backup in the SAN transport mode is not supported, and will fallback to the NBD transport mode. Also, if replica cannot be presented to the backup host when SAN transportation mode is selected, VM backup will fall back to NBD transportation mode.

- **Replica provision type**

To manage storage resources you can select thin provisioned or fully allocated replica provision type to ensure enough free disk space on a storage system.

- **Presentation of LUN**

For zero downtime backups of virtual machines from the 3PAR replica, ensure that the LUN that is used to create the source data store (where the virtual machine to be backed-up resides) is not presented to the system that is configured as the mount proxy host.

- **Backup to StoreOnce Catalyst device**

From 9.07 onwards, all VEPA backups to StoreOnce Catalyst device are performed with "Single Object Per Media Store" mode. Even if this option is not selected on the StoreOnce Catalyst device, this mode will be enforced.

Any value that you enter in the "Store Media Size Threshold (GB)" field will be ignored to enable Cached GRE or Power On and Live Migrate from the backups done to StoreOnce Catalyst device.

- **Data Protector licenses**

The following licenses are not required for performing zero downtime backups of virtual machines from the 3PAR replica:

- HPE Data Protector instant recovery extension for UNIX - 1 TB
- HPE Data Protector instant recovery extension for Linux - 1 TB
- HPE Data Protector instant recovery extension for Windows - 1 TB

Restore concepts

You can restore VMware objects backed up with vStorage Image backup method in different ways.

Restore of VMware objects backed up with vStorage Image method

Virtual machines, virtual machine disks, and virtual machines templates backed up with the **vStorage Image** method can be restored:

- To a datacenter
- To a directory on a backup host

Restore to a datacenter

By default, virtual machines are restored to the original datacenter and the original datastore, but you can select a different datacenter if you want.

By default, Data Protector deletes a virtual machine (if it exists) before it is restored, even if it resides in a different datacenter from the datacenter you restore to.

NOTE:

If you select an ESX(i) Server client in the Restore client option (destination client) in the restore wizard, the migrated virtual machine will not be deleted, as the ESX(i) Server client is not able to detect virtual machines that are located on different ESX(i) Server clients (only a vCenter client can do that). Consequently, you end up with two virtual machines having the same UUID.

Alternatively, you can choose to restore virtual machines only if they do not exist, leaving existing virtual machines intact.

For the restore, you can also specify the following:

- Whether the restored virtual machines should be registered in the datacenter
- Whether the restored virtual machine snapshots should be consolidated when the restore completes
- Whether the restored virtual machines should be powered on

The restore options provided are, by default, set to restore virtual machines to the original datacenter.

You can restore virtual machines and virtual machine disks, from a replica to a data center. Note that restore sessions from replica to a directory is not supported. In **Disk + Tape** backups, if the replica is rotated or removed by the administrator, the restore happens from the tape.

Restore of individual virtual machine disks

To be able to restore individual virtual machine disks to a datacenter, the original virtual machine must still exist. Otherwise, the restore fails.

Here is the progress of a restore session:

1. The virtual machine is powered off.
2. If the disks to be restored still exist, they are removed.
3. The disks are restored from the backup.

NOTE:

After restore, virtual disks that are part of a dynamic disk set or virtual disks from different points in time may require additional user action (for example, mounting, resignaturing, or recovery) in the guest operating system and/or applications running inside.

Restore to a directory

When restoring to a directory (restore outside a datacenter) all the files of virtual machines are restored to a directory of your choice (for example, C:\tmp) on a backup host.

In the directory you specified, subdirectories are created with names corresponding to those of the datastores where the virtual machines (their virtual disks) resided at the time of backup. The files related to the virtual disks are restored to the respective subdirectories.

After such a restore, the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX(i) Server system, using the VMware Converter as described in [Recovering virtual machines after restore to a directory, on page 381](#).

Restore of Nova Instances and Shadow VMs backed up with vStorage Image + Openstack method

Progress of a restore session:

1. IDB is queried to obtain the list of Shadow VMs attached to the Nova Instance.
2. vCenter is queried to create the map of Nova Instance and Shadow VMs.
3. Related Shadow VMs are added into the restore object list.
4. vCenter is queried to create the map ([Step 3](#)) and validated if the Shadow VMs are attached to the original instance.
5. Queried the restore version from IDB to obtain the restore chain.
6. Shadow VM files are removed from the vCenter.
7. Nova Instance files are removed from the vCenter.
8. Restored the Nova Instance and disk files in the same folder structure.
9. Restored the Shadow VM configuration files.
10. Registered the virtual machine and restored the network.

NOTE:

If the selected Nova Instance is available in the vCenter during the restore, the Shadow VMs are detached before deleting the Nova Instance.

After the Data Protector restore, to bring back the OpenStack instance in the OpenStack Horizon dashboard to the correct state, perform the following steps:

1. Restart the Nova Compute service in the Nova Proxy node by executing the following command:

```
Service nova-compute restart
```
2. Refresh the Horizon Dashboard to check if the Nova Instance is available or not. If the Nova Instance is not available, connect to the OpenStack Management node and execute the following commands:
 - `Nova list`: Lists the Nova Instances.
 - `Nova reset-state -active "instance-uuid"`: Resets the state of the Nova Instance to active.
 - `Nova reboot "instance-uuid"`: Reboots the Nova Instance.

Once the reboot is done, refresh the Horizon Dashboard and check for the availability of the Nova Instance.

Restore chain

When you restore a virtual machine from a backup created in an incremental or differential session, Data Protector automatically restores the complete backup chain, starting with the last full backup, which is then followed by the differential and all subsequent incremental backups (if they exist) up to the selected session.

Power On and Live Migrate

Power On

Virtual machines can be powered on instantly within seconds from the Data Protector backup image that resides on the 3PAR replica (local or remote copy), Smart Cache, and StoreOnce Catalyst devices. Previously, virtual machines had to be powered on only after the complete data migration to the production data center. Use this feature if you want to verify the sanity of the backup. Note that the changes done to the virtual machine once it is powered on, will be available until you perform the clean up operation.

When you power on a virtual machine, the backup image is presented to the destination ESX server. A new virtual machine is created, whose data disks point to the Data Protector backup image. The other files reside on the destination data center.

Live Migrate

This option will power on the virtual machine from the backup image, and will simultaneously start the data migration to the destination datastore. During this process, the virtual machine will continue to be accessible. Since the data movement is a back end operation, it will have minimum impact on the usage and accessibility of the powered on virtual machine. Any modifications done to the virtual machine data will be consolidated, and the migrated virtual machine will have all the modified content on top of the restored image from the backup.

Once the data migration is complete, the virtual machine functions from the destination datastore, and has no dependency on the backup image. Also, the backup image presentation is removed.

NOTE:

When using Power On and Live Migrate on windows backup host, ensure that Data Protector INET service and Data Protector Filter listener service are running under same user credentials.

NOTE:

Power On and Live Migrate operations are available, only if the original backup resides on one of the supported devices. If the backup resides on a non-supported device and an object copy is performed to a supported device, then these operations are not available until the original session exists.

Only when the original backup has expired and the copy is promoted as original, the Power On and Live Migrate operations are available from the copied session (if on a supported device).

For the list of supported devices for Power On and Live Migrate, see [Supported backup types with GRE, Power On, and Live Migrate, on page 334](#).

StoreOnce Catalyst device only

- Power On and Live Migrate operation from full, incremental, and differential backups are supported.
- Power On and Live Migrate operation from object copy is supported for 9.05 and 9.06 backups. The object copy should be performed on per session basis to ensure data consistency.

NOTE:

If you want to migrate your Data Protector 9.05 or 9.06 version backups to StoreOnce Catalyst to use the feature of Power On and Live Migrate, it is recommended to perform

object operations at individual session. If you choose multiple sessions at once, data consistency will not be there.

- If the virtual machine is powered on from the StoreOnce Catalyst device, the virtual machine must be cleaned up before performing the Live Migrate operation.

Cleanup/Power Off

Data Protector stores the list of all the powered on virtual machines. It stores the detailed information about all powered on virtual machines as an XML file in the Cell Server. The cleanup and power off actions are listed below:

- If the virtual machine is already powered on for more than 24 hours, it will be powered off, and the related storage will be cleaned up.
- If the virtual machine is powered on from the replica, the data store will be un-mounted and the replica created during the power on process will be removed from the array.
- If the virtual machine is powered on from the Smart Cache or StoreOnce Catalyst device, during the clean up process, the data store will be removed, and the NFS share will be deleted.

Note that the above actions are applicable to virtual machines that are powered on using the **Power On** feature.

The virtual machines that are powered on for more than 24 hours will undergo a cleanup operation during the next daily maintenance job. All the virtual machines that are cleaned up in daily maintenance job are logged in the `poweronvms_cleanup.log` file.

For more information on the procedure of Power On and Live Migrate, see [Restoring using the Data Protector GUI](#).

Note that Power On and Live Migrate is not supported with incremental and differential backups on a Smart Cache device.

Restore considerations

• Concurrent sessions

Restore sessions that use the same devices cannot run concurrently.

• Failed restore sessions

Sometimes, when a virtual machine restore fails, Data Protector creates extra files on the datastore which you need to clean up manually when the session completes. Otherwise, corrupt virtual machine backups may be created in subsequent sessions and restore from a such a backup also fails. For details, see [Cleaning up a datastore after a failed restore, on page 388](#).

When restoring a virtual machine to a non-original datastore whose block size is not compatible with the virtual machine disks' sizes (that is, a `.vmdk` file size is not a multiple of the datastore block size), the restore fails.

• Virtual machines in vApp

When you restore a virtual machine that resided in a vApp container at the time of backup, the virtual machine is not restored back to the vApp container, but to the ESX(i) Server root level. If the virtual machine in the vApp container still exists, it is deleted or the restore is skipped, depending on what you select in the Existing virtual machine handling option.

• Partial restore from a vStorage Image backup

When performing a partial restore from a vStorage Image backup (for example, when restoring only some out of many backed up VM disks), the default option **Delete after restore** is ignored and the **Delete before restore** option is used instead.

- **Transportation modes**

The following recommendations for specific virtual machine transportation modes apply:

- SAN transportation mode: To use the SAN transportation mode for restore:
 - Select a backup host for a restore session.
 - Ensure that the storage volumes that are presented to both the backup host and ESX(i) Server systems are not read-only. For details on how to check storage volume properties, see [A restore session using SAN transportation mode fails](#).
 - Ensure that the storage volume size is a multiple of the underlying VMFS block size. Otherwise, the Write operation to the remainder fails. For example, if the storage volume size is 16.3 MB and the block size 1 MB, writing to the remaining 0.3 MB fails. For details, see the VMware Knowledge Base at:
<http://kb.vmware.com/selfservice/microsites/searchEntry.do>.
Search for “Best practices when using SAN transport for backup and restore”.
- It is recommended to use CBT for making incremental/differential backups because it is faster and uses less space on the backup device.
- Hotadd transportation mode: Hotadd transportation mode is available for restore, however VMware does not support multiple disks. Therefore, in a HotAdd environment use the omnirc option `OB2_VEAGENT_RESTORE_TRANSPORT_METHOD` to set the restore transport mode to NBD.

Power On considerations

- Install the following NFS packages on the backup host :
 - For Smart Cache backups: NFS version 3 or later
 - For StoreOnce Catalyst backups: NFS version 4 or later
- Data Protector uses the Windows PowerShell script `nfsServiceCheck.ps1` to initiate the NFS service, which is required for the Power On process. Execution of this script requires the execution policy to be set to *RemoteSigned*. If you need the policies to be *Restricted*, set the omnirc variable `OB2_NO_NFSSERVICE_CHECK` to 1. In case the script fails, you need to execute it manually. Some of the possible reasons for the failure of the NFS service installation may be:
 - NFS port is used by another application.
 - Powershell may require a reboot.
 - NFS module not present in the Powershell repository.

Use the following command to execute the NFS service manually:

- Windows Command Line: `powershell.exe NFSServiceCheck.ps1`
- PowerShell Command Line: `NFSServiceCheck.ps1`

- The backups of non-persistent virtual machines that are created as part of the Power On feature will be skipped. The following message is displayed when you try to perform a backup operation: *Non Persistent Powered on Virtual Machine found, Skipping Backup.*
- Network will be disabled in the virtual machine that is powered on from the Smart Cache or StoreOnce Catalyst devices.
- To manually clean up all the VMs that are powered on in your Cell Manager, perform the following:
 1. Set the omnirc variable `FORCE_PURGE_POWERON_VMS` to 1 on all the backup hosts.
 2. Run the following command in the Cell Manager:

```
/opt/omni/sbin/omnidbutil -purge_expired_poweron_vms
```

NOTE:

All powered on vm's will be powered off and removed.

HPE 3PAR storage systems only

- You must install the `vmfs-tools-0.2.5` in the Linux mount proxy host. .
- You must ensure that multipath services are running in the Linux mount proxy host.
- You must install the `sg3_utils rpm` package in the Linux mount proxy host.
- To perform GRE operations from the 3PAR replica, the GRE mount proxy host and the source ESX server must be present in the same 3PAR zone.

StoreOnce Recovery Manager Central Integration

StoreOnce Recovery Manager Central (RMC) software integrates HPE 3PAR StoreServ primary storage with HPE StoreOnce Backup systems. RMC integrates 3PAR StoreServ primary storage and StoreOnce Backup for converged data protection that delivers assured recovery of application-consistent recovery points with flexible recovery options.

The Express Protect feature offers a second-tier of data protection by facilitating direct backup from 3PAR StoreServ to StoreOnce device, independent of backup software. Backups to StoreOnce are self-contained volumes, deduplicated to save space, and can be used to recover back to the original or a different 3PAR StoreServ array, even if the original base volume is lost. The Express Protect feature facilitates direct backup from the primary storage to the backup storage completely removing the application server from the data path.

HPE StoreOnce RMC for VMware enables you to protect VMware Virtual Machine Disks (VMDKs) and data stores using application-consistent snapshots for rapid online recovery.

With Data Protector, you can backup virtual machines snapshots created by RMC to Data Protector supported secondary storage devices. You can then perform a restore operation to the required destination. Note that Granular Recovery Extension (GRE) operations are possible only for Snapshot+Tape backups.

Data Protector supports the following types of backup:

- **Snapshot backups:** With Snapshot backups, you can create snapshots of the original volumes.
- **Snapshots+Tape:** With Snapshot+Tape backups, you can create snapshots, and backup the data to Data Protector supported secondary storage devices.

- **Express Protect backups:** With Express Protect backups, you can backup snapshots from HPE 3PAR StoreServ to HPE StoreOnce.

Supported backup types with GRE, Power On, and Live Migrate

The following table lists the supported backup types, backups that support GRE, Power On, and Live Migrate operations.

Supported backups	Backup type	GRE	Power On and Live Migrate
Snapshot	<ul style="list-style-type: none"> • Full 	Not supported	Not supported
Snapshot+Tape	<ul style="list-style-type: none"> • Full 	If the tape device is Smart Cache or StoreOnce Catalyst, cached GRE is supported. For any other type of tape devices, non-cached GRE is supported.	Supported, if the tape device is Smart Cache or StoreOnce Catalyst.
Snapshot+Tape	<ul style="list-style-type: none"> • Incremental • Differential 	If the tape device is StoreOnce Catalyst, cached GRE is supported. For any other type of tape devices, non-cached GRE is supported.	Supported, if the tape device is StoreOnce Catalyst.
Express Protect	<ul style="list-style-type: none"> • Full • Incremental 	Not supported	Not supported

RMC integration considerations

- Single datastores consisting of multiple LUNs are not supported for RMC backups.
- The RMC server details that are provided when you create the backup specification cannot be modified once the specification is saved.
- For RMC integration backups, if automation scripts are used for creating barlists, ensure that the recovery set name, specified for barlist creation, is unique.
- For RMC Express Protect backups, Data Protector maintains 2 snapshots internally. This is required for executing incremental backups. When the snapshot count exceeds 2, Data Protector rotates the older snapshot.
- RMC Express Protect backups are supported only from Inform OS version 3.2.1 MU1 onwards.
- The RMC schedules and backup reports are supported only in Windows and Linux Cell Manager platforms.
- For RMC restore sessions from snapshot, array credentials must be added using the `omnidbzd` command.
- For RMC restore sessions, only Data Protector tape restores can be restored to a directory.

- For RMC Snapshot and Express Protect backups, the session information is not displayed in the GRE GUI.
- Backup and Application hostname should match the name of the host on 3PAR as seen from 3PAR Management console or by running 3PAR CLI command showhost. The hostname is case-sensitive.

Prerequisites

You should be familiar with the StoreOnce RMC concepts and procedures. For more information on RMC, see the *HPE StoreOnce Recovery Manager Central User Guide* and *HPE StoreOnce Recovery Manager Central for VMware User Guide*.

Limitations

Express protect restore: If a restore session is aborted when express protect is restoring data from StoreOnce to a new snapshot, the Data Protector session is aborted. However, restore from StoreOnce to a snapshot continues in the background. You need to wait until the RMC restore operation is complete, before another Data Protector restore operation from the same session is triggered.

RMC Integration Process

1. **Adding the RMC server details:** Add the RMC server using the Command Line Interface. Execute the command `omnidb -addhost -servername <rmcservername> -user <username> -passwd <password>`.
2. **Backup:** Perform the backup operation by selecting the **Backup Type** as StoreOnce RMC in the Data Protector GUI.
3. **Restore:** Perform the restore operation or the GRE operation using the GUI.

Configuring the integration

Configure the integration as follows:

- Import VMware clients into the Data Protector Cell.
- Configure virtual machines you want to back up.

Recommendations

- HPE recommends not to use the percent sign in names of virtual machines that will be backed up, restored and recovered with Data Protector. If a virtual machine name contains %, the name is displayed incorrectly in the Data Protector GUI and in the Data Protector session messages.

Prerequisites

- Make sure that you have a correctly installed and configured VMware vSphere environment. For supported versions, platforms, devices, and other information, see the latest support matrices at <https://softwaresupport.hpe.com/>.

- Ensure you have **Disable Methods** and **Enable Methods** Global privileges granted to the user account that is used to connect to the vCenter Server.
- Make sure you have the necessary VMware vSphere privileges granted to the user account that is used to connect to the vCenter Server. For details, see [Importing and configuring VMware clients, below](#).
- Make sure that you have correctly installed Data Protector. For information on installing Data Protector in various architectures, see the *HPE Data Protector Installation Guide*.

Data Protector Virtual Environment ZDB integration for VMware must be correctly installed and configured. Make sure that you have at least one client with the ZDB Agent (Storage Provider component for the non-HPE Storage Array or the HPE 3PAR Storage Provider component) and Virtual Environment Integration component installed (**backup system**) in your environment. For configuration, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

IMPORTANT:

The client you intend to use as the backup system should not have the VMware Consolidated Backup (VCB) software installed.

If you intend to restore virtual machine files to a directory on a backup host, also install the Disk Agent component on the backup system. Otherwise, you will not be able to use the **Browse** button to specify the target directory (however, you will still be able to type the directory by yourself).

Before you begin

- Configure devices and media for use with Data Protector.

Importing and configuring VMware clients

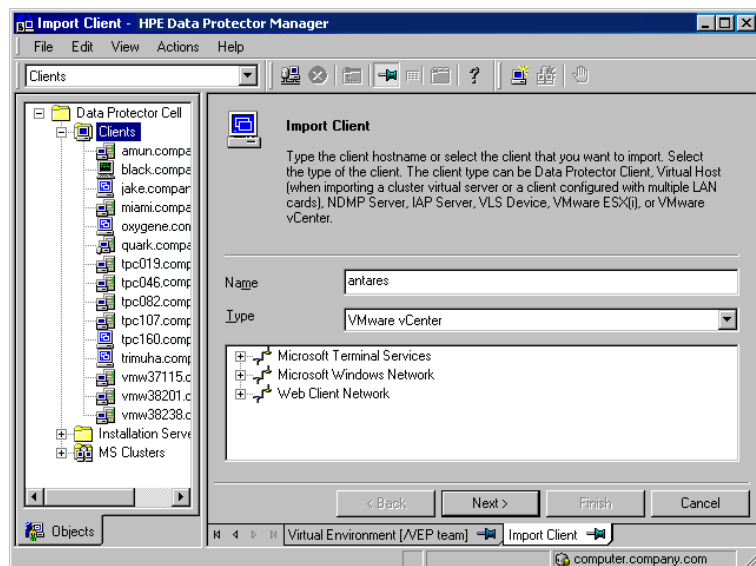
With the Data Protector Virtual Environment ZDB integration for VMware, it is not necessary to install any Data Protector components on the VMware clients (vCenter Server systems, ESX(i) Server systems). To make them Data Protector clients, the VMware clients must be properly imported into the Data Protector cell and configured.

Procedure

To import a client into a Data Protector cell:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell**, right-click **Clients**, and select **Import Client**.
3. In the Import client page, enter the client name in the **Name** option, select the appropriate client type (**VMware ESX(i)**, **VMware vCenter**) from the **Type** drop-down list and click **Next**.

Importing a VMware vCenter Server client (Name and Type)



- If you select **Standard security**, you need to manually specify which login credentials Data Protector should use to connect to the VMware client:

Port : Specify the port that VMware vSphere is using. By default, VMware uses the port 443. For VMware ESX(i) Server that you will use as a mount host, specify the port 443.

Username and Password: Specify an operating system user account that has the following VMware vSphere privileges on root vCenter level:

The privileges shown in the table are as seen in VMware vSphere 5.5. and VMware vSphere 6.

Datastore -> Allocate space
Datastore -> Browse datastore
Datastore -> Low level file operations
Datastore -> Remove file
Datastore -> Rename datastore
Extension -> Register extension
Extension -> Unregister extension
Extension -> Update extension
Folder -> Delete folder
Folder -> Rename folder
Global -> Disable methods
Global -> Enable methods

Global -> Licenses
Host -> Configuration -> Maintenance
Host -> Configuration -> Storage partition configuration
Host -> Inventory -> Add standalone host
Network -> Assign network
Resource -> Assign virtual machine to resource pool
Resource -> Remove resource pool
Resource -> Rename resource pool
Sessions -> Validate session
vApp -> Delete
vApp -> Rename
vApp -> Add virtual machine
Virtual machine -> Snapshot Management -> Revert to snapshot
Virtual machine -> Configuration *
Virtual machine -> Interaction -> Answer question
Virtual machine -> Interaction -> Power Off
Virtual machine -> Interaction -> Power On
Virtual machine -> Inventory -> Create new
Virtual machine -> Inventory -> Register
Virtual machine -> Inventory -> Remove
Virtual machine -> Inventory -> Unregister
Virtual machine -> Provisioning *
Virtual machine -> Snapshot Management -> Create snapshot
Virtual machine -> Snapshot Management -> Remove snapshot

- **Web service** : Optionally, change the web service entry point URI. Default: /sdk

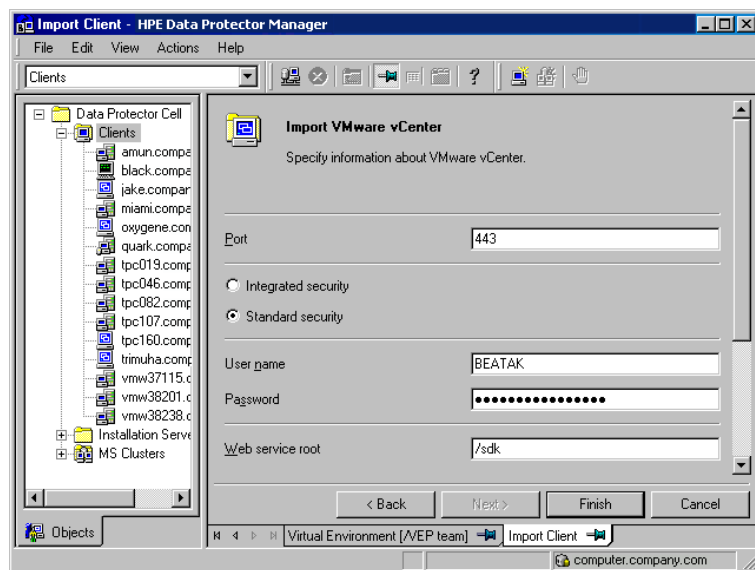
If you select **Integrated security**, which is only available for VMware vCenter Server systems with both the application client and the backup host are Windows systems, Data Protector connects to the VMware vCenter Server system with the user account under which the Data Protector Inet service on the backup system is running. Ensure that this user account has appropriate VMware vSphere rights to connect to the VMware vCenter Server system and the Data Protector Inet service on the backup host is configured for user impersonation.

For details on setting accounts for the Inet service user impersonation, see the *HPE Data Protector Help* index: "Inet user impersonation".

For the **Port** and **Web service root** options, Data Protector uses the values that are currently specified for the standard security. Integrated security is based on Security Support Provider Interface (SSPI).

Click **Finish**.

Importing a VMware vCenter Server client (login credentials)



NOTE:

For details on how to change or check the parameters later, see [Changing the configuration of VMware clients, below](#) and [Checking the configuration of VMware clients, on page 341](#).

Changing the configuration of VMware clients

When you update the credentials for connecting to a VMware client (vCenter Server, ESX(i) Server client), you actually update the `cell_info` file which resides on the Data Protector Cell Manager. Therefore, you can only change the login credentials if you have the `Data ProtectorClients` configuration user right. For details on the Data Protector user rights, see the *HPE Data Protector Help* index: "user groups".

To update the credentials, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can update the credentials in two different places: in the Clients or in the Backup context.

Clients context

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand Clients, and then select the client for which you want to change the login credentials.


3. In the Results Area, click the **Login** tab.
4. Update the credentials and click **Apply**.

Backup context

It is assumed that a backup specification for the VMware client for which you want to change the login credentials already exists.

1. In the Context List, click **Backup**.
2. Open a backup specification for the VMware client for which you want to change the login credentials.
3. In the Source page, right-click the client at the top, and select **Configure**.
4. In the Configure Virtual Environment dialog box, update the values and click **OK**.

Changing the configuration of a VMware vCenter Server or VMware ESX(i) Server client



Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.
2. Execute:
For **Integrated security**:


```
vepa_util.exe  
command  
--config  
--virtual-environment vmware  
--host VMwareClient  
--security-model 1
```

For **Standard security**:

VMware vCenter Server or VMware ESX(i) Server client

```
vepa_util.exe  
command  
--config  
--virtual-environment vmware  
--host VMwareClient  
--security-model 0  
--username Username  
{--password Password | --encoded-password Password}  
[--webroot WebServiceRoot]  
        [--port WebServicePort]
```

The message *RETVAL*0 indicates successful configuration.

For descriptions of the options, see the `vepa_util.exe` man page or the *HPE Data Protector Command Line Interface Reference*.

Checking the configuration of VMware clients

During the configuration check, Data Protector tries to connect to a VMware client using the login credentials from the `cell_info` file on the Data Protector Cell Manager.

To verify the connection, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can verify the connection to a VMware client after you have created at least one backup specification for this client.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Virtual Environment**. Click the backup specification for the VMware client to be checked.
3. In the Source page, right-click the VMware client and select **Check configuration**.

Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.
2. Execute:
VMware vCenter Server or VMware ESX(i) Server client

```
vepa_util.exe  
command  
--check-config  
--virtual-environment vmware  
--host VMwareClient
```

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

For option description, see the `vepa_util.exe` man page or the *HPE Data Protector Command Line Interface Reference*.

Configuring virtual machines

To configure a virtual machine means to specify how the virtual machine should be backed up.

You can specify the following:

- (Windows virtual machines only) Whether a quiescence snapshot should be taken to make applications running inside a virtual machine consistent for backup.
- Which transportation mode should be used during backups.

For each datacenter, you can specify:

- Common settings that apply to all virtual machines in the datacenter.
- Virtual machine-specific settings that override the common settings. If there are no virtual machine-specific settings, the common settings are used for that particular virtual machine.

All these settings are saved in a datacenter-specific file `VMwareClient%DatacenterPath` on the Cell Manager. The file is used for all backup sessions that use any of the backup specifications for this datacenter.

Similarly, backup sessions using any of the backup specifications for `All` datacenters, use the settings from the file `VMwareClient%ALLDatacenters`.

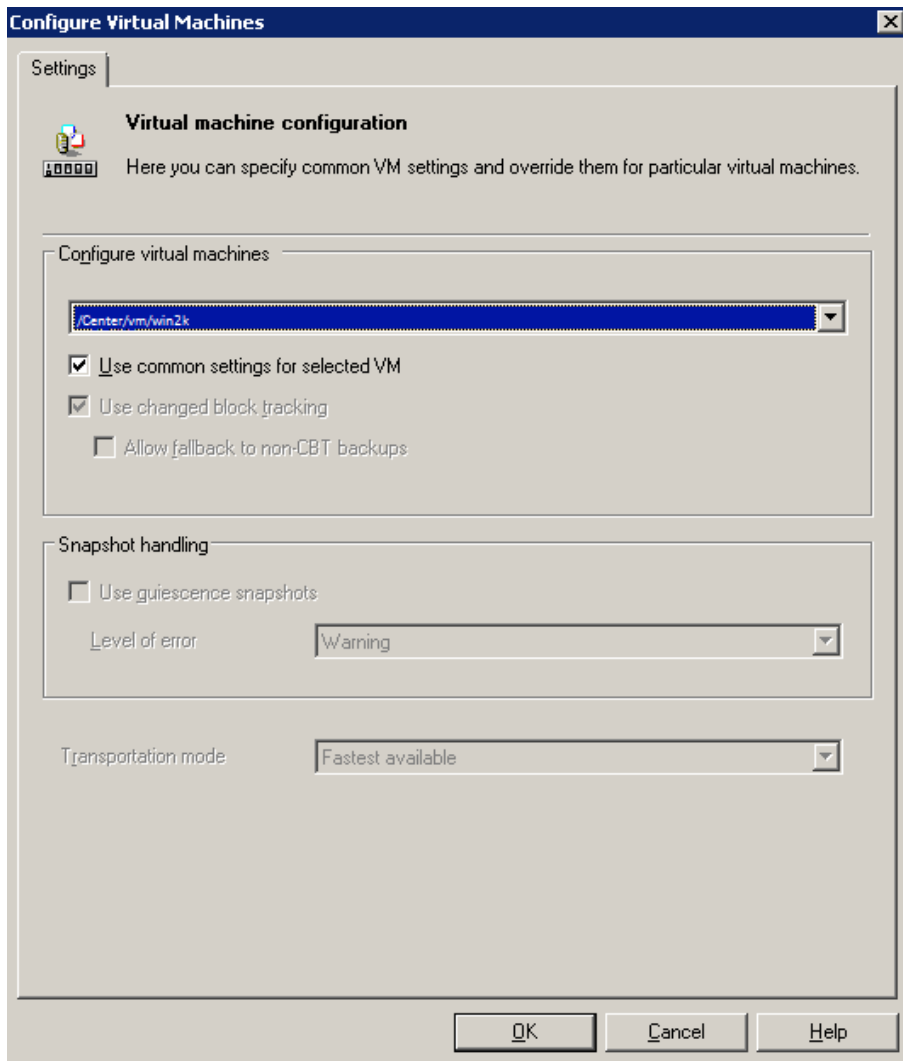
The files `VMwareClient%DatacenterPath` and `VMwareClient%ALLDatacenters` are created or updated when you create or update a backup specification for a specific datacenter or all datacenters respectively.

To configure virtual machines, use the Data Protector GUI or CLI.

Using the Data Protector GUI

You can configure virtual machines when you create or modify a backup specification. In the Source page of a backup specification, right-click the client system at the top or any of the virtual machines listed below, and select **Configure Virtual Machines**.

Configure virtual machines (settings for VMware vCenter Server client)



In the Configure Virtual Machines dialog box **Settings** page, specify the following settings:

Virtual machine settings

Options Available	Description / Action
Select whether you want to specify common virtual machine settings (Common VM Settings) or settings for a specific virtual machine. Virtual machine specific settings override the common virtual machine settings.	
Configure virtual machines	
Use common settings for selected VM	Available only if a virtual machine is selected. Select this option if you want the common settings to apply to the selected virtual machine.

Options Available	Description / Action
	Default: selected
Use default settings	Available only if Common VM Settings is selected. Select this option to set default values for the common virtual machine settings. Default: selected
Enable changed block tracking	The VMware changed block tracking functionality is enabled for the selected virtual machines. Default: selected and greyed out
Allow fallback to non-CBT backups	Enabled only if Use changed block tracking is selected. Select this option to continue backup in a non-CBT mode for successful backup of Data Protector. For more information on Non-CBT, refer to Non-Changed Block Tracking (Non-CBT) backup, on page 321 topic. Default: not selected.
Snapshot handling	
Use quiescence snapshots	Applicable for Windows virtual machines. Not available if Use default settings or Use common settings for selected VM are selected. Select this option to use Microsoft Volume Shadow Copy Service (VSS) functionality to quiesce all applications with VSS writers before performing the VEPA backup. This produces application-consistent backups. Default: not selected. For details, see Quiescence .
Level of error	Available only if Use quiescence snapshots is selected. Specify the level of error to be reported if a quiescence snapshot fails. Default: Warning. For details, see Quiescence .
Transportation mode	
	Select the transportation mode to be used when backing up a virtual machine. <ul style="list-style-type: none"> • NBD : Use this mode when your ESX(i) Server systems do not have access to a SAN, but use local storage devices or NAS to store virtual machine disks. This is an unencrypted transportation mode over a

Options Available	Description / Action
	<p>local area network that uses the Network Block Device (NBD) driver protocol. This transportation mode is usually slower than Fibre Channel.</p> <ul style="list-style-type: none"> • NBD (SSL) : Same as NBD except that the communication over the network is encrypted using the Secure Socket Layer (SSL) cryptographic protocol. • Hotadd : Use this mode if your backup host (a client with the Data ProtectorVirtual Environment Integration component installed) is a virtual machine. Such a configuration enables you to back up other virtual machines residing on datastores visible to the ESX(i) Server that hosts the backup host. • SAN : Use this mode when your ESX(i) Server systems store their virtual machine disks on Fibre Channel SAN or iSCSI SAN. This is an unencrypted transportation mode over Fibre Channel or iSCSI. <ul style="list-style-type: none"> ◦ Caution: Do not reformat these storage volumes. Otherwise, you will delete all your virtual machines. <p>For details on these VMware transportation modes, see the VMware documentation.</p> <p>If you are not concerned which mode is used, select Fastest available.</p> <p>Default: Fastest available</p>

Using the Data Protector CLI

1. Log in to the backup host, open the command prompt and change to the directory in which the `vepa_util.exe` command is located.
2. Execute:

```
vepa_util.exe
    command
    --configvm
    --virtual-environment { vmware | vCD }
    --host AppHostName
    --instance DatacenterPath
    --vm VMpathVM_OPTIONSVM_OPTIONS
    --transportation-mode {san | nbd | nbdssl | hotadd | fastest}
    --quiescence { 0 | 1 }
    --quiescenceErrLvl { 0 | 1 }
    --uuid UUID_of_VM
```

For details, see the `vepa_util.exe` man page or the *HPE Data Protector Command Line Interface Reference*.

To change the virtual machine specific settings back to the common virtual machine settings, execute:

```
vepa_util.exe
    command
    --configvm
    --virtual-environment { vmware | vCD }
    --host AppHostName
    --instance DatacenterPath
    --vm VMpath
    --uuid UUID_of_VM
    --default
```

The message *RETVAl*0 indicates successful configuration.

Example

To configure the virtual machine with the virtual machine path /MyDatacenter/MyVM and UUID 42375365-ebe1-e9da-7068-7beb727cab19 that resides in the datacenter /MyDatacenter and registered in the vCenter Server system vc.company.com, with the fastest available transportation mode.

execute:

```
vepa_util.exe
    command
    --configvm
    --virtual-environment vmware
    --host vc.company.com
    --instance /MyDatacenter
    --vm /MyDatacenter/MyVM
    --quiescence 1
    --quiescenceErrLvl 0
    --transportation-mode fastest
    --uuid 42375365-ebe1-e9da-7068-7beb727cab19
```

Customizing the Data Protector behavior with omnirc options

The omnirc options are useful for troubleshooting or overriding other settings affecting the behavior of the Data Protector client. The options that apply to the Virtual Environment ZDB integration for VMware have the prefix OB2_VEAGENT.

For details on how to use Data Protector omnirc options, see the *HPE Data Protector Help* index: “omnirc options”.

Adding the RMC Server details in Data Protector using the Command Line Interface

You can add the RMC server details using the omnidb command. Execute:

```
omnidb -addhost -servername <clientname> -user <username> -psswd <password>
```

You can also use the following commands to list and remove the required servers. Execute:

```
omnidb -listhost -servername <clientname>
```

```
omnidb -removehost -servername <clientname> -user <username>
```

Backup

This section contains procedures that are needed to perform zero downtime backup of virtual machines. For details on backup concepts, see [Backup concepts, on page 315](#).

With the Virtual Environment ZDB integration for VMware, you can perform the zero downtime backups for:

- **NetApp storage:** ZDB to tape
- **3PAR storage:** ZDB to disk, ZDB to tape, and ZDB to disk+tape

For more information on the ZDB types, see the *HPE Data Protector Concepts Guide*.

You should be familiar with the NetApp and HPE 3PAR storage concepts and procedures and basic Data Protector ZDB functionality. See the storage-related documentation and the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

For information about the supported configurations, ZDB types and replication techniques available on this storage system, and storage system-specific ZDB considerations, see the *HPE Data Protector Concepts Guide*.

Backup limitations

- Before performing backups, ensure that you use only the supported characters in the names of any VMware vCenter and VMware ESX(i) objects (for example, virtual machines, datastores, datacenters, vApps, and so on), as special characters are not supported.

The following list includes the supported characters:

- Letters a-z without any special characters
- Numbers 0-9
- Single Quotes (')
- Spaces
- Underscores (_)
- Hyphens (-)

A **partial list** of special characters that are **not supported** is as follows:

- % (percentage)
- + (plus)
- = (equals)

- @ (at)
- < (less than)
- > (greater than)
- : (colon)
- " (double quote)
- / (forward slash)
- \ (backslash)
- | (vertical bar or pipe)
- ? (question mark)
- * (asterisk)

Special characters are not supported due to a known VMware issue described at the following URLs:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd;=displayKC&externalId;=2017661

<http://www.vmware.com/support/developer/vddk/VDDK-1.2.1-RelnoteOPs.html>

- After you upgrade, you cannot restart the failed backup object versions from earlier Data Protector versions below 8.11.
- After you upgrade from Data Protector 7.03 or earlier versions, you cannot run incremental or differential backups without running a full backup.
- After an upgrade to Data Protector 9.05, the SAN transport mode falls back to NBDSSL on the vSphere versions 5.1 and 5.5. This is a VMware VDDK 6.0 Update1 issue. For more details, see the following URL: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2135621.
- For RMC integrations, backups that have virtual machines that are presented through iSCSI to the ESXi server is not supported.
- In RMC integrations, for **Snapshot + Tape** backups, the name of the configured mount ESXi host in the 3PAR array must be the same as its hostname. Also, the configured name of the mount ESXi in the 3PAR array must be in lower case.
- For datastores not located on VMFS volume, the full backups take full size of the disk and not only the changed blocks. For more details, see the following URL:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1020128
- Consider all limitations that apply to the Data Protector NetApp Storage and HPE 3PAR Storage integration. The following limitations apply to Virtual Environment ZDB integration for VMware:
 - ZDB backup is supported on NetApp storage and HPE 3PAR storage system. Multiple storage providers are not supported in a single backup session.

- You cannot perform ZDB backup on virtual machines or virtual disks that do not reside on a NetApp storage or 3PAR storage.
- A template backup in the SAN transport mode is not supported, and will fall back to the NDB transport mode.
- If you want to back up a single disk from a virtual machine that contains multiple disks from different data stores, you must select the disk that belongs to the data store containing the configuration files.
- Only one snapshot type for target volumes can be created during a ZDB session.
- When cloning process for a source volume is in progress, another snapshot (any type) of that source volume cannot be created.
- You cannot back up replicas (target volumes from existing and currently recorded backup sessions).
- If there is not enough space for a fully allocated replica creation, the session fails.
- Do not enable the vCenter option **disable-datastore-web**. This may result in unsuccessful backups of vmdk and vmx files. For more information, see [vSphere Hardening Guides](#).

vStorage Image + OpenStack backup method limitations

- ZDB backup and restore are not supported.
- Non-CBT backups are not supported.

Creating backup specifications

Create a backup specification using the Data Protector GUI. The ZDB backup steps for NetApp storage and HPE 3PAR storage systems are similar.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Virtual Environment**, and select **Add Backup**.
3. In the Create New Backup dialog box, select **Snapshot or split mirror backup** as Backup type and **Storage Provider Plugin** as Sub type. For description of options, press **F1**. Click **OK**.
4. Specify the application to be backed up:
 - In the Client drop-down list, select a VMware client.

NOTE:

The drop-down list contains all clients that have been imported into the Data Protector cell as VMware vCenter or VMware ESX(i) clients. They have a corresponding label appended at the end of their names, such as **(VMware vCenter)** or **(VMware ESX(i))**.

If the selected VMware client is not configured correctly, a warning is displayed. Click **OK** to open the Configure Virtual Environment dialog box and provide the connection parameters as described in [Importing and configuring VMware clients, on page 336](#).

- In the Backup host drop-down list, select a VMware vCenter system to be used to control the backup. The list contains all clients that have the Data Protector related ZDB integration component (NetApp Storage Provider or 3PAR Storage Provider) and Virtual Environment Integration component installed.
- In Datacenter/Organization, select a datacenter to back up from.

NOTE:

If you have selected a standalone ESX(i) Server system in the **Client** option, there is only one datacenter available – **/ha-datacenter**.

If you have selected a vCenter Server system in the **Client** option, you can select **All Datacenters** to back up virtual machines from different datacenters.

- In Mount host, select the ESX(i) Server system to be used for mounting the replicas.

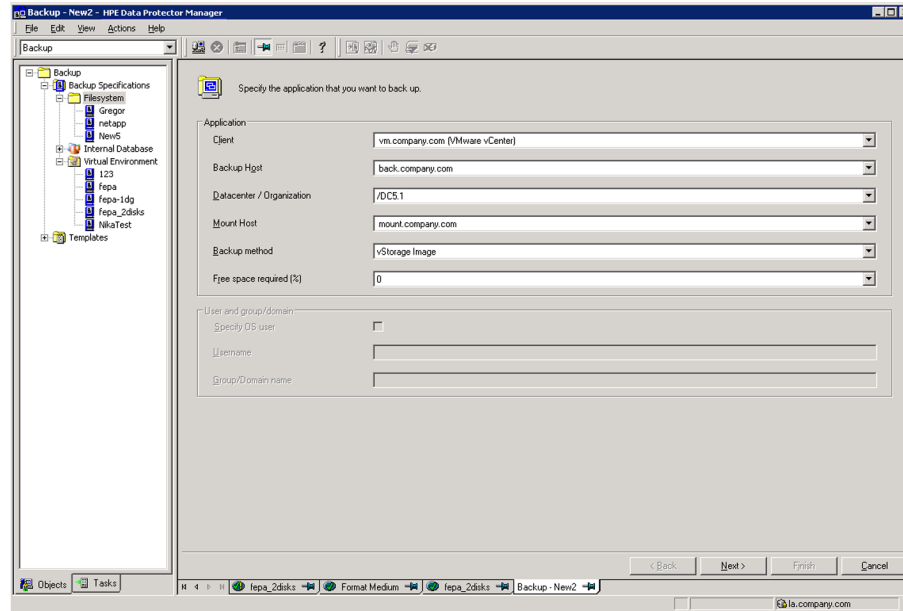
NOTE:

If the mount proxy host with ESXi version 5.5 U2 belongs to the source data center, Remote Copy failover backup will fail. Choose the ESXi mount proxy host from a different data center that belongs to the same vCenter.

- In Backup method, the backup method is displayed:
 - vStorage Image for VMware vCenter and VMware ESX(i) clients
 - vStorage Image + OpenStack for VMware vCenter
- In Free space required [%], specify the percentage of disk space that should be free on a datastore before a virtual machine is backed up. The free space is calculated based on the size of the datastore where the virtual machine disks reside.

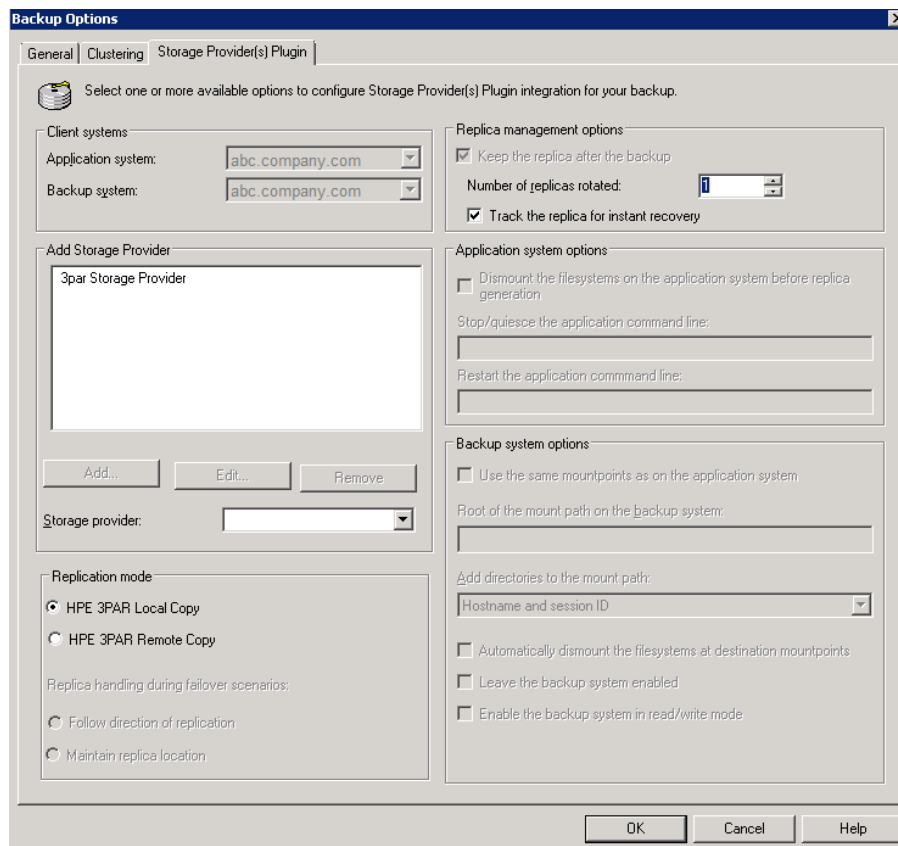
The check is performed separately for each virtual machine. For details, see [Free space required option, on page 324](#).

Selecting a VMware vCenter Server client, backup host, mount host, and datacenter



Click **Next**.

Selecting a Replication mode



Click **OK**.

5. Under Add Storage Provider, select the desired storage provider from the Storage Provider drop-down list and then click **Add**. The storage provider specific options dialog opens. Select the replica provision type, enter a replica description for NetApp Storage system residing in cluster, enter the destination array and Vserver names. Specify other storage provider specific options, if available:
 - Replica provision type: **Thin provisioned** or **Fully allocated** replica provision type are available for NetApp and 3PAR storage providers.
 - Transportation mode: **SAN** and **NBD** transportation modes are available for NetApp storage provider.

Click **OK**. The storage provider is added to a list. You can later change its options by clicking **Edit** or remove it from the list by clicking **Remove**. For more information, press **F1**. Click **Next**.

6. Under Replication mode, the **HPE 3PAR Local Copy** is selected by default. You can select the **HPE 3PAR Remote Copy** as well. For more information, see the *HPE Zero Downtime Backup Administrator's Guide*.
7. For 3PAR storage systems, select the **Keep the replica after the backup** check box. Choose the required value for **Number of replicas** rotated. Optionally, select the **Track the replica for instant recovery** check box.
8. Select the objects that you want to back up. In the Show drop-down list, simplify your selection by choosing the **Hosts and Clusters** or **VMs and Templates** view. By default, Hosts and Clusters is displayed.

NOTE:

If you switch the view after you have already selected one or more objects for backup, a warning dialog is displayed. Its confirmation clears the already selected objects, clicking **No** does no change to the view.

You can make your selections at different levels:

- For VMware vCenter and VMware ESX(i) clients:
 - ESX/ESXi Servers systems
 - Pools
 - vApps
 - VM folders
 - Individual VMs
 - VM disks
 - VM templates


If you select any level above individual VMs (for example, a vApp), all VMs and VM disks contained in the selected item will be included in the backup specification. If VMs are added within the item after the backup specification is saved, they will also be backed up.

NOTE:

If you clear the check box against an object, that object will be excluded from the backup specification. Thereafter, if you add a new virtual machine, a pool, or a vApp to an existing logical object, it is automatically included in the backup. You are not required to create a

new backup specification. The check box for excluded objects is marked with a red cross symbol in the selected view.

Selecting VMware objects (vCenter Server client)



NOTE:
During the **vStorage Image + Openstack** backup specification creation, the Shadow VMs are not available for selection during backup and restore operations.

IMPORTANT:
In the object tree of a particular VMware client, a virtual machine may be displayed as selected in two different ways:

- The *blue* check mark indicates that the virtual machine is selected for backup in its entirety, including its configuration and all its virtual disks.
If such a virtual machine is backed up, it can be restored even if the original virtual machine does not exist anymore.
- The *gray or black* check mark indicates that some or all virtual disks belonging to the virtual machine are selected. The virtual machine itself and its configuration is omitted from the backup.
If such a virtual machine is backed up, its disks can only be restored if the original virtual machine is still configured at the time of restore.

If your virtual machines are not configured yet, right-click the client system at the top or any of the virtual machines listed below, and select **Configure Virtual Machines**. For details, see [Configuring virtual machines, on page 342](#).

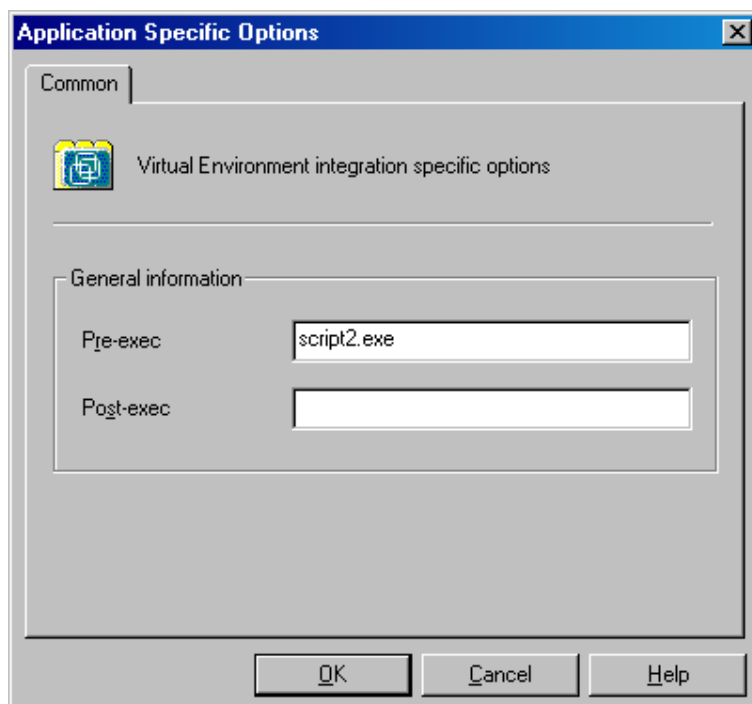
Click **Next**.

9. Select which devices to use for the backup.
To specify device options, right-click the device and select **Properties**. Specify the number of parallel backup streams in the Concurrency tab and which media pool to use.

Click **Next**.

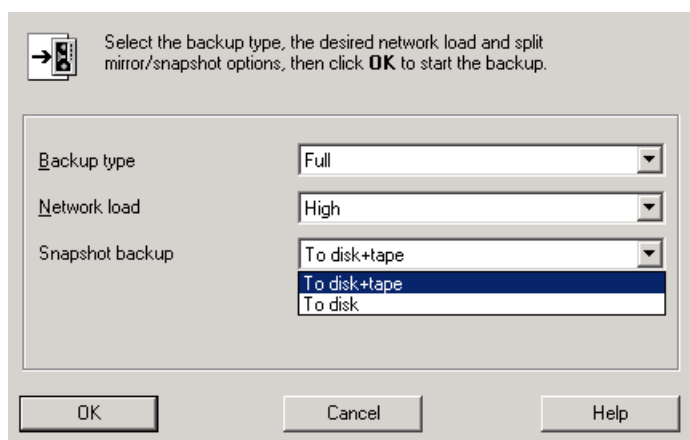
10. Set backup options.
11. For information on application-specific backup options, see [VMware backup options](#) , on the next page.

Application-specific options



Click **Next**.

12. Click **Save As** to save the backup specification, specify a name and a backup specification group. Optionally, you can click **Save and Schedule** to save, and schedule the backup specification. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.
13. Click **Start Backup** to start the backup session. If you are using the 3PAR storage system, you will have the option of selecting the **Backup type**, **Network load**, and the **Snapshot backup** options.



TIP:
Preview your backup specification before using it for a real backup. See [Previewing backup sessions, on page 359](#).

VMware backup options

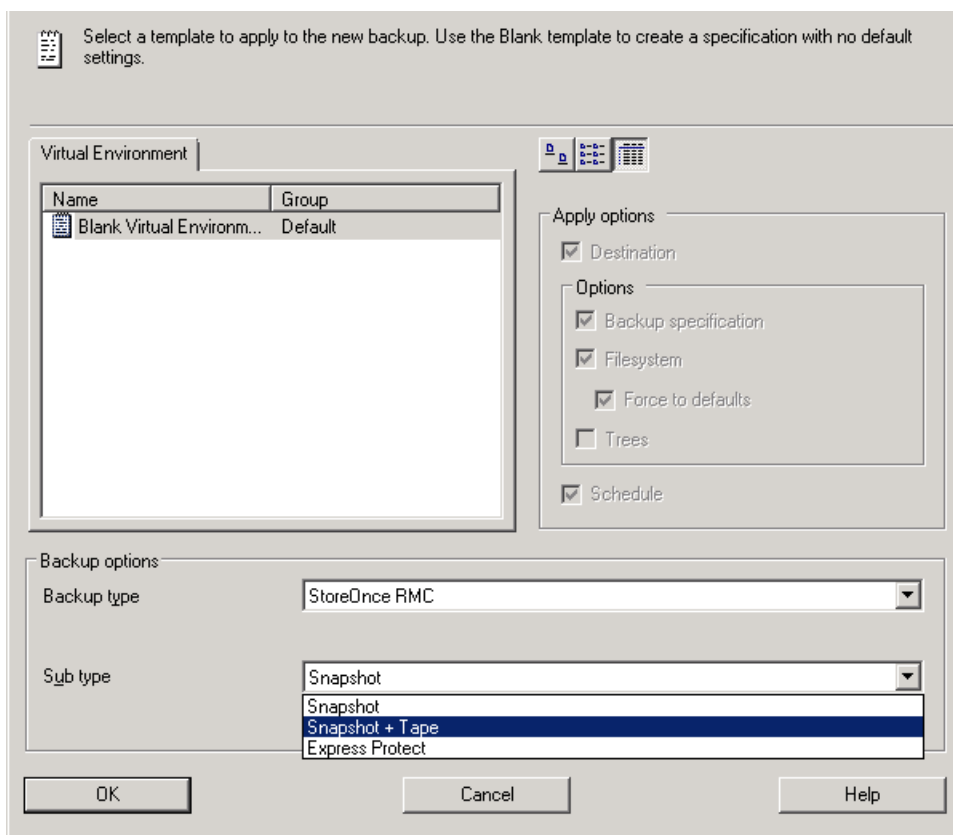
Option	Description
Pre-exec , Post-exec	<p>Specifies which command line to run on the backup host before (<i>pre-exec</i>) or after (<i>post-exec</i>) the backup.</p> <p>Do not use double quotes. Type only the name of the command and ensure that the command resides in the default Data Protector administrative commands directory on the backup host.</p> <p><i>Windows systems:</i> Data_Protector_home\bin</p> <p><i>Linux systems:</i> /opt/omni/bin</p>

Creating backup specifications for RMC backups

Consider all the prerequisites and limitations listed in the [StoreOnce Recovery Manager Central Integration](#) section.

Create a backup specification using the Data Protector GUI.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Virtual Environment**, and select **Add Backup**.
3. In the Create New Backup dialog box, select **StoreOnce RMC** as Backup type. Select **Snapshot**, **Snapshot+Tape**, or **Express Protect** as as Sub type, based on your requirements. Click **OK**.
 - **Snapshot:** If you select this, replicas of the source volumes are created.
 - **Snapshot+Tape:** If you select this, replicas of the source volumes are created, and the data is backed up to the chosen tape device. Ensure that you specify the **Mount Host**, and the required **Devices** in the subsequent steps.
 - **Express Protect:** If you select this, you can backup snapshots from HPE 3PAR StoreServ to HPE StoreOnce. The backup proceeds as per the backup policy created in RMC.




4. Specify the application to be backed up:
 - In the Client drop-down list, select a VMware client. If the selected VMware client is not configured correctly, a warning is displayed. Click **OK** to open the Configure Virtual Environment dialog box and provide the connection parameters as described in [Importing and configuring VMware clients, on page 336](#).
 - In the Backup host drop-down list, select a VMware vCenter system to be used to control the backup.
 - In Datacenter/Organization, select a datacenter to back up from.
 - In Mount host, select the ESX(i) Server system to be used for mounting the replicas. This field is mandatory for **Snapshot+Tape** backups.
 - In Backup method, select the method to be used for the backup.

NOTE:
RMC does not support the **vStorage Image + OpenStack** backup method.

5. Specify the RMC configuration details:

- **RMC Server:** Select the RMC server for backup.
- **Backup Policy** (*available only for Express Protect backups*): Select the policy name that you have created in RMC. The backup policy usually contains the backup system and the backup store. For more information, see the *HPE StoreOnce RMC User Guide*.
- **Snapshot Count:** Specify the maximum number of snapshot count that should be maintained in the storage array. A maximum of 1000 and a minimum of 1 can be created. The default value is 10.
- **Express Protect Count** (*available only for Express Protect backups*): Specify the maximum number of backups that should be maintained in StoreOnce for the backup specification. The minimum value that can be set is 2. The default value is 10.



6. Select the objects that you want to back up. In the Show drop-down list, simplify your selection by choosing the **Hosts and Clusters** or **VMs and Templates** view. By default, Hosts and Clusters is displayed.
7. Select the devices that you want to use for the backup. If you have specified the mount host, then the device is selected by default.

Click **Next**.

NOTE: This step is not applicable for **Snapshot** and **Express Protect** backups.

8. Set backup options. Click **Next**.
9. Click **Save As** to save the backup specification, specify a name and a backup specification group. Optionally, you can click **Save and Schedule** to save, and then schedule the backup

specification. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

10. Click **Start Backup** to start the backup session.

Modifying backup specifications

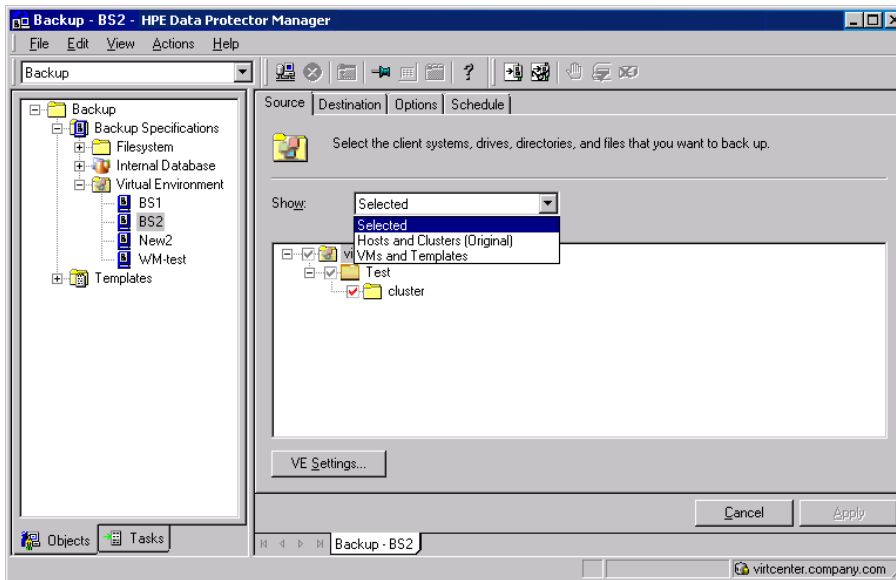
To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

In the Source page, you can modify the backup objects by using the Show drop-down list. In the drop-down list, the **Hosts and Clusters** view or the **VMs and Templates** view are available. The view that you have used during the creation of your backup specification has the string (Original) appended. If you switch the view, a warning dialog is displayed and its confirmation clears the already selected objects.

NOTE:

When modifying a backup specification created with one of the previous versions of Data Protector, you can switch between **Selected**, **All**, **Hosts and Clusters**, and **VMs and Templates** views. The view **All** is available for the legacy backup specifications only and provides a legacy browsing mechanism. Selecting **Hosts and Clusters** or **VMs and Templates** and clicking **Yes** in the warning dialog clears the previous backup object selection and upgrades your browsing mechanism. After saving the backup specification only the **Hosts and Clusters** and **VMs and Templates** views are available.

Modifying a backup specification (VMware vCenter Server client)



To display the virtual environment settings, in the Results Area, click **VE Settings**. Not all settings can be modified.

Scheduling backup sessions

You can run unattended backups at specific times or periodically. For more information on how to create and edit schedules, see Scheduler in Data Protector in *HPE Data Protector Administrator's Guide*.

Previewing backup sessions

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

Using the Data Protector GUI

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

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

Using the Data Protector CLI

1. Log in to any client with the `Data ProtectorUser Interface` component installed.
2. Open the command prompt and change to the directory in which the `omnib` command is located.
3. Execute:

```
omnib -veagent_list BackupSpecificationName -test_bar
```

What happens during the preview?

The following are tested:

- Communication between the backup host and Data Protector
- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

Starting backup sessions

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

To start a backup interactively, use the Data Protector GUI or CLI.

Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **Virtual Environment**. Right-click the backup specification you want to use and select **Start Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

Using the Data Protector CLI

1. Log in to any client with the `Data ProtectorUser Interface` component installed.
2. Open the command prompt and change to the directory in which the `omnib` command is located.
3. Execute:

```
omnib -veagent_list BackupSpecificationName [-barmode VirtualEnvironmentMode]
[ListOptions]
```

where *VirtualEnvironmentMode* is one of the following backup types:

```
full|diff|incr
```

The default is `full`.

For *ListOptions*, see the `omnib` man page or the *HPE Data Protector Command Line Interface Reference*.

Examples

To start a full backup using the backup specification `MyVirtualMachines`, execute:

```
omnib -veagent_list MyVirtualMachines -barmode full
```

To start a differential backup using the same backup specification, execute:

```
omnib -veagent_list MyVirtualMachines -barmode diff
```

Preparing for disaster recovery

To do a disaster recovery, you need backups of the following VMware objects:

What must be backed up

VMware object	How to back it up
ESX/ESXi Server console	<p>ESX Server systems:</p> <ol style="list-style-type: none"> 1. Ensure that the <code>Data ProtectorDisk Agent</code> component is installed on the backup system. 2. In the Backup context of the Data Protector GUI, right-click Filesystem and select Add backup to create a backup specification of the filesystem type. In the Source page of the backup specification, select ESX Server consoles of all ESX

VMware object	How to back it up
	<p>Server systems.</p> <p>For details on what to back up, see the topic “ESX Server Configuration Backup and Restore procedure” at http://kb.vmware.com/selfservice/microsites/microsite.do.</p> <p>3. Start a backup using the newly created backup specification.</p> <p>ESXi Server systems:</p> <p>With ESXi Server systems, it is not possible to install a Data ProtectorDisk Agent, so you will need to use VMware utilities to back up the configuration.</p> <p>A tool, <code>esxcfg-cfgbackup</code> is available from VMware. For information, see the VMware website.</p>
<p>vCenter configuration database</p> <p><i>(applicable only for VirtualCenter environments)</i></p>	<p>The vCenter configuration database can be an Oracle database or Microsoft SQL Server database. To back up the database, use the corresponding Data Protector integration. For example, if it is an Oracle database, follow these steps:</p> <ol style="list-style-type: none"> 1. Ensure that the Data ProtectorOracle Integration component is installed on the vCenter Server system. 2. In the Backup context of the Data Protector GUI, right-click Oracle Server and select Add backup to create a backup specification of the Oracle type. In Application database, type the name of the vCenter configuration database. <p>Continue with the backup specification creation as described in the <i>HPE Data Protector Integration Guide</i>.</p> <p>3. Start a backup using the newly created backup specification.</p>
<p>VMware virtual machines</p>	<p>Back up the virtual machines as described in this chapter.</p>

Restore

This section contains procedures that are needed to restore virtual machines. For details on restore concepts, see [Restore concepts, on page 327](#).

Restore limitations

- When restoring a VM to a datacenter with a datastore shared among several inventory objects (hosts or clusters) or several cluster nodes, and the restore option for its subsequent registration is selected, the restored VM may not register at the original inventory location:

- If several inventory objects share the datastore, the VM is registered with the first available host or cluster.
- If several cluster nodes share the datastore, the VM is registered with the first available cluster node.

If the VM must be registered at its original inventory location, migrate it appropriately after the restore session completes.

- Before performing a restore, ensure that you use only the supported characters in the names of any VMware vCenter, or VMware ESX(i), objects (for example, virtual machines, datastores, datacenters, vApps, and so on), as special characters are not supported. For more details on the supported characters, see Backup Limitations.
- vDatacenters that are disabled do not allow adding or removing of vApps. Consequently restore to a disabled vDatacenter is not supported.
- In a vSphere environment, do not use the left parenthesis (the (character) in names of vSphere distributed port groups. Data Protector is unable to connect a restored virtual machine to a non-original vSphere network that uses distributed switches and is assigned such a distributed port group. In this case, after the restore, you need to manually connect the restored virtual machine to the desired network using vSphere Client.
- Recovering virtual machines after a restore to a directory is not supported for incremental or differential backups. The VMware Converter that is used for moving the restored VM images to an ESX Server or ESXi Server system recognizes only the full backup type.
- Recovering virtual machines from a backup chain (for example, full, incremental, incremental, incremental...) is supported only for a restore to a datacenter.
- All the media related to a session should be exported or imported together. If a media from the full backup object is missing then the Incremental/Differential backup will not detect a missing media and continue to run in the selected mode. A restore from such a session will not be successful.
- A virtual machine can be seen as different objects in the Internal Database (IDB) and in the restore context based on how the backup specification was created.
- It is not possible to restore a suspended Virtual Machine.
- In earlier versions of Data Protector, all disks are read even if you select only one disk for restore. From Data Protector 8.11 onwards, only the selected disks are read and restored. This is applicable for backups created with Data Protector 8.11 and not with its earlier versions.
- Movement from one disk to another on a tape is not possible. If you select consecutive disks (scsi0:1 and scsi0:2) for restore, then both the disks are read and restored. If you select scsi0:1 and scsi0:4 for restore then the following combinations occur:
 - scsi0:0 is read.
 - scsi0:1 is read and restored.
 - scsi0:2 is read.
 - scsi0:3 is read.
 - scsi0:4 is read and restored.

So, random selection of disks could result in higher restore window when compared to single or consecutive disk selection.

- During a VEPA (Virtualization Environment Agent) restore process, the Virtual Environment Integration Agent makes a connection to the vCenter / ESX to upload the virtual machine configuration files to VC / ESX. Such a connection could continue to exist in vCenter even after the restore session is complete. This idle session is cleared by VMware after a default time-out value of 30 minutes. Alternatively this session can be cleared manually from the vCenter.

vStorage Image + OpenStack restore limitations

- Restore from Replica is not supported.
- Features, such as, Restore-as, Delete-after, and Forensics are removed from the Data Protector Restore context.
- Granular Recovery Extension is not supported for Nova Instance and Shadow VM.
- If the Instances are deleted from the OpenStack dashboard, Recovery of instance from the Data Protector to Dashboard is not possible.
- If the Shadow VM is attached to any other instance other than the original instance to which it was attached at the time of backup, then the restore fails.
- Restore to a different cluster is not supported.

NOTE:

The Power On and Live Migrate functionalities are not supported for **vStorage Image + OpenStack** method.

Power On and Live Migrate limitations

- The Power On and Live Migrate options are not supported for Smart Cache device that are created on CIFS and NFS shares.
- Power On and Live Migrate operations from StoreOnce catalyst are supported on Windows and Linux backup hosts.
- For virtual machines that are created on Smart Cache and StoreOnce Catalyst devices, only eight virtual machines can be powered on at once. If you want to increase this number, see [Increasing the default value that defines the maximum number of NFS mounts on an ESXi/ESX host](#).
- If you have backed up a virtual machine using the **Disk Only** option, restore to an ESX server is not supported.
- If you have backed up a virtual machine using the **Disk + Tape** option, restore to an ESX server is supported from the tape.
- If you Power On a virtual machine from Smart Cache and StoreOnce Catalyst devices, the virtual machine will be powered on without the virtual RDM disks.
- If you perform a Live Migrate procedure from Smart Cache and StoreOnce Catalyst devices, the virtual RDM disks will be migrated as thick disks, i.e. after migration they will reside as eager zeroed thick provisioned vmdk disks on the target datastore. If you want to restore virtual machines that have virtual RDM disks, use the normal Restore procedure.
- Power On and Live Migrate of virtual machines to ESX servers outside the vCenter, is not

supported.

- Live Migrate of backups that are performed with an ESX server as the backup client, is supported only through the CLI.
- You can perform only one operation (Restore / GRE / Power On / Live Migrate) at a time on the virtual machines that have backups on StoreOnce Catalyst device.
- If the virtual machine has ongoing restore operations (Power On, Live Migrate, or GRE) from the StoreOnce Catalyst device, clean up the virtual machine before performing any of the restore operations from the same backup session. This is also applicable for the backup sessions that are part of the restore chain.
- StoreOnce Catalyst devices does not support Power On and Live Migrate options for Non-CBT backup sessions.
- Power On and Live Migrate operations from the StoreOnce Catalyst devices are not supported if the selected backup host is RHEL 6.6.
- Power On and Live Migrate operations are not supported if the backup to StoreOnce Catalyst device is performed using software compression or AES encryption.
- If Power On and Live Migrate of a virtual machine is performed by selecting vApp as destination, then the virtual machine will be powered on and live migrated outside the vApp. You can manually move it inside the vApp if required.
- If the mount proxy host is rebooted with any active Power On or Live Migrate request, then this request will be inaccessible for 4 hours after the reboot.
- If the data backed up with Data Protector 9.04 or earlier versions is transferred to a StoreOnce Catalyst through object copy in 9.07 or later version, then Live Migrate and Power On operations are not supported.
- Full, differential, and incremental backups on different devices are not supported for Power On and Live Migrate operations.

For example:

Full backups of the same virtual machine are done to the StoreOnce Catalyst and Smart Cache devices respectively.

Incremental backup is performed to the StoreOnce Catalyst device. This incremental backup is not eligible for Power On and Live Migrate operations because the full backup preference for this incremental chain resides on Smart Cache.

Finding information for restore

You can find information about backup objects in the Data Protector IDB, such as which backup type and media were used, and which messages were displayed during the backup. To retrieve this information, use the Data Protector GUI or CLI.

Using the Data Protector GUI

In the Internal Database context, expand **Objects** or **Sessions**.

If you expand **Objects**, backup objects are sorted according to the virtual machine for which they were created. For example:

- In a vCenter environment, the backup objects for the virtual machine /vm/mach1 are listed under /4/vCenterName%2FvmInstanceUUID.

Where,


vCenterName is the name of the virtual center.

vmInstanceUUID is a unique identifier for virtual machine /vm/mach1 on vCenter.

To view sessions in the vCenter environment, double-click the

/object1/vCenter%2FvmInstanceUUID.

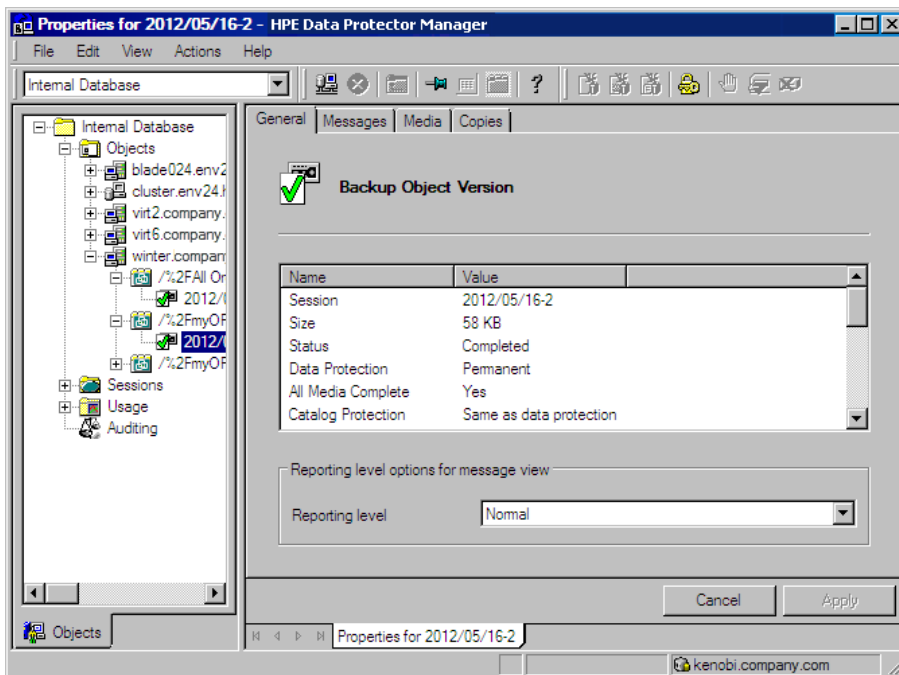
VEAgent Object details



If you expand **Sessions**, backup objects are sorted according to the session in which they were created. For example, backup objects created in the session 2012/07/10-82 are listed under 2012/07/10-82.

To view details on a backup object, right-click the backup object and select **Properties**.

Backup object information



TIP:
To view the messages displayed during the session, click the **Messages** tab.

Using the Data Protector CLI

1. Log in to any client with the Data ProtectorUser Interface component installed.
2. Open the command prompt and change to the directory in which the omnidb command is located.
3. Get a list of VMware backup objects created in a backup session with the session ID *SessionID*:
`omnidb -session SessionID`
4. Get details on a backup object with the backup object name *BackupObjectName*:
`omnidb -veagent BackupObjectName -session SessionID -catalog`

Here is one example of a backup object name:

`gabriel.company.com::/%2FE1Datacentro/0/%2Fvm%2Fharbour`

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

Restoring using the Data Protector GUI

Use this procedure to Restore, Power On and Live Migrate virtual machines.

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Virtual Environment**, expand the relevant client and click the datacenter from which you backed up.
3. In the Source page, specify the following:

- a. From the **Backup method** drop-down list, select any of the following backup methods:
 - vStorage Image for VMware vCenter and VMware ESX(i) clients
 - vCD vStorage Image for VMware vCloud Director clients
 - vStorage Image + OpenStack for VMware vCenter in OpenStack environment
- b. From the **From** and **To** drop-down lists, you can narrow the scope of displayed virtual machines to those backed up within the specified time interval.
- c. In the VM Filter text box, enter filter text for the VM and press **Enter** , or click **Apply Filter** . The filter hides the VMs, vApps, and resource pools that do not match the filter pattern allowing you to find the desired object easily.

After you select the VMware objects, you can choose to **Restore**, **Power On** or **Live Migrate** them. Choose the required option from the **VM Options** drop down. Note that the Power On and Live Migrate options are available only when one object is selected.

NOTE:

Filters are case-sensitive and apply to VMware Virtual Machine objects, VMware Virtual Application (vApp) objects, and resource pools. If a matching sub-node is found (such as, another VM, vApp, or resource pool), they are displayed. If you leave the VM Filter text box empty, all the VMs, vApps, and resource pools are displayed. However, if you enter filter text, only the matching sub-nodes (if any) are displayed. If you modify the filter values using the From, or To drop-down lists, filtering is re-applied. Filter does not apply to already selected VMs, vApps, or resource pools. This means that you can filter machines using one filter, select objects, and again change the filter. In the new filter, the previously marked objects remain visible.


The following are the types of filter usage:

- **Simple substring usage** - If you enter part of a VM, vApp, or resource pool name, all VMs, vApps, or resource pools in the object tree that have the entered string in their name remain visible. All other objects are filtered out.
- **Wildcards and question mark usage** - The following are the filtering options:
 - `<filter string>*` - Filters VM, vApp, or resource pool names that start with the `<filter string>` and end with any set of characters.
 - `*<filter string>*` - Filters VM, vApp, or resource pool names that start and end with any set of characters, and have the `<filter string>` in between.
 - `*<filter string>` - Filters VM, vApp, or resource pool names that start with any set of characters and end with the `<filter string>`.
 - `<filter string>*01` - Filters VM, vApp, or resource pool names that start with the `<filter string>`, end with a "01", and have any set of characters in place of the wildcards (*). For example, `Production_VM01`.
 - `<filter string>*0?` - Filters VM, vApp, or resource pool names that start with the `<filter string>`, end with a "0" and have a character, number, or letter in place of the question mark (?). For example, `Production_VM01` and `Production_VM0A`, but not `Production_VM11`.

Select the objects that you want to restore.

NOTE:
Data Protector restores the full restore chain for each selected VMware object, beginning with the last full backup session (even if that full backup is outside the specified time interval) and ending with the last backup session performed during the specified time interval.

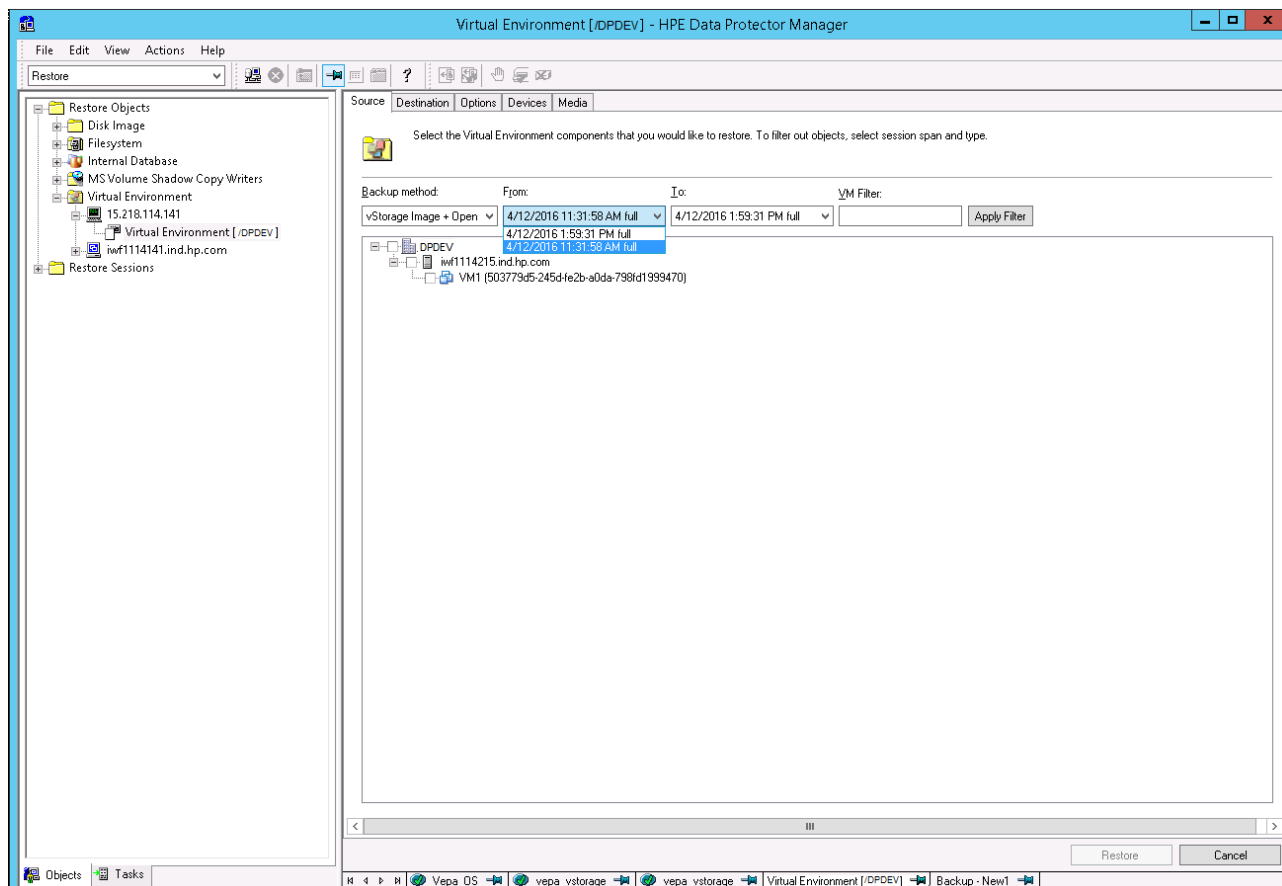
Selecting VMware objects for restore (vCenter Server client)



Power On (PowerOnOption ON): Power on the virtual machine from the backed up image, Smart Cache, or StoreOnce Catalyst device.

Live Migrate (PowerOnOption MIGRATE): Live Migration of the virtual machine from the backed up image, Smart Cache, or StoreOnce Catalyst device.

Selecting vStorage Image + OpenStack backup method for restore




Restoring as a new virtual machine

Right-click the selected virtual machine and then click **Restore As / Into** to restore it as a new virtual machine for vStorage Image backup method. A new dialog box opens to specify the name of the virtual machine.

NOTE:
The **Restore As/Into** option is specific to vStorage Image backup method.

When you select the **vStorage Image + OpenStack** backup method, you get to view the Nova instances that were backed up and its versions. The Shadow VM objects are not displayed during restore.



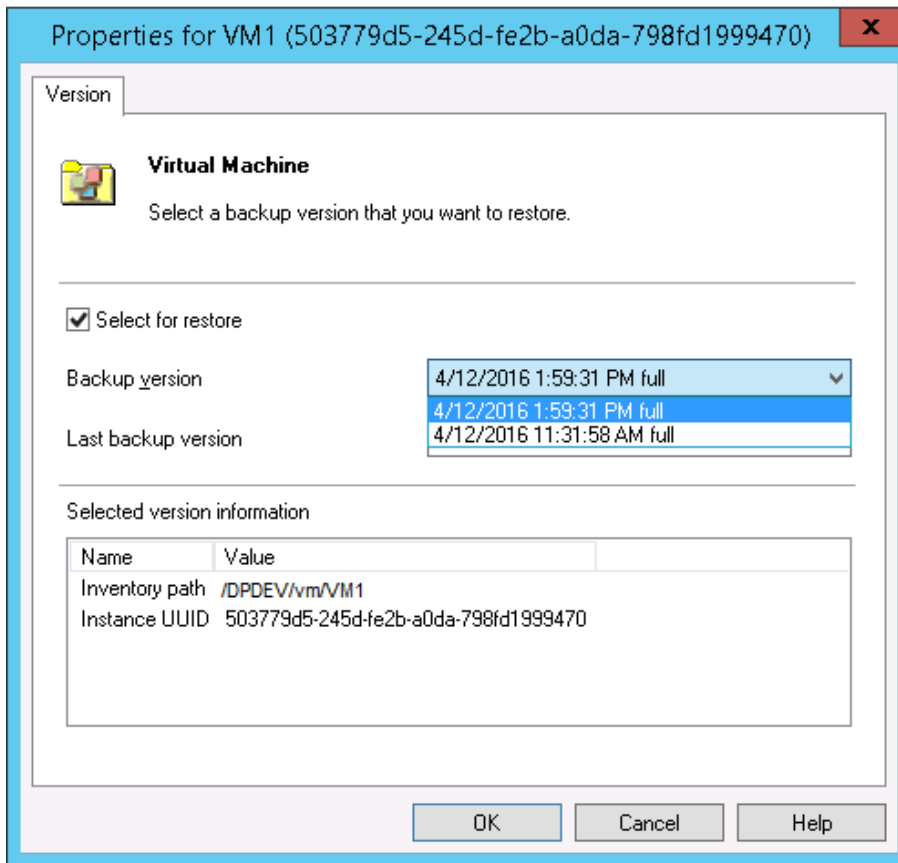
Restoring selected backup version

Right-click the selected virtual machine and click **Restore Version** to select a backup version that you want to restore.

A new dialog box opens for selecting the backup version. The object versions for the selected Nova Instances are displayed.


NOTE:

(Specific to **vStorage Image** backup method) A warning message VM is being restored to a different data center is displayed in the dialog box when the backup version is selected from a different datacenter.




4. In the Destination page, specify the restore destination. For details, see [Restore destination \(VMware vCenter Server and VMware ESX\(i\) Server clients\)](#), on page 373.

Restore destination (VMware vCenter Server client)

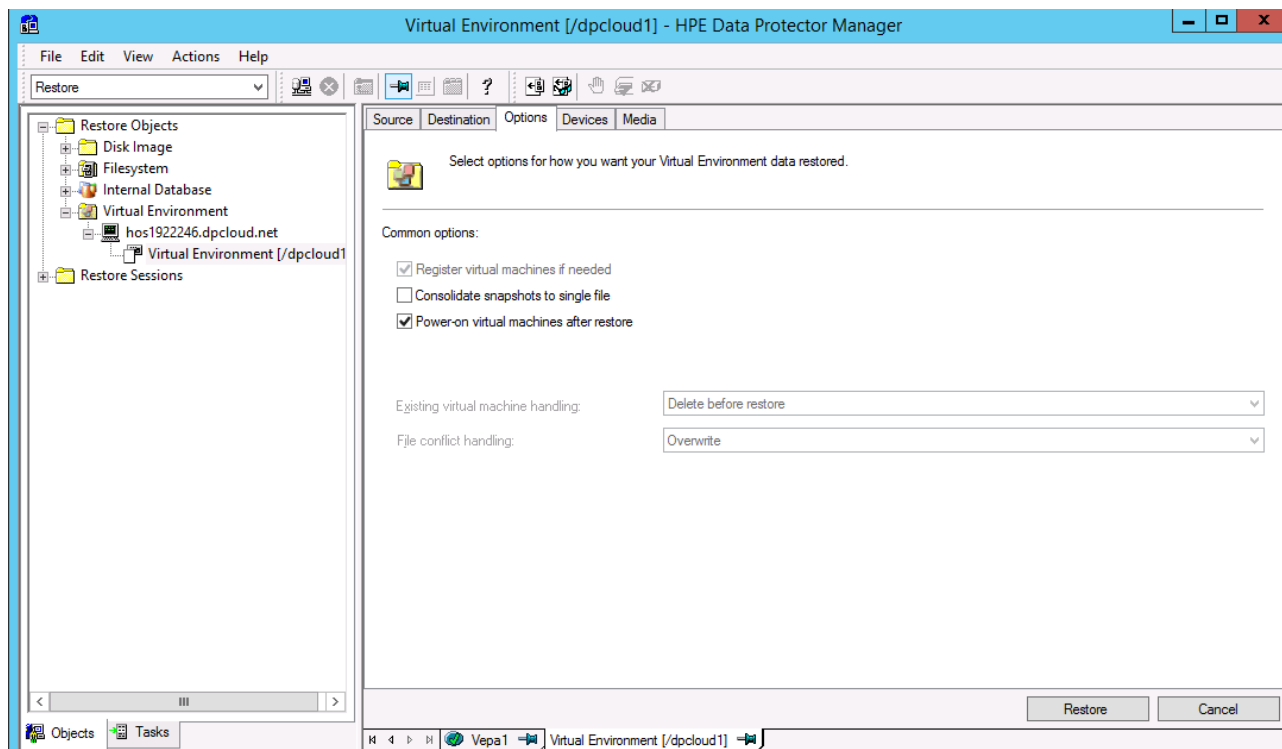


5. If you have selected Power On or Live Migrate, click the respective button to complete the operation.
6. In the Options page, specify the VMware restore options. For details, see [Restore options, on page 374](#).

Restore options (VMware vCenter Server and VMware ESX(i) Server clients)



Restore options (vStorage Image + OpenStack)



Here, VMs are deleted before restore, and Nova Instance and Shadow VMs are registered in the vCenter after restore.

7. In the Devices page, select devices to use for restore.
8. Click **Restore**.
9. In the Start Restore Session dialog box, click **Next**.
10. Specify **Report level** and **Network load**.

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

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

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

NOTE:
If the restore fails, see [Cleaning up a datastore after a failed restore, on page 388](#).

RMC integration restore

- For RMC backups, if the backup type is **Snapshot+Tape**, preference is given for the restore operation from the snapshot. If the snapshot is absent, the restore happens from the tape device.
- If the backup type is **Express Protect**, restore happens to the 3PAR LUNs first, and subsequently, to the location specified.

Restore destination (VMware vCenter Server and VMware ESX(i) Server clients)

GUI/CLI option	Description
Backup host / -barhost	Specifies the client with the Virtual Environment ZDB Integration for VMware installed to control the restore session. By default, the same client as used for the backup is selected.
Restore client / -apphost	Specifies the client that the selected virtual machine objects should be registered and restored to. By default, the client from which the virtual machines were backed up is selected. To change the client configuration, click the Connect button.
Restore to datacenter / -instance - newinstance	Select this option to restore a virtual machine to a datacenter. By default, virtual machines are restored to the original datacenter. You can restore from a 3PAR array to a datacenter.
Host/Cluster / -host/cluster	Select the ESX(i) Server system or the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original ESX(i) Server system or cluster.
Specific host / -specificHost	Select the specific ESX(i) Server system in the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original ESX(i) Server system.
Resource pool / -resourcePool	Select the resource pool on the ESX(i) Server system or the cluster to which virtual machines should be restored. By default, virtual machines are restored to the original resource pool.
Datastore / -store	Specifies the datastore to which the virtual machines should be restored. You can choose among all datastores that are accessible from the selected restore target host. If you leave this option empty, the virtual machines are restored to the original datastore.
Restore to directory / -directory	Select this option to restore virtual-machine files to a directory (outside of the datacenter) on the backup host. You can use the Browse button to find the target directory. After such a restore, the virtual machines are not functional. You need to manually move the restored virtual machine images to an ESX Server or ESXi Server system, using the VMware Converter as described in Recovering virtual machines after restore to a directory, on page 381 .

Restore options

GUI/CLI option	Description
Register virtual machines if needed /	Available if Restore to datacenter

GUI/CLI option	Description
<p>-register</p>	<p>is selected.</p> <p>Select this option to register restored virtual machines.</p> <p>If this option is not selected, you need to manually recover the restored virtual machines as described in Recovering virtual machines manually, on page 381.</p> <p>Default: selected.</p> <div style="border: 1px solid black; padding: 5px; background-color: #f0f0f0;"> <p>NOTE: The Register virtual machines if needed option is unavailable for selection with vStorage Image + OpenStack backup method.</p> </div>
<p>Consolidate snapshots to single file / -consolidate</p>	<p>Select this option to commit all snapshots (including non-Data Protector ones) to the virtual machine base once a virtual machine is restored.</p> <p>Available if Restore to datacenter is selected.</p>
<p>Power-on virtual machines after restore / -poweron</p>	<p>Select this option to power the virtual machines on once they are restored.</p> <p>Available if Restore to datacenter is selected.</p>

GUI/CLI option	Description
<p data-bbox="245 310 553 373">Existing virtual machine handling</p> <div data-bbox="245 394 578 709" style="border: 1px solid black; padding: 5px;"><p data-bbox="269 405 521 699">NOTE: This option is unavailable for selection with vStorage Image + OpenStack backup method. The Delete before restore option is used.</p></div>	<p data-bbox="613 310 1365 373">Specifies Data Protector's behavior when restoring existing virtual machines.</p>

GUI/CLI option	Description
<p>Delete before restore / -deletebefore</p>	<p>Select this option to delete an existing virtual machine before it is restored, and then restore it from new. The existing virtual machine is deleted even if it resides in a different datacenter from your target datacenter.</p> <p>This is the space-efficient option, but is less secure, since the old virtual machine is not available if the restore fails, so select it with caution.</p>
<p>Skip restore / -skip</p>	<p>Select this option to skip the restore of an existing virtual machine. This allows you to restore missing virtual machines without affecting existing ones.</p>
<p>Delete after restore / -deleteafter</p>	<p>Select this option to delete an existing virtual machine after it is restored. The existing virtual machine is deleted even if it resides in a different datacenter than your target datacenter. If the restore fails, the existing virtual machine is not deleted.</p> <p>Default: selected.</p> <p>This option is not available for VMware vCloud Director client.</p> <p>This option cannot be used if the virtual machine is in a suspended state. If the virtual machine is in a suspended state, do any of the following:</p> <ul style="list-style-type: none"> • Restore to a different location. • Select the Delete before restore option. • Power On or Off the virtual machine.
<p>Keep for forensics / -keep_for_forensics</p>	<p>Select this option to mark an existing virtual machine with a</p>

GUI/CLI option	Description	
		<p>timestamp. The virtual machine which is kept for forensics is powered off after the restore and remains at the original location. It does not affect consecutive backups of the original virtual machine.</p> <p>This option is not available for VMware vCloud Director and Microsoft Hyper-V clients.</p>
<p>File conflict handling</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>NOTE: This option is unavailable for selection with vStorage Image + OpenStack backup method. The Overwrite option is used.</p> </div>	<p>Specifies Data Protector's behavior when restoring existing files.</p>	
	<p>Overwrite / -overwrite</p>	<p>Select this option to overwrite existing files with those from the backup.</p> <p>Default: selected.</p>
	<p>Keep latest / -latest</p>	<p>Select this option to leave an existing file intact if it is more recent than the one from the backup. Otherwise, the existing file is overwritten with the one from the backup.</p>
	<p>Skip / -skip</p>	<p>Select this option to preserve an existing file (the file is not restored from the backup).</p>

Restoring using the Data Protector CLI

1. Log in to any client with the Data ProtectorUser Interface component installed.
2. Open the command prompt and change to the directory in which the omnir command is located.
3. Execute:

VMware vCenter Server or VMware ESX(i) Server client

```

omnir -veagent
      -virtual-environment vmware
      -barhost BackupHost
      -apphost OriginalVMwareClient

      -instance OriginalDatacenter
      -method vStorageImage| vStorageImageOpenStack
    
```

```
[-session BackupID]  
VirtualMachine [VirtualMachine...]  
[VMwareClient | Directory]  
  
VirtualMachine  
-vm VMPATH -instanceUUID vmInstanceUUID [-new_name NewVirtualMachineName][  
disk DiskName [-disk Disk...]]  
  
VMwareClient  
[-newinstance TargetDatacenter]  
[-store TargetDatastore]  
[-network_name TargetNetwork]  
[-destination TargetVMwareClient]  
[-consolidate] [-register][-poweron]  
[-deletebefore | -deleteafter | -skip | -keep_for_forensics]  
  
Directory  
-directory RestoreDirectory  
[-overwrite | -skip | -latest]  
  
Restore Options  
  
[-consolidate] [-register] [-poweron] [-PowerOnOption { ON | MIGRATE }] [  
deletebefore | -skip | -keep_for_forensics]
```

For a description of all the options, see the `omnir` man page or the *HPE Data Protector Command Line Interface Reference*.

NOTE:

You should not specify the `instanceUUID` parameter for `restore` while restoring the virtual machine backed up from Data Protector 8.1 and below.

IMPORTANT:

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you have to specify the session ID of the original backup session (that is, the backup ID) and not the session ID of the object copy session.

The `omnir` syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

NOTE:

If the restore fails, see [Cleaning up a datastore after a failed restore, on page 388](#).

Example (Restoring virtual machines to a datacenter)

Suppose you want to restore the virtual machine `/vm/machineA` and the individual disks (`scsi0:0` and `scsi0:1`) of the virtual machine `/vm/machineB`. At the time of backup, the virtual machines were running on the ESX Server systems that belonged to the datacenter `/MyDatacenter` managed by the vCenter Server system `vcenter.company.com`. The virtual machines were backed up with the `vStorage Image` backup method.

To restore them to the original location, using the backup session `2011/01/11-1`, and to ensure that the newly restored virtual machines are put online when the session completes, execute:

```
omnir -veagent -virtual-environment vmware -barhost backuphost.company.com -apphost vcenter.company.com -instance /MyDatacenter -method vStorageImage -session 2011/1/11-1 -vm /vm/machineA -vm /vm/machineB -disk scsi0:0 -disk scsi0:1 -poweron
```

To restore the virtual machine `/vm/machineA` with an instanceUUID `503eeaac-6fae-7898-73e1-93b722a0517c`, execute:

```
omnir -veagent -virtual-environment vmware -barhost backuphost.company.com -apphost vcenter.company.com -instance /MyDatacenter -method vStorageImage -session 2011/1/11-1 -vm /vm/machineA -instanceUUID 503eeaac-6fae-7898-73e1-93b722a0517c -disk scsi0:0 -disk scsi0:1 -poweron
```

Example (Restoring virtual machines to a directory)

Suppose the virtual machines `/MyVirtualMachines/machineA` and `/MyVirtualMachines/machineB` were backed up in the session `2011/02/12-5` from the datacenter `/MyDatacenter` that is managed by the vCenter Server system `vcenter.company.com`, using the `vStorage Image` backup method. To restore the virtual machines outside of the datacenter, to the directory `C:\tmp` on the backup host `backuphost.company.com`, execute:

```
omnir -veagent -virtual-environment vmware -barhost backuphost.company.com -apphost vcenter.company.com -instance /MyDatacenter -method vStorageImage -session 2011/2/12-5 -vm /MyVirtualMachines/machineA -vm /MyVirtualMachines/machineB -directory c:\tmp
```

Example (Restoring object names with instanceUUID in its name)

To support restore of object names with instanceUUID in its name, execute:

```
omnir.exe -veagent -virtual-environment vmware -barhost barHostName -apphost appHostName -instance instanceName -method vStorageImage -session sessionID -vm vmPath -instanceUUID vminstanceUUID -register -poweron -deletebefore
```

Example (Restoring Nova Instance to its original location)

To restore object names with instanceUUID in its name, execute:

```
omnir -veagent -virtual-environment vmware -barhost barHostName -apphost appHostName -instance /Datacenter -method vStorageImageOpenStack -session sessionID -vm vmPath -instanceUUID novainstanceUUID -register -deletebefore
```


Recovering virtual machines manually

There are two different scenarios in which you need to recover virtual machines manually after they have been restored with Data Protector:

- If you have restored the virtual machines to a directory on a backup host (**Restore to directory**).
For details, see [Recovering virtual machines after restore to a directory, below](#).
- If you have restored the virtual machines to a datacenter (**Restore to datacenter**) without selecting the restore option Register virtual machines if needed.
For details, see [Recovering virtual machines after restore to a datacenter, on page 387](#).

Recovering virtual machines after restore to a directory

The steps for recovering virtual machines after restore to a directory depend on the format in which the virtual machine configuration file was backed up.

Recovering with the VM configuration file in the VMX format

Suppose the virtual machine `helios` was restored to the directory `C:\tmp\helios` on the backup host using the following backup session:

- Backup method: **vStorage Image**
- Backup type: Incremental
- CBT: Enabled and used

To move the virtual machine files manually to the ESX(i) Server system `dioxide.company.com` managed by the vCenter Server system `bmwvc2.company.com`, using the VMware Converter:

1. Display the contents of the directory `C:\tmp\helios`:

```
helios.vmdk
helios.vmx
helios.vmdk
scsi0-0.cbt
scsi0-0.meta
helios-flat.vmdk
helios.vmx-1
helios.vmdk-1
scsi0-0.cbt-1
scsi0-0.meta-1
helios.vmx-2
helios.vmdk-2
scsi0-0.cbt-2
scsi0-0.meta-2
```

Note that all files backed up in the last full, differential, and the selected incremental session are restored.

2. Share the folder `C:\tmp\helios` so that it can be accessed from the system with the VMware Converter installed.

3. Log in to the system with the VMware Converter installed and open the VMware Converter user interface.
4. Click **Convert Machine** to open the Conversion wizard.
5. In the Source System page, select **VMware Workstation or other VMware virtual machine** for the source type, browse to the C:\tmp\helios directory, and select the helios.vmx file.

Click **Next**.

Conversion (Source System)

Conversion

Source System
Select the source system you want to convert

Source System
Destination System
Destination Virtual Machine
Destination Location
Options
Summary

Source: C:\helios\helios.vmx **Destination:** none

Select source type: VMware Workstation or other VMware virtual machine

Convert a virtual machine from VMware Workstation, VMware Player, VMware product.

Browse for source virtual machine or image

Virtual machine file: C:\helios\helios.vmx Browse...

[View source details...](#)

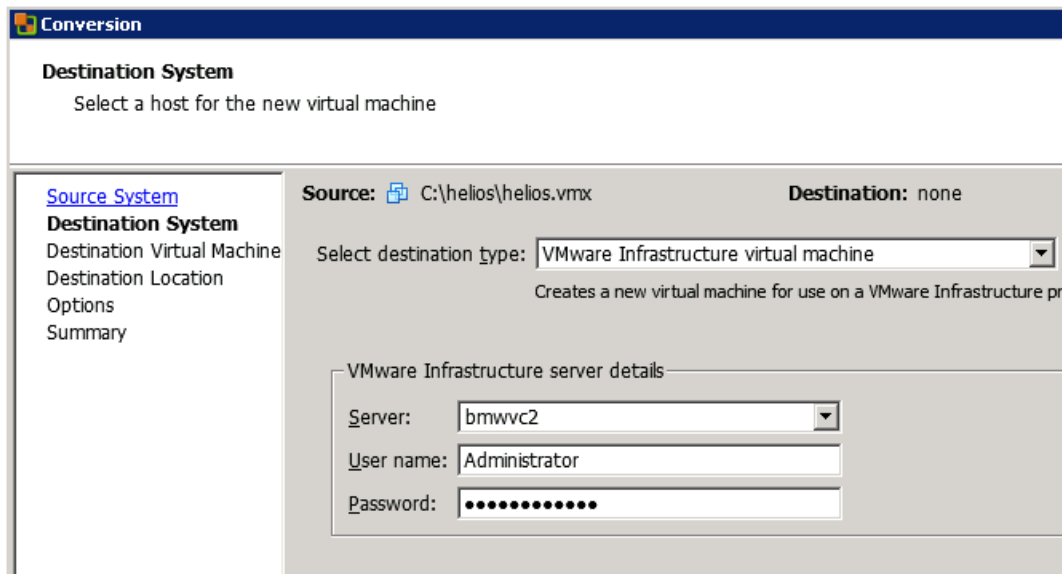
NOTE:

In our example, the VMware Converter is installed on the backup host.

6. In the Destination System page, select **VMware Infrastructure virtual machine** for the destination type and provide the login credentials for the vCenter Server system.

Click **Next**.

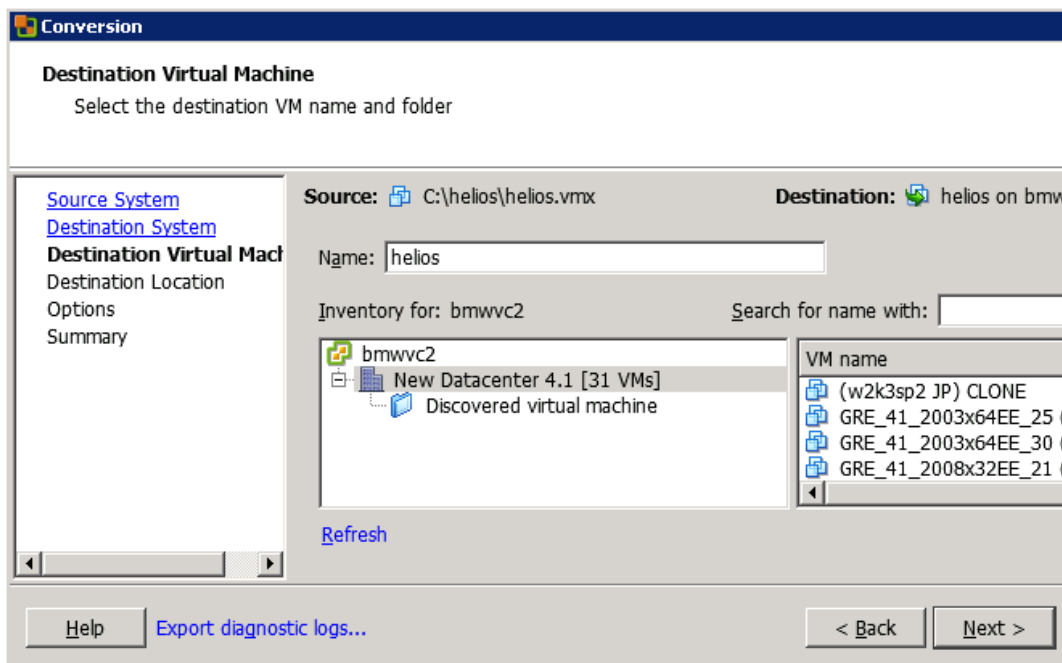
Conversion (Destination System)



7. In the Destination Virtual Machine page, specify the name under which the virtual machine should be recovered.


Click **Next**.

Conversion (Destination Virtual Machine)




8. In the Destination Location page, select the destination ESX(i) Server system and datastore.

Conversion (Destination Location)




9. In the Options page, edit the options and click **Next**.

Conversion (Options)



10. In the Summary page, review your selection and click **Finish**.

Conversion (Summary)



11. Open the Datastore Browser and upload the files created in the incremental and differential backup sessions to the virtual machine directory:

```
helios.vmx-1  
helios.vmdk-1  
scsi0-0.cbt-1  
scsi0-0.meta-1  
helios.vmx-2  
helios.vmdk-2  
scsi0-0.cbt-2  
scsi0-0.meta-2
```

12. Power the virtual machine on.

Recovering with the VM configuration file in the XML format

Follow the procedure:

1. Open vSphere Client and log in to an ESX(i) Server or vCenter Server system.
If the virtual machine is still configured, remove all its hard disks:
 - a. In the inventory object tree, right-click the virtual machine and select **Edit Settings**.
 - b. In the Virtual Machine Properties window, in the **Hardware** tab, select each hard disk and click **Remove**.
 - c. Click **OK** to confirm the removal.If the virtual machine is no longer present, configure a new virtual machine without hard disks, and use the name of the original virtual machine.
In either case, remember the associated datastore name.
2. Upload the virtual machine files that were created during the backup session:
 - a. In the inventory objects tree, select the ESX(i) Server system that hosts the virtual machine.
 - b. Click the **Configuration** tab and select **Storage** under Hardware.
 - c. Right-click the datastore name and select **Browse Datastore**.
 - d. In the Datastore Browser window, in the folder tree, select the virtual machine folder, and click a corresponding icon on the window toolbar. Select **Upload File** or **Upload Folder** as appropriate.
 - e. Select all applicable files and complete the upload.
3. Add hard disks to the virtual machine by reusing their backup copies:
 - a. In the inventory object tree, right-click the virtual machine and select **Edit Settings**.
 - b. In the Virtual Machine Properties window, click **Add**.
 - c. In the Add Hardware window, select **Hard Disk** and click **Next**.
 - d. Select **Use an existing virtual disk** and click **Next**.
 - e. Click **Browse**.
 - f. In the Browse Datastores window, browse to the appropriate datastore and open the virtual machine folder. Select the virtual disk file and click **OK**.
 - g. Follow the Add Hardware wizard to complete the process.
 - h. Repeat the substeps from b to g for each additional hard disk for which a backup copy exists.
4. Power the virtual machine on.

Recovering virtual machines after restore to a datacenter

If you have restored a virtual machine to a datacenter without selecting the option Register virtual machines if needed:

1. Open the Datastore Browser and browse to the restored virtual machine directory.
2. Right-click the virtual machine *.vmtx file and select **Add to Inventory**.
3. Follow the wizard and click **Finish**.

Restoring using another device

You can restore using a device other than that used for backup. For details, see the *HPE Data Protector Help* index: “restore, selecting devices for”.

Cleaning up a datastore after a failed restore

Sometimes, when a virtual machine restore fails, Data Protector creates extra files on the virtual machine datastore. If these files are not deleted, corrupt virtual machine backups may be created in subsequent sessions and, consequently, restore from such a backup also fails.

Suppose the virtual machine `MyVirtualMachine` failed to be restored. To clean up the datastore after the restore:

1. Open the VMware vSphere client.
2. Right-click the virtual machine and select **Delete from disk**.
3. Open the **Datastore Browser**.

The directory `MyVirtualMachine` should no longer be there.

Check if there are any extra directories:

`MyVirtualMachine_1`

`MyVirtualMachine_2`

and so on.

Right-click each such directory and select **Delete from disk**.

Disaster recovery

Disaster recovery is very complex, involving different products from different vendors. Check the instructions from the guest operating systems and VMware on how to prepare for it.

The following are the main steps needed to recover a virtual machine after a disaster:

1. Reinstall the VMware environment. The configuration should be the same as during the backup.
2. Install Data Protector in the newly configured environment.
3. Restore the service console of the ESX Server system on which the virtual machine was running to the newly configured ESX Server system from a Data Protector filesystem backup.

For details on what to restore, see the topic “ESX Server Configuration Backup and Restore procedure” at <http://kb.vmware.com/selfservice/microsites/microsite.do>.

For details on how to restore from a filesystem backup, see the *HPE Data Protector Help*.

4. Restore the original vCenter database (if needed). For details, see the Data Protector integration that was used to back up the database.
5. Restore the virtual machine from a Data Protector Virtual Environment backup as described in this chapter.

Instant recovery

For general information on instant recovery, see the *HPE Data Protector Concepts Guide* and the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Prerequisites

- ESX Server
- 3PAR arrays
- vCenter
- Application host on virtual machine
- Backup host on physical machine

NOTE:

Both the hosts including the ESX Server should have access to 3PAR arrays.

HPE 3PAR ZDB Instant recovery

Data Protector offers HPE 3PAR ZDB instant recovery for agents within a VMware Virtual Machine (VM). This is supported only for physical Raw device mapping (RDM).

To perform an instant recovery on Linux and Windows:

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

For example:

```
/sqlplus /nolog
connect sys/oracle@APPN as sysdba
sql> shutdown immediate
sql> exit
```

2. Enable the `omnirc` option on Application host:

```
ZDB_IR_MANUAL_AS_PREPARATION=1
```

3. **For Linux:** Dismount volume:

```
# umount /dev/3PAR_ESX2/lvo10
```

For Windows: Offline the disks.

4. Prepare Application Host to remove volume (export, deactivate, and backup Volume Group).
5. Remove hard disks from Application host on vCenter server.
6. Rescan volumes on VM, and confirm that the disk is not presented any more on the Application Host.
7. Execute Instant Recovery.

NOTE:

If you are using the Oracle integration, ensure that you deselect the **Recovery** checkbox.

8. Add the hard disks back to the application host from the vCenter server.
9. Rescan Application Host for new volume.

10. Add the exported Volume Group.
11. **For Linux:** Mount volume.
For Windows: Online the disks.
12. Follow the steps as mentioned in the [Oracle database recovery after the instant recovery](#) section.

Monitoring sessions

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

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

To monitor a session, see the *HPE Data Protector Help* index: “viewing currently running sessions”.

Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Virtual Environment ZDB integration for VMware.

For general Data Protector troubleshooting information, see the *HPE Data Protector Troubleshooting Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. For more information, see the *HPE Data Protector Help* index: “patches”.
- See the HPE Data Protector Product Announcements, Software Notes, and References for general Data Protector limitations, as well as recognized issues and workarounds.
- See <https://softwaresupport.hpe.com/> for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

If your configuration, backup, or restore failed:

- Examine system errors reported in the `debug.log` on the backup host.
- Check if you can do a filesystem backup and restore on the problematic client. For information, see the *HPE Data Protector Help*.

Additionally, if your backup failed:

- Check the configuration of the vCenter Server or standalone ESX(i) Server system as described in [Checking the configuration of VMware clients, on page 341](#).

Additional actions:

- After performing the file delete operation in a guest VM, the full CBT backup session has the similar backed up data as the previous full CBT backup session. The reason is that the VMware does not reclaim the disk space after the file delete operation is performed in the guest VM.

Workaround-1:

1. Reclaim the unused space of a virtual disk in ESXi/ESX 4.1 or later. For details, see

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004155

2. Disable and re-enable the CBT. For details, see

http://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=1031873

Workaround-2:

Perform a Storage vMotion on the virtual machine or VMDK to a datastore formatted with a different block size. For example, if the VMDK is on a datastore formatted with 2 MB blocks, format the target VMFS datastore with a 1 MB, 4 MB, or 8 MB block size. Once this is done, the CBT has to be reset.

- Before enabling CBT and running a backup, ensure that user-created snapshots are not available in the virtual machine. Otherwise, the following error message is displayed:

```
[Warning] From: VEPALIB_VMWARE@hostname "<DataCenter>" Time: Date Time
[172:390] Virtual Machine 'VM': User Snapshot(s) found. Changed Block Tracking
cannot be enabled.
```

For details, see http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1020128

Workaround:

- Delete all user-created snapshots associated with the virtual machine before enabling CBT.
- Consolidate changes, if the changes in snapshot needs to be backed up using CBT.
- An incremental or differential CBT backup session falls back to a FULL CBT backup session when performing an incremental or differential CBT backup session after a virtual disk migration using the VMware Storage VMotion, the backup falls back to a FULL CBT session.

Workaround:

The issue is resolved with ESX 5.5 update 2. The reason is that the CBT is reset after a virtual disk migration using the storage VMotion. For details, see

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2048201

Problems

Problem

An incremental or differential CBT backup session fails

When performing an incremental or differential backup session with the Use changed block tracking option enabled, the session fails with an error similar to the following:

```
[Critical] From: OB2BAR_VEPA_BAR@droid.company.com "/New Datacenter 4.1"  
Time: 2/10/2011 11:14:52 AM  
Virtual Machine 'ddd': Could not gather changed blocks on disk scsi0:0 ...
```

The reason may be that you performed a restore session and forgot to run a full backup session to start a new backup chain.

To make sure changed block tracking is working properly, certain requirements need to be met. For details, see

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd;=displayKC&externalId;=1020128

Action

1. Run a full backup session.
2. Run an incremental or differential backup session.

Problem

After a restore or move to a different folder, backups are not performed correctly

After you have restored or moved a virtual machine to a different folder, the virtual machine is not backed up correctly. For example, instead of an incremental backup, a full backup is performed.

The reason for this is that the datacenter configuration file has been updated. As a result, it contains two virtual machine sections with the same UUID; this is from where the inconsistencies arise.

Action

Re-configure the virtual machine:

1. Open the backup specification.
2. In the Source page, right-click the VMware client and select **Configure Virtual Machines**.
3. Click **OK**.

If addition of disks is done after the backup on the backed up Nova instance, the following error message is displayed:

```
Error: Virtual Machine '9d28cd95-c158-45ec-b606-53f7c63a2a78': Cannot perform  
Restore. Found new disk attached to the Nova Instance.
```

When a user needs a point in time restore, and a new disk was added after that point in the backup, restore is not allowed. User has to detach the new volumes attached to the instance from the OpenStack and perform the restore again.

It is required to perform a full backup after the addition of the new disks.

Problem

A restore session using SAN transportation mode fails

A restore session that is using SAN transportation mode fails with a message similar to the following:

```
[Critical] From: OB2BAR_VEAgent@dpi00019.company.com  
"/BlrVirtual01_ESX401" Time: 13-03-2011 12:22:57  
Virtual Machine 'Win2k3_x64_dpi00002': Error restoring item
```

```
\a1f9f4e3-482d-4b7f-afcb-cb16babe1980\%2FB1rVirtual01_ESX401\vm  
\%2FB1rVirtual01_ESX401%2Fhost%2Fclus01%2Fwin2k3_x64_dpi00002\  
images\3\scsi2:15.
```

This may happen if a storage volume that is shared between the backup host and an ESX(i) Server system is read-only.

Action

1. Log in to the backup host and open the command prompt.
2. Execute diskpart.

```
C:\Users\Administrator>diskpart
```

```
Microsoft DiskPart version 6.1.7600  
Copyright (C) 1999-2008 Microsoft Corporation.  
On computer: TPC134
```

3. Set the SAN policy to onlineAll.

```
DISKPART> san policy=onlineAll
```

```
DiskPart successfully changed the SAN policy for the current  
operating system.
```

4. Select the disk (storage volume) that should be used for restore.

```
DISKPART> list disk
```

Disk ###	Status	Size	Free	Dyn	Gpt
Disk 0	Online	136 GB	1024 KB		
Disk 1	Offline	14 GB	14 GB		
Disk 2	Offline	14 GB	14 GB		
Disk 3	Offline	14 GB	14 GB		
Disk 4	Offline	14 GB	14 GB		
Disk 5	Offline	50 GB	50 GB		
Disk 6	Offline	14 GB	14 GB		
Disk 7	Offline	14 GB	14 GB		

```
DISKPART> select disk 1
```

5. Bring the disk online.

```
DISKPART> online disk
```

```
DiskPart successfully onlined the selected disk.
```

6. Ensure the disk is not read-only.

List the disk properties.

```
DISKPART> detail disk
```

```
HP OPEN-V SCSI Disk Device
Disk ID: 00000000
Type   : FIBRE
Status : Online
Path   : 0
Target : 0
LUN ID : 0
Location Path : UNAVAILABLE
Current Read-only State : Yes
Read-only   : Yes
Boot Disk   : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
```

There are no volumes.

Clear the read-only attribute.

```
DISKPART> attribute disk clear readonly
```

Disk attributes cleared successfully.

List the disk properties again.

```
DISKPART> detail disk
```

```
HP OPEN-V SCSI Disk Device
Disk ID: 00000000
Type   : FIBRE
Status : Online
Path   : 0
Target : 0
LUN ID : 0
Location Path : UNAVAILABLE
Current Read-only State : No
Read-only   : No
Boot Disk   : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
```

There are no volumes.

Exit the session.

```
DISKPART> exit
```

7. Restart the restore session.

Problem

Slow performance of vepa_util.exe browse command on newer Red Hat Enterprise Linux (RHEL) versions

When executing the `vepa_util.exe browse` command on a newer RHEL version its performance is significantly slower than on other operating systems.

Action

The root cause of the problem is that on newer versions of RHEL systems name service cache daemon is not enabled by default.

Start the name service cache daemon by invoking the following command: `/etc/init.d/nscd start`. To enable automatic daemon start-up during system start-up, execute: `chkconfig nscd on`.

Problem

A restore job fails when restoring virtual machines to ESX(i) hosts managed by vCenter Server 5.x or higher

When restoring a virtual machine to an ESX(i) host, which is managed by a vCenter Server 5.x or higher, the restore job fails and the virtual machine is not restored successfully. Starting with ESX(i) 5.0, VMware has blocked the ability to restore a virtual machine to an ESX(i) which is managed by vCenter.

Action

To resolve this issue, either restore virtual machines through the vCenter Server or restore them through ESX/ESX(i) after unmanaging the host from the vCenter Server.

For more details on how to unmanage an ESX/ESXi host from vCenter Server, see <http://kb.vmware.com/kb/2038838>

Problem

Restore of a virtual machine using an ESX(i) Server system ends with a corrupted VM guest operating system

In a vSphere environment, when you restore a virtual machine to a `/ha-datacenter` using an ESX(i) 5.0 Server system or later as a restore client, the restore session ends successfully, but the guest operating system on the virtual machine is corrupted.

Action

When selecting the objects for restore, in the **Destination** page, in the **Restore client** drop-down list, select a vCenter Server instead of an ESX(i) Server system.

Problem

In some cases, when backup is started and the backup specification contains multiple Virtual Machines, the following message appears:

```
[Major] From: BSM@company.name.com "backup_spec_name" Time: 4/10/2014 2:55:07 PM  
[61:2052] Bar backup session was started but no client connected in 600 seconds.  
Aborting session!
```

The VEAgent collects metadata information for all the Virtual Machines before starting the backup process. While the VEAgent is busy collecting information, the BSM waits for 10 minutes (default timeout) for the VEAgent to get connected. After the default time has elapsed the BSM times-out with the message shown above.

Action

Increase the value of the timeout variable available in Data Protector Global Options.

```
SmWaitForFirstBackupClient=WaitForInMinutes
```

```
SmWaitForFirstBackupClientSec=WaitForInSeconds
```

Problem

Backup of Virtual Environments can fail with the following error messages:

- Exception occurred while creating VM snapshot
- Backup of object failed

These errors may occur if the timeout values on vCenter and ESX hosts are low.

Action

Increase the timeout values in the vCenter configuration file (`vpzd.cfg`) and the ESX hosts configuration file (`vpza.cfg`).

For more details on the resolution, see the following:

<http://kb.vmware.com/kb/1017253>

<http://kb.vmware.com/kb/1005757>

NOTE: NOTE: You can change the values suggested here based on your requirements.

In addition, check the timeout values and operational timeout values in the vCenter:

Increase the operational timeout values on vCenter.

1. Log in to the vCenter server with administrator credentials.
2. Go to the following location:
vCenter Server Settings > Timeout Settings.
3. Under Client Connection Timeout, set the following values:
 - Normal Operations timeout: 600
 - Long Operations timeout: 2000

Increasing the timeout value in vCenter

To increase the number of idle connections between ESX and the vCenter host, add the following values in the `vpzd.cfg` file within the `<config>` and `</config>` tags:

```
<vpzd>  
<maxHostPooledConnections>20000</maxHostPooledConnections>  
</vpzd>
```


NOTE: NOTE: This will reduce the creation of additional TCP connections to the proxy during host synchronization.

Problem

After upgrading, if the backup specification is created using Data Protector 7.00, 7.01, or earlier versions, the VEAgent backup fails with the following error:

```
[Critical] From: VEPALIB_VMWARE@<hostname> "<Datacenter>" Time: <Date Time>  
No Objects found for backup
```

Action

After upgrading, perform the following steps:

1. Re-create the backup specification with the same VM selection and options as it was earlier.
2. Run the backup again.

Problem

The Virtual Environment Integration agent (VEPA) and the Session Manager stall while waiting for the time-out value to elapse.

While running multiple parallel VEPA backup sessions, you could encounter a scenario where a small fraction (1 or 2) of the total number of sessions may stall in such a way that the VEPA and BSM are not responsive until the time-out period. This could be because the vCenter is loaded with multiple concurrent connection requests from the VEPA agent, from parallel VM backup sessions.

The stalled VEPA and BSM processes for these sessions may finally time-out with the following message:

```
[Major] From: BSM@machineName "barlist7" Time: 6/13/2014 2:05:41 AM  
[61:1002] The OB2BAR Backup DA named "/Datacenter" on host machineName reached its  
inactivity timeout of xxxx seconds. The agent on host will be shutdown.
```

Action

After the time-out period:

1. Stop the vepa_bar processes manually and wait for the BSM process associated with it to close as well. (The BSM closes after the vepa_bar exits.)
2. Restart the backup specification that contains the failed VM objects.

Note: If the time-out period elapses before the backup starts (while resolving the objects), then increase the `SmWaitForFirstBackupClient` parameter in Internal Database -> Global Options.

Problem

Parallel backup sessions fail

If the VMs in the backup specification belong to the same LUN (already being backed up), and the mount proxy hosts that are used for the backups are the same then, for parallel backups the second backup session will fail.

Action

Use a different mount proxy host for the second backup session.

Problem

The zero downtime backup of virtual machines from the 3PAR replica fails with the following error message:

If the backup specification has large number of virtual machines or the load on the virtual center or the ESX server is high, the zero downtime backups may fail with following error message:

```
[Major] From: BSM@hostname.com "New2" Time: MM/DD/YYYY HH:MM:SS AM
```

```
[61:2052] Bar backup session was started but no client connected
```

```
in XX seconds.
```

```
Aborting session!
```

Action

Extend the timeout value by resetting the Data Protector `SmWaitForFirstBackupClient` global options variable. For details on how to set the variable, see the Data Protector online help.

Problem

For Linux virtual machines, the IP configured is lost after performing a restore operation.

This is likely to happen in the following scenario:

Before performing a Restore operation, if you click **Keep for forensics** or **Delete after restore** option in the restore work flow, then the restored VM will have a new NIC set to DHCP, and its original NIC goes into hiding.

Action

This is a known VMware / Linux limitation. For a virtual machine that is running a Linux guest operating system, while restoring that virtual machine, the ESX server may assign the virtual machine a new (virtual) MAC address. After restarting the virtual machine, you may have to configure its MAC address. For example, the original MAC address of the virtual machine may be in a configuration file that has to be updated as mentioned in the following KB articles.

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002767

and

<http://www.uptimemadeeasy.com/vmware/fixing-eth0-mac-address-vmware-clone-restore/>

Problem

Virtual machine restore: Could not find object on three disks in virtual environment.

While restoring a virtual machine, the object could not be found on the three disks in the virtual environment.

Action

To enable partial restore of a virtual machine that is no longer available on the data store, proceed as follows:

1. Create a temporary virtual machine with the same UUID as the original backup.
2. Restore partial data from backup to the temporary virtual machine.

Problem

Data Protector Virtual Environment integration (VEPA) backup sessions that use Backup to Disk (B2D) gateways may fail.

In virtual environment backups that use B2D gateways, when the Cell Manager, VEPA agent, Media Agent gateway for the server side deduplication are on the same machine, the backups may fail.

Action

You may perform the following tasks to solve this problem:

- Move the Media Agent, and the VEPA agent to another host.
- Reduce the number of concurrent streams in the Media Agent.
- Set the omnirc variable OB2BMAUPDT to 10000.
- Increase the process priority of the Backup Session Manager.
- Switch to source-side deduplication.
- Switch the VEPA backups from SAN mode to NBD mode.
- Avoid monitoring sessions in the Data Protector GUI.

Problem

Backups of VMware virtual machine disk may fail

Backups of VMware virtual machine disk may fail with the following error message:

```
[ 20] [VddkUtil::diskLibWarning] VixDiskLib: Failed to load vixDiskLibVim.dll :  
ErrorCode = 0x7f.
```

Action

This is a known VMware issue, and may be solved if you perform the following tasks:

1. Install the vCenter and the VDDK on separate servers.
2. Copy the libldap_r.dll and liblber.dll from \Program Files\OmniBack\lib\vddk\AMD64 to \Program Files\OmniBack\bin.
3. Copy the libcurl.dll from \Program Files\OmniBack\lib\x8664 to \Program Files\OmniBack\bin.

For more details on the resolution, see <http://kb.vmware.com/kb/2120818>.

Problem

Booting of the Windows virtual machine fails after restore

When the Windows virtual machine is booted after a restore operation, the booting process fails with the following error: No operating system found.

Action

This is a known VMware problem. This error may occur if the operating system disks and the data disks are from different controller types, such as SCSI and IDE. For more details, see

<http://kb.vmware.com/kb/1023592>.

You must change the boot order of the virtual machine using the options `bios.bootOrder` and `bios.hddOrder`. For the detailed procedure, see <http://kb.vmware.com/kb/2011654>.

Problem

VMware ZDB backup, Power On and Live Migrate may fail

Some times, the VMware ZDB backup session, Power On and Live Migrate sessions may fail with the following error message:

```
[Critical] From: VEPALIB_VMWARE@<HostName> "<AppName>" Time: <Timestamp>
Error mounting datastores
```

Action

This error may occur because the source ESX host or the mount ESX may contain some erroneous unresolved disk volumes that cannot be resolved to the data store. You must identify such volumes on the source ESX server or the mount proxy ESX server host, and remove the presentation of those volumes from the hosts.

Problem

Error message appears during Power On and Live Migrate

The following error message appears during Power On and Live Migrate sessions:

```
Failed to install/start NFS service.
```

Action

Open the Windows PowerShell window on the mount proxy host, and run the following command:

```
PS C:\Program Files\OmniBack\bin> .\nfsServiceCheck.ps1
True
PS C:\Program Files\OmniBack\bin>
```

Problem

Error while creating virtual machine snapshots and object backup fails

An error occurs when creating the snapshots of a virtual machine :

```
[Critical] From: VEPALIB_VMWARE@<HostName> "<AppName>" Time: <Timestamp>
Error mounting datastores

[Normal] From: VEPALIB_VMWARE@BACKUPHOSTNAME "/DATACENTERNAME"
Virtual Machine 'VMNAME': Creating snapshot ...

[Major] From: VEPALIB_VMWARE@BACKUPHOSTNAME "/DATACENTERNAME "
Virtual Machine 'VMNAME': Error removing snapshot

[Critical] From: VEPALIB_VMWARE@BACKUPHOSTNAME "/CPD2" Time: 19/03/2016 8:01:48
Backup of object failed.
```

```
Name: VMNAME  
Path: / DATACENTERNAME /DATASTORE/ VMNAME  
InstanceUUID: IUUIDOFVM
```

This issue occurs when the virtual machines are located in Site Recovery Manager (SRM).

Action

The create snapshot API is disabled for virtual machines in SRM and therefore, backup operation is not supported.

Problem

Virtual Machines with virtual hardware version 4 fails to boot after the restore to Data Center

When the Virtual Machine is restored to datacenter from the Data Protector GUI, the virtual machine fails to boot although the restore is successful.

Action

1. From the Data Protector GUI, restore virtual machine to a folder on to backup host.
2. From the vCenter, create a new virtual machine without any disks attached.
3. Go to the vCenter datastore browser, and upload the vmdk files from the folder created in [Step 1](#).
4. Edit the VM Settings from the vCenter to attach the uploaded vmdk files.
5. Restart the virtual machine to boot.

Problem

GRE, Power On, and Live Migrate operation fails

GRE, PowerOn and Live Migrate operation fails with the following error message:

```
Object locked: The VM <VM Name> could be locked by another process for  
recovery/power on/live migrate.
```

Please retry after the process is either done or cancelled.

Action

Ensure the following:

- Recover Files window displaying the Browse option is closed.
- The virtual machine does not have ongoing restore operations (Object Copy, Restore, Power On, Live Migrate, or GRE) from the StoreOnce Catalyst device.
- Clean up the powered on virtual machine.

Problem

Restore cannot be performed, if a new disk found is attached to the Nova instance

If a new disk found is attached to the Nova instance, the following error message is displayed:

```
Virtual machine '9d28cd95-c158-45ec-b606-53f7c63a2a78': Cannot perform Restore.  
Found new disk attached to the Nova instance.
```

Action

When a user needs a point in time restore, and a new disk was added after that point in the backup, restore is not allowed. User has to detach the new volumes attached to the instance from the OpenStack and perform the restore again.

Problem

At the time of restore, if the Shadow VM part of the backed up Nova Instance is attached to another Nova instance, the following error message is displayed in the session logs:

```
Virtual machine '9d28cd95-c158-45ec-b606-53f7c63a2a78': Can not perform restore, as the related shadow VM is attached to another Nov Instance '8d28ad65-c158-45ec-b606-53f7c63a4578'
```

Action

Detach the Shadow VM from the new instance and start the Restore.

Problem

The restored instance in the OpenStack dashboard does not reflect the correct state and remains in error state

Action

After the restore completes, restart the Compute Nova proxy service, and manually reset the error state to Active. For more information, see [Restore of Nova Instances and Shadow VMs backed up with vStorage Image + Openstack method](#).

Problem

Restore or object operation might fail if GRE, Power On, or Live Migrate operation is being performed for the same VEPA backup session to StoreOnce Catalyst

The object operations (object copy) will fail for a backup session "X" if GRE is being performed for the same object which is backed up in session "X". Following error is displayed in Data Protector object copy session:

```
[Major] From: RMA@hostname <DataCenter> Time: <Timestamp>
```

```
Cannot open device (StoreOnce error: The object is already locked and multiple open sessions not supported, or the server is unable to lock any more objects due to resource constraints)
```

Action

Perform the following:

- Ensure that the virtual machine does not have ongoing restore operations (Object Copy, Restore, Power On, Live Migrate, or GRE) from the StoreOnce Catalyst device.
- Clean up the powered on virtual machine.
- Retry after the active GRE, Power On, or Live Migrate operation is completed.

Problem

A backup session to StoreOnce Catalyst is not eligible for Power On and Live Migrate if:

- "No Logs" option is selected in the backup specification or
- The following error message is seen in the backup session report:
[Major] From: BSM@hostname.com <VMname> Time: <Timestamp>
[61:4039] Following error occurred while storing detail catalog information for device <Catalyst_device> with loaded medium <Catalyst_medium> to Data Protector Internal Database: There is no more space available in any of the Detail Catalog directories. From this point on, all objects on this medium will have logging switched to "No Log".

Action

The following actions can be taken:

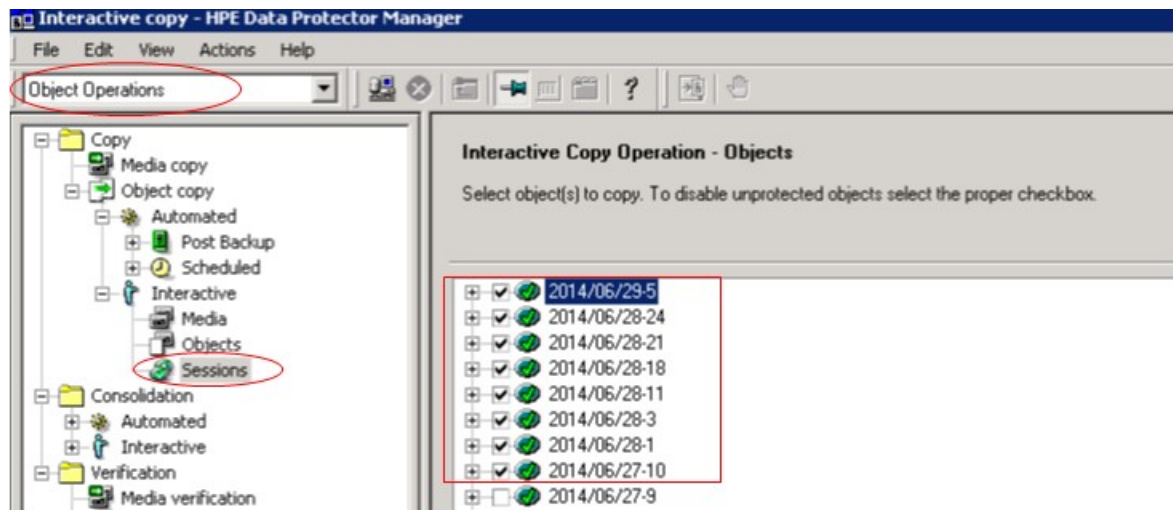
- Remove the "No Logs" option or
- Create a space on the IDB drive on the cell manager and perform a single session copy to another device. Ensure that replication option is not selected when single session copy is being performed to another StoreOnce Catalyst device.

Problem

Data consistency issues during Power On and Live Migrate operations

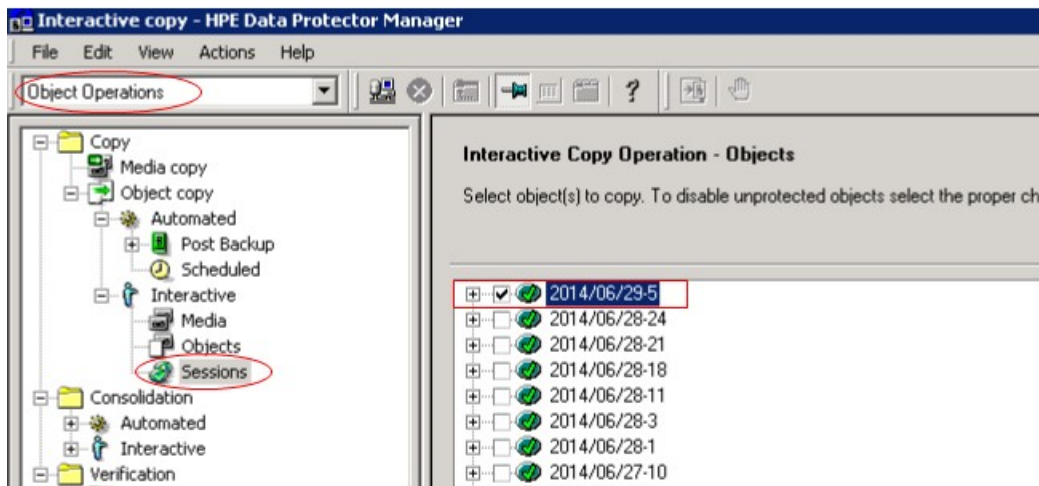
Data consistency issues during Power On and Live Migrate operations due to following reasons:

1. User has chosen a session which is a result of object copy.
2. Object copy is performed by aggregating many backup sessions into single object copy session as shown in the figure below:



Action

Choose individual session while performing object copy as shown in the figure below:



Problem

Restore of Linux virtual machine completes successfully, but ifconfig displays missing NIC

After performing restore of Linux virtual machine, the network adapters are given a new MAC address. Then the `ifconfig` command running on restored virtual machine displays missing NIC.

For more information, see https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002767

Action

To details on updating the MAC address of the Linux operating system, see the VMware KB article: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002767

Problem

During backup of virtual machine, the following error message is displayed:

```
[Warning] From: VEPALIB_VMWARE@<backup_host> <datacenter> Time: <date> <time>  
Virtual Machine <virtual_machine_name>: Download of disk descriptor file (<disk_vmdk_file_name>) failed.
```

Action

Check whether the `vpxd.enableHttpDatastoreAccess` and `vpxd.enableDebugBrowse` are enabled in vCenter configuration file, as these settings might be disabled due to security reason. For the detailed procedure, see

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2101567

Problem

Shadow VM is not attached to Nova Instance after restore

At the time of restore, if the backed up Shadow VM is not attached to the original backedup Nova instance, then the restored shadow VM is detached after the restore.

Action

To have the consistency with the Openstack dashboard, the restored shadow VM is detached during the restore process. To resolve, re-attach the shadow VM to Nova Instance from the openstack Dashboard.

Chapter A: Appendix

This appendix contains information on the following topics:

- Reconfiguring an Oracle instance for instant recovery
- ZDB integrations omnirc options

Reconfiguring an Oracle instance for instant recovery

If the control files or redo logs are located on the same volume group (if LVM is used) or source volume as the database files, the control files and online redo logs are overwritten during instant recovery. In such case, you may want to reconfigure the Oracle instance.

For details on the required configuration, see [Oracle backup set ZDB concepts , on page 22](#) and to [Oracle proxy-copy ZDB concepts , on page 26](#).

For additional examples on how to move the redo logs and control files, see [Examples for moving the control files and redo logs to different locations, on the next page](#).

Moving online redo logs

To move the *online redo log files* from the source volumes to be replicated, to other locations:

1. List the online redo log files:

```
$ sqlplus
SQL> select member from v$logfile;
```

2. Shut down the database:

```
SQL>connect user/password@service as sysdba;
SQL> shutdown
SQL> exit
```

3. Move the log files to a different location using operating system tools.

4. Start the database in mount mode:

```
$ sqlplus
SQL> connect user/password@service as sysdba;
SQL> startup mount;
```

5. Register the new locations for each moved file:

```
SQL> alter database rename file 'OldPathName' to 'NewPathName';
where OldPathName and NewPathName are full paths to the log file.
```

6. Open the database in normal mode:

```
SQL> alter database open;
```

Moving control files

To move the *control files* from the source volumes to be replicated, to other locations:

1. Determine if the database uses the SPFILE parameter:

```
SQL> show parameter SPFILE
```

2. If the database does not use SPFILE:

- a. Shut down the database.

```
SQL> shutdown
```

- b. Move the control files to a different location using operating system tools.
- c. Edit the `CONTROL_FILES` parameter in the database's initialization parameter file (usually located in the `$ORACLE_HOME/dbs/initSID.ora` directory) to change the existing control file names:

```
control_files = ("NewPathName", ...)
```

- d. Restart the database:

```
SQL> startup
```

If the database uses SPFILE:

- a. Specify the new location for control files by running the following command:

```
SQL> alter system set control_files='NewPathName1',  
'NewPathName2',..., scope=spfile
```

- b. Shut down the database.

```
SQL> shutdown
```

- c. Move the control files to a different location.

- d. Restart the database:

```
SQL> startup
```

Examples for moving the control files and redo logs to different locations

Example - moving online redo logs

In the following example for Oracle10g on HP-UX, the data files are on the same source volume as the control files and redo logs, which is `/opt/oracle/product/10.2.0`.

To move the *online redo log files* from `/opt/oracle/product/10.2.0` to `/oracle/logs` (which is not replicated):

1. List the online redo log files:

```
$ sqlplus
```

```
SQL> select member from v$logfile;
```

```
/opt/oracle/product/10.2.0/oradata/redo01.log
```

```
/opt/oracle/product/10.2.0/oradata/redo02.log
```

```
/opt/oracle/product/10.2.0/oradata/redo03.log
```

List the filenames and tablespaces to check whether they are on the same source volumes as the control files:

```
SQL> select FILE_NAME, TABLESPACE_NAME, BYTES from dba_data_files;
```

```
FILE_NAME
-----
TABLESPACE_NAME          BYTES
-----
/opt/oracle/product/10.2.0/oradata/system01.dbf
SYSTEM                   419430400
/opt/oracle/product/10.2.0/oradata/undotbs01.dbf
UNDOTBS1                 377487360
/opt/oracle/product/10.2.0/oradata/cwmlite01.dbf
CWMLITE                  20971520
```

2. Shut down the database:

```
SQL>connect user/password@service as sysdba;
SQL> shutdown
SQL> exit
```

3. Move the log files to a different location.

```
$ mv /opt/oracle/product/10.2.0/oradata/redo* /oracle/logs
```

4. Start the database in mount mode:

```
$ sqlplus
SQL> connect user/ password@service as sysdba;
SQL> startup mount;
```

5. Rename the new locations for each moved file:

```
alter database rename file '/opt/oracle/product/10.2.0/oradata/redo01.log' to
'/oracle/logs/redo01.log';
Database altered.

alter database rename file '/opt/oracle/product/10.2.0/oradata/redo02.log' to
'/oracle/logs/redo01.log';
Database altered.

alter database rename file '/opt/oracle/product/10.2.0/oradata/redo03.log' to
'/oracle/logs/redo01.log';
Database altered.
```

6. Open the database in normal mode:

```
SQL> alter database open;
```

Example - moving control files for Oracle 10g

In the following example, the Oracle 10g database uses SPFILE. To move the control files from /opt/oracle/product/10.2.0/ to /oracle/oractl:

1. Determine if the database uses the SPFILE parameter:

```
SQL> show parameter spfile;

NAME                                TYPE          VALUE
-----
spfile                               string       ?/dbs/spfile@.ora
```

2. Specify the new location for the control files by running the following command (in a single line and without the "\" characters):

```
SQL> alter system setcontrol_files='/oracle/logs/RCVCAT \
/control01.ctl', '/oracle/logs/RCVCAT/control02', '/oracle \
/logs/RCVCAT/control03.ctl' scope=spfile;
```

3. Shut down the database:

```
SQL>shutdown
```

4. Move the control files to the new location:

```
mv /opt/oracle/product/10.2.0/oradata/control* /oracle/oractl
```

5. Restart the database:

```
SQL>startup
```

ZDB integrations omnirc options

The Data Protector ZDB integrations use `omnirc` options, which can be set on both the application and backup systems. These options are used for Data Protector ZDB integrations customizing. For information on how to use the `omnirc` options, see the *HPE Data Protector Help* index: “omnirc options”.

For information on Data Protector ZDB agents `omnirc` file options, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

This section explains the `omnirc` file options that can be set for the Data Protector ZDB integrations.

ZDB_ORA_INCLUDE_CF_OLF : Data Protector Oracle Server integration and Data Protector SAP R/3 integration-related options.

NOTE:

This option is not supported on EMC.

The default value is 0. Possible values are 0 and 1.

If an offline backup is performed using the Data Protector SAP R/3 integration, the option is ignored and the integration behaves as if the option was set to 1.

Instant recovery

The instant recovery process depends on whether the control file and redo logs reside on the same disk array source volume as datafiles or not:

- If this option is set to 0 (default), during a ZDB session, Data Protector creates target volumes only for the source volumes containing Oracle datafiles. Target volumes for source volumes containing Oracle control file and Oracle online redo logs are not created.

For Oracle proxy-copy or backup set ZDB and restore concepts when this option is set to 0, see [Oracle backup set ZDB concepts](#) , on page 22 and [Oracle proxy-copy ZDB concepts](#) , on page 26. For SAP R/3 backup and restore concept when this option is set to 0, see [Integration concepts](#) , on page 18.

- If this option is set to 1, Data Protector creates target volumes for all source volumes containing Oracle datafiles, Oracle control file, and if Oracle integration is used, Oracle online redo logs.

IMPORTANT:
 If the ZDB_ORA_INCLUDE_CF_OLF option is set to 1 the control files and redo logs are overwritten during instant recovery.

Opening the database on the backup system

To successfully open the database on the backup system for *other* purposes than Data Protector,

- With Oracle proxy-copy ZDB method, set this option to 1.
- With Oracle backup set ZDB method, used with the Oracle integration, you can always open the database on the backup system.

Prerequisites

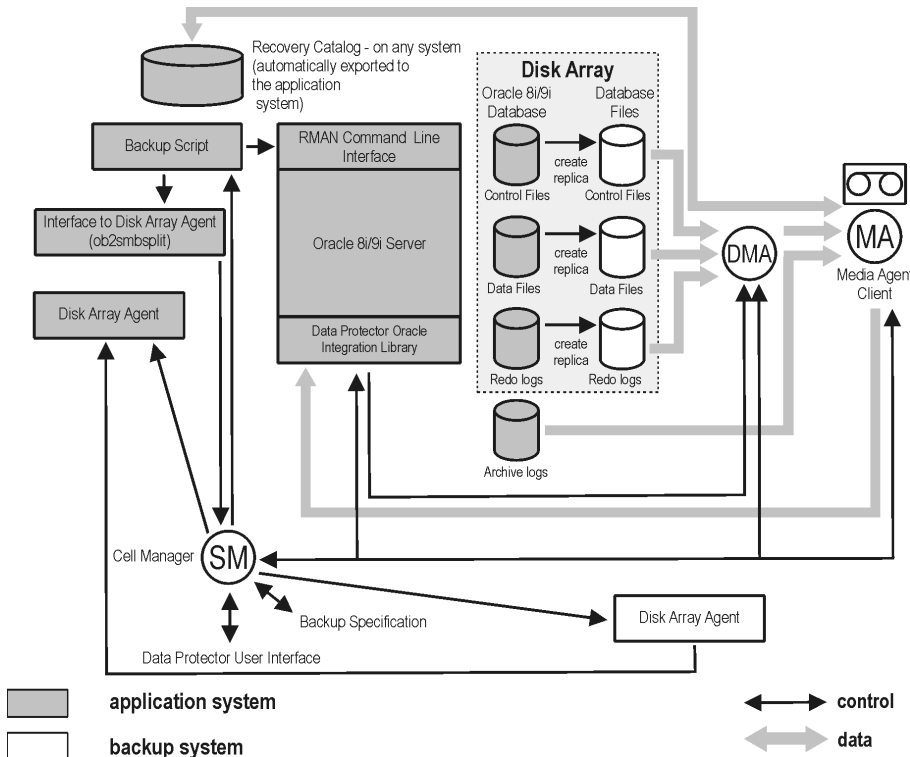
The prerequisites for this option to be set to 1 are:

- Data Protector Oracle Server integration: Oracle datafiles, Oracle control file, and Oracle online redo logs must be installed on a disk array.
- Data Protector SAP R/3 integration: Oracle datafiles and Oracle control file must be installed on a disk array.


See [Oracle proxy-copy ZDB and restore concepts when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1, below](#) and [Oracle backup set ZDB and restore concept when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1, on the next page](#) for Oracle backup and restore concepts when this option is set to 1.

See [SAP R/3 backup and restore concept when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1 with online backup, or in case of offline backup, on the next page](#) for SAP R/3 backup and restore concepts when this option is set to 1.


Oracle proxy-copy ZDB and restore concepts when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1



Oracle backup set ZDB and restore concept when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1



SAP R/3 backup and restore concept when the ZDB_ORA_INCLUDE_CF_OLF option is set to 1 with online backup, or in case of offline backup



ZDB_ORA_INCLUDE_SPF : Data Protector Oracle Server integration-related options.

The default value is 0. Possible values are 0 and 1.

The option is ignored and the integration behaves as if the option was set to 1 if offline backup is performed using the Data Protector SAP R/3 integration.

If this option is set to 0, during ZDB sessions, Data Protector checks if Oracle Server datafiles and the Oracle Server SPFILE are located on the same source volume. If the datafiles and the SPFILE are located on the same volume and instant recovery is enabled in the backup specification, the ZDB session fails. If this option is set to 1, Data Protector skips the check. To enable instant recovery, leave this option set to the default value.

CAUTION:
 If this option is set to 1 and the datafiles are located on the same volume as the SPFILE, the SPFILE is overwritten during instant recovery, potentially resulting in a data loss.

ZDB_ORA_NO_CHECKCONF_IR:Data Protector Oracle Server integration-related option.

The default value is 0. Possible values are 0 and 1.

By default, the Oracle Server configuration is checked whether it is instant recovery-enabled or not (whether the Oracle control file, the Oracle Server SPFILE, and the Oracle Server online redo logs are located on volumes of a different volume group than Oracle Server datafiles or not). For Oracle Server configuration check, the Data Protector command `omniresolve` is used internally. On UNIX systems,

the `omniresolve` file must have the `setuid` bit set. When this option is set to 1, the configuration check is omitted.

CAUTION:

Checking the Oracle Server configuration for instant recovery suitability is an essential step to ensure the instant recovery session does not result in a data loss. It is therefore not recommended to set this option to 1 unless you ensure that the Oracle Server is and remains configured appropriately for instant recovery.

OB2MARAWREAD_KB: This option sets the read block size for Oracle and SAP R/3 ZDB integrations on UNIX systems with Oracle tablespaces or datafiles installed on disk images and when using the proxy-copy method (when using DMA).

The default value is 64 kB. The specified value must be in the range between 1 kB and 1 MB.

The specified size is automatically adjusted to a size which is a multiple of the block size. The values above 256 kB could cause the DMA to fail.

ZDB_TAKE_CLUSRES_ONLINE: This option specifies how many times Data Protector tries to connect to Microsoft SQL Server in case the first connection fails. A reconnection is triggered every 30 seconds. This means that Data Protector waits up to $ZDB_TAKE_CLUSRES_ONLINE \times 30$ seconds for the Microsoft SQL Server resources to start up.

SSEA_ATOMIC_SPLIT: Determines if the Data Protector HPE P9000 XP Agent should use the atomic split configuration of a disk array of the HPE P9000 XP Disk Array Family to ensure data consistency of replicas of the Oracle Server data in configurations where Automatic Storage Management (ASM) is used.

Default: 0 (disabled). Possible: 0|1.

Send documentation feedback

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

Feedback on Zero Downtime Backup Integration Guide (Data Protector 10.02)

Add your feedback to the email and click **Send**.

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

We appreciate your feedback!