

NetBackup™ for Oracle Administrator's Guide

Release 11.0

Last updated: 2025-03-18

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

Cohesity, Veritas, the Cohesity Logo, Veritas Logo, Veritas Alta, Cohesity Alta, and NetBackup are trademarks or registered trademarks of Cohesity, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Cohesity is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Cohesity product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Cohesity, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Cohesity, Inc. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Cohesity as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Cohesity, Inc.
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Cohesity account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Cohesity website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Cohesity community site:

<http://www.veritas.com/community/>

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introduction	13
	About NetBackup for Oracle	13
	NetBackup for Oracle features	14
Chapter 2	Installing NetBackup for Oracle	18
	Planning the installation of NetBackup for Oracle	18
	Verifying the operating system and platform compatibility	19
	License for NetBackup for Oracle	20
	NetBackup server and client requirements for NetBackup for Oracle	20
	Requirements for using NetBackup for Oracle in a NetBackup cluster	21
	About linking Oracle RMAN with NetBackup for UNIX	21
	Verifying environment variables and shutting down Oracle	22
	Linking Oracle RMAN with NetBackup on UNIX platforms	23
Chapter 3	Configuring RBAC for the Oracle administrator	29
	RBAC roles for the Oracle administrator	29
Chapter 4	Managing Oracle discovery and databases	31
	About Oracle discovery	31
	Viewing the Oracle repository	32
	Configure discovery of Oracle clients	33
	Manually discover databases	33
	Manually discover pluggable databases	34
	Managing Oracle databases	34
	View and edit Oracle database details	35
	Associate an RMAN catalog with a database	36
	Remove an RMAN catalog from a database	37
	Activate or deactivate a database	37
	Remove a database or pluggable database	37
	Manually add a database	38
	Manually add an Oracle instance	39

	Edit Oracle instance details	40
	Copy a RAC instance	40
	Remove an Oracle instance	40
	Load balance Oracle RAC instances	41
	Managing RMAN catalogs	42
	Manage a Data Guard	43
Chapter 5	Managing Oracle credentials	45
	About Oracle credentials	45
	Oracle credential types in the NetBackup Credential management	46
	Add a credential for an Oracle database or RMAN catalog	48
	View the credential name that is applied to an asset	49
	Edit or delete a named credential	49
	Register credentials for an Oracle database or RMAN catalog	50
	Manually validate database credentials	51
	Configure an Oracle Wallet with RAC within NetBackup	51
Chapter 6	Configuring Oracle policies	54
	Preparing for NetBackup for Oracle configuration	54
	About Oracle policy configuration	54
	Oracle Home User permissions when NetBackup SAN Client is used	56
	Oracle backup policy types	56
	Configuring the Maximum jobs per client	58
	About Oracle Intelligent Policies (OIP)	59
	Create an Oracle Intelligent Policy (OIP)	60
	Oracle database upgrade effect on Oracle Intelligent Policies	61
	Configuring NetBackup for Oracle automatic backup schedules	62
	About NetBackup for Oracle schedule properties using Oracle Intelligent Policy	62
	Oracle Intelligent Policy - Storage and Retention	64
	About Oracle Intelligent Policy primary server behavior	66
	Instances and databases tab	66
	Backup selections tab	67
	Oracle tab	68
	Create a Data Guard policy	74
	Create <code>bpstart_notify</code> and <code>bpend_notify</code> scripts for an Oracle Intelligent Policy (OIP)	75
	Logging the RMAN input and output on a client	76

About script-based Oracle policies	77
Adding a new script-based Oracle policy	78
About policy attributes	79
Schedule properties for Oracle script-based policies	80
Script-based policy - Storage and Retention	81
Add clients to a policy	84
About adding backup selections to an Oracle policy	84
About configuring the run-time environment	86
About creating shell scripts	93
Configuring the logon account for the NetBackup Client Service for NetBackup for Oracle	97
Testing configuration settings for NetBackup for Oracle	98

Chapter 7 Performing backups and restores of Oracle 100

Overview of using NetBackup for Oracle	100
Maintaining the RMAN repository	101
Querying the RMAN repository	105
About NetBackup for Oracle backups	106
Running the NetBackup for Oracle shell script	107
Running RMAN	108
Browsing backups using the bplist command	108
Managing expired backup images	109
About preventing the direct expiration of backup images	110
About NetBackup for Oracle restores	111
About the Oracle recovery shell script on the client	111
About catalog names with RAC and Data Guard aware intelligent policies	112
Running RMAN on the client	113
About Oracle multistream restore for proxy backup	114
Redirecting a restore to a different client	115
Using NetBackup for Oracle in a Windows Server Failover Cluster (WSFC)	121
About backups of an Oracle clustered database on Windows	121
Bringing the database instance offline on Windows	122
Bringing the database instance online on Windows	123
User-directed backup or restore from the Windows client	123

Chapter 8 Oracle cloning 125

About cloning an Oracle database or a pluggable database	125
Create a custom RBAC role for Oracle cloning	128
Clone an Oracle database	130

	Clone a pluggable database	131
	Cloning a database from an Oracle Intelligent Policy (OIP) that is Data Guard aware	133
	Additional functionality in NetBackup APIs	133
Chapter 9	NetBackup Copilot for Oracle	135
	About Oracle Copilot	135
	Configuring an OIP using a share on the NetBackup appliance (Oracle Copilot)	137
	Creating an instant recovery point from an Oracle Copilot image (NetBackup Appliance share)	139
	Deleting an instant recovery point for Oracle Copilot instant recovery	142
	Configuring an OIP using universal shares (Oracle Copilot)	143
	Recommendations when you configure an Oracle Copilot with universal shares	146
	Troubleshooting issues related to Oracle Copilot with universal shares	146
	Managing an instant access mount from an Oracle Copilot universal share image	149
	Cleaning up the Oracle Copilot share after point in time restore of database	149
	Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount	156
	About restoring from a data file copy to ASM storage using RMAN	160
Chapter 10	Oracle Copilot with instant access	162
	Prerequisites when you configure an instant access Oracle database	162
	Hardware configuration requirement of instant access	164
	Things to consider before you configure an instant access mount point	164
	Backing up an Oracle database using Oracle Copilot policy with a universal share	165
	Configure an instant access mount	165
	View the livemount details of an instant access mount	166
	Configuring Auto Image Replication for Oracle instant access backups	168
	Delete an instant access mount	169
	NetBackup for Oracle terms	169
	Frequently asked questions	171

Chapter 11	NetBackup for Oracle with Snapshot Client	174
	About NetBackup for Oracle with Snapshot Client	174
	Proxy copy	176
	NetBackup for Oracle stream-based operations	177
	NetBackup for Oracle file-based operations	178
	How NetBackup for Oracle with Snapshot Client works	179
	About the NetBackup for Oracle backup and restore operations	180
	Database objects supported by advanced backup methods	180
	About NetBackup multistreaming	181
	RMAN multiple channels	181
	Restoring data files to a new location	182
	Redirecting a restore to a different client	182
	Symbolic links and raw data files (UNIX)	183
	Quick I/O data files (UNIX)	183
	RMAN incremental backups	183
	Proxy backup examples	184
	About configuring Snapshot Client with NetBackup for Oracle	187
	Configuration requirements for snapshot backups with NetBackup for Oracle	187
	Configuring a snapshot policy for NetBackup for Oracle	188
	Configuring a snapshot policy using a share on the NetBackup appliance (Oracle Copilot)	191
	Restoring NetBackup for Oracle from a snapshot backup	193
	About restoring individual files from a NetBackup for Oracle snapshot backup	193
	About NetBackup for Oracle restores of volumes and file systems using snapshot rollback	194
	About configuring NetBackup for Oracle block-level incremental backups on UNIX	196
	How BLI works with NetBackup for Oracle (UNIX)	197
	About the Storage Checkpoint facility and NetBackup for Oracle	198
	Configuration requirements for BLI backups with NetBackup for Oracle	198
	Configuring policies for BLI backups with NetBackup for Oracle	199
	About Snapshot Client effects	200
	How Snapshot Client software affects backup types	201
	How Snapshot Client software affects schedule properties	201
	How Snapshot Client software affects scripts	202
	Oracle with Snapshot Client environment variables	202

About Oracle support for Replication Director	204
Configuring an Oracle Intelligent Policy using Replication Director	205
Configure a script-based Oracle policy	209
Chapter 12 NetBackup Dedupe Direct for Oracle	212
About NetBackup Dedupe Direct for Oracle plug-in	212
Verifying the operating system and platform compatibility	213
Before you install NetBackup Dedupe Direct for Oracle plug-in	213
Installing NetBackup Dedupe Direct for Oracle plug-in on the Oracle database server	214
Using the NetBackup Dedupe Direct for Oracle plug-in	215
Configuring the NetBackup Dedupe Direct for Oracle plug-in	216
NetBackup Dedupe Direct for Oracle rman.cfg file parameters	216
Configuring External Certificate Authority (ECA)	218
ECA msdp_app.cfg file parameters	219
Configuring Oracle RAC	220
Disabling Oracle stream handler	220
Performing the RMAN backup directly to MSDP server	221
Performing the RMAN restore directly from the MSDP storage server	222
Creating an SLP to import the Oracle backups to NetBackup on Flex WORM	222
Uninstalling the NetBackup Dedupe Direct for Oracle plug-in	223
Chapter 13 Troubleshooting	224
About troubleshooting NetBackup for Oracle	225
About NetBackup for Oracle troubleshooting steps	225
NetBackup debug logs and reports	227
Enabling the debug logs manually (Windows) (NetBackup for Oracle)	228
Enabling the debug logs manually (UNIX) (NetBackup for Oracle)	229
About the NetBackup for Oracle log files	230
Setting the debug level on a Windows client	232
Setting the debug level on a UNIX client	232
About RMAN utility logs	233
Troubleshooting RMAN backup or restore errors	233
Verifying the RMAN script on UNIX	233
Troubleshooting each stage of the backup or restore	234
Troubleshooting NetBackup for Oracle with Snapshot Client	236

	Error: Unable to re-create online log	237
	Minimizing time-out failures on large database restores	237
	Minimizing the loading and unloading of tapes for database backups	238
	Enabling the debug logs manually for NetBackup Dedupe Direct for Oracle plug-in	239
Appendix A	Deduplication best practices	240
	Optimizing and deduplicating stream-based and proxy copy Oracle backups	240
	Oracle stream handler	242
	Configuring a stream-based Oracle backup	245
	Example RMAN script for a stream-based backup	247
	Editing the RMAN script and configuring NetBackup for Oracle for a proxy copy backup	248
	Example RMAN script for a proxy copy backup	249
Appendix B	Snapshot Client support of SFRAC	252
	About Snapshot Client support of SFRAC	252
	NetBackup configuration for an SFRAC environment	253
	Configuring the SFRAC environment for a backup operation	253
	Performing a rollback restore in an SFRAC environment	255
	Troubleshooting NetBackup in an SFRAC environment	256
Appendix C	Script-based Block-Level Incremental (BLI) Backups without RMAN on UNIX and Linux systems	259
	About script-based Block-Level Incremental (BLI) Backups without RMAN	259
	About BLI backup and restore operations	260
	Verifying installation requirements for BLI backups without RMAN	260
	File system and Storage Checkpoint space management	261
	Creating NetBackup policies for script-based BLI backup	263
	Number of policies required for BLI backup	263
	About BLI policy attributes	265
	About the BLI client list	266
	Backup selections list for BLI backups	266
	About schedules for BLI backup policies	267
	Example Oracle BLI backup policy	268
	Setting the Maximum jobs per client global attribute	269

About BLI backup methods	270
Creating notify scripts for BLI backups	271
Performing backups and restores	276
About NetBackup for Oracle agent automatic backups	276
About NetBackup for Oracle manual backups	277
Backing up Quick I/O files	278
Restoring BLI backup images	279
About NetBackup backup and restore logs	280
About troubleshooting backup or restore errors	281
Troubleshooting stages of backup and restore operations	281
NetBackup restore and backup status codes	282
Improving NetBackup backup performance	284
About BLI backup and database recovery	285

Appendix D	XML Archiver	286
	NetBackup for Oracle XML export and XML import	286
	NetBackup for Oracle XML export and import archiving features	287
	XML export archive process	288
	Sequence of operation: XML export archive	289
	XML import restore process	291
	Sequence of operation: XML import restore	292
	About the environment variables set by a user in the XML export	
	parameter file	294
	About XML export shell scripts	295
	Creating XML export scripts manually	295
	Performing an XML export archive	297
	Running the NetBackup for Oracle XML export script on the client	297
	Running bporaexp on the client as an Oracle user	298
	Writing to a directory versus writing to a storage unit	299
	About bporaexp parameters	300
	Browsing XML export archives using bporaimp parameters	303
	Browsing XML export archives using bplist	304
	Restoring an XML export archive	306
	Running an XML import script on the client	306
	Running bporaimp on the client in NetBackup for Oracle	307
	About bporaimp parameters	307
	About redirecting a restore of an XML export archive to a different	
	client	311
	Troubleshooting XML export or XML import errors	313
	Checking the logs to determine the source of an error	314

	Troubleshooting each stage of the XML export or XML import	315
	Additional XML export and import logs	318
Appendix E	Register authorized locations	319
	Registering authorized locations used by a NetBackup database script-based policy	319

Introduction

This chapter includes the following topics:


- [About NetBackup for Oracle](#)
- [NetBackup for Oracle features](#)

About NetBackup for Oracle

NetBackup integrates the database backup and recovery capabilities of the Oracle Recovery Manager (RMAN) with the backup and recovery management capabilities of NetBackup.

[Figure 1-1](#) shows the major components in a NetBackup configuration. The server that hosts the Oracle database must be a NetBackup client. The primary server must also have NetBackup for Oracle licensed.

Figure 1-1 NetBackup for Oracle on a sample network



NetBackup for Oracle features

The following table shows the NetBackup for Oracle main features.

Table 1-1 NetBackup for Oracle features

Feature	Description
Integration with NetBackup role-based access control (RBAC)	RBAC allows the NetBackup administrator to delegate the management of Oracle operations and of Oracle assets within NetBackup to specific DBAs, which decreases the configuration burden for the NetBackup administrator. See "RBAC roles for the Oracle administrator" on page 29.

Table 1-1 NetBackup for Oracle features (*continued*)

Feature	Description
Oracle database management	NetBackup automatically discovers Oracle databases in the environment. Manual discovery can also be performed. When Oracle DBAs are assigned the Default Oracle Administrator RBAC role, they can use NetBackup web UI to manage Oracle assets and their credentials. After databases are registered with credentials, an Oracle DBA can create a policy to protect Oracle assets.
Oracle CMS credentials	NetBackup supports named credentials from both the internal and the external providers for use with Oracle assets. See “Oracle credential types in the NetBackup Credential management” on page 46.
Oracle Intelligent Policies	The Oracle Intelligent Policy (OIP) protects multiple Oracle databases and Data Guard databases that are spread over multiple clients. You select databases for a policy from a list of databases that are automatically discovered in the NetBackup environment. OIP does not require you to know how RMAN functions or how to use scripts. The feature is database-based and not scripting based.
Script-based policies	Script-based policies are available as an alternative to OIP. If you use the command-line interface or scripts, you must use script-based Oracle policies. These policies use Oracle's Recovery Manager (RMAN) as if NetBackup were not present.
Scheduling facilities	NetBackup scheduling facilities on the primary server can be used to schedule automatic and unattended Oracle backups. This feature also lets you choose the times when these operations can occur. For example, to prevent interference with normal daytime operations, you can schedule your database backups to occur only at night.
Transparent Oracle and regular file system backup and restore operations	All backups and restores run simultaneously and transparently without any action from the NetBackup administrator. The database administrator, administrator, or any authorized user can run database backup and restore operations through NetBackup.
Centralized and networked backup operations	From the NetBackup primary server, you can schedule database backups or start them manually for any client or database. The Oracle databases can also reside on the hosts that are different from the devices on which NetBackup stores the backups.
Recovery points	The web UI provides recovery points from which you can easily create a clone from a database or pluggable database. If instant access is configured, you can also create an instant access database from a recovery point.
Compression	Compression increases backup performance over the network and reduces the size of the backup image that NetBackup writes to the storage unit.
Multiplexed backups and restores	NetBackup for Oracle lets you take advantage of NetBackup's multiplexing capabilities. Multiplexing directs multiple data streams to one backup device, thereby reducing the time necessary to complete the operation.

Table 1-1 NetBackup for Oracle features (*continued*)

Feature	Description
Parallel backup and restore operations	NetBackup for Oracle supports the parallel backup and restore RMAN capabilities. For example, you can run more than one tape device at a time for a single Oracle backup or restore. This capability reduces the time necessary to complete the operation.
Immediate backup	An Oracle DBA can perform an immediate backup of a database using the <code>nboraacmd</code> command. See the NetBackup Commands Guide for details.
Support for Container and Pluggable databases	Oracle 12c introduced the container databases (CDB) and pluggable databases (PDB). The Oracle Intelligent Policy is enhanced and allows a backup to include single or multiple PDBs. This feature also lets you select one or more Oracle 12c instances along with non-Oracle 12c instances in OIP.
Support for Oracle Real Application Clusters (RAC)	The NetBackup web UI provides full support for protection of Oracle Real Application Clusters (RAC) with OIP and clients at version NetBackup 8.3 or later. NetBackup can be set up to load balance the Oracle instances. Use this feature to distribute the backup load across all of the instances and to exclude specific Oracle instances from backups.
Support for Data Guard	NetBackup has full support for the primary and the standby databases that are members of a Data Guard configuration. This feature enables a unified image cataloging system across any type of intelligent policy that protects the Data Guard configuration. A primary or a standby preference can also be configured in the policy. See “Manage a Data Guard” on page 43.
Sharing the same storage units that are used for other file backups	You can share the same devices and media that are used for other backups or give Oracle exclusive use of certain devices and media. NetBackup for Oracle can use the Media Manager, disk, or Media Server Deduplication Pool (MSDP) storage units.
NetBackup Dedupe Direct for Oracle	NetBackup Dedupe Direct for Oracle offers a lightweight plug-in for RMAN backups to MSDP storage. The Oracle database administrator can control the whole protection and life-cycle stages without the NetBackup client. It also enables client-side deduplication to minimize network traffic and improves overall backup speed.

Table 1-1 NetBackup for Oracle features (*continued*)

Feature	Description
Support on a NetBackup appliance for backup to an appliance share (NetBackup Copilot for Oracle)	Note: This feature requires a NetBackup appliance that runs software version 2.7.1 or later. Oracle Copilot enhances the Oracle Intelligent Policy by giving you options for protecting an Oracle database using a share on a NetBackup appliance. This feature gives you better control of backups when an Oracle database backup is placed in a database share by the DBA. This feature also lets you choose a database share as the destination for the first backup copy. The backup copy is a full set of database data file copies created, incrementally updated, and protected by NetBackup. You must create a share on the appliance for this option using the NetBackup Appliance Shell Menu. Oracle backups only work on an NFS share on the NetBackup appliance. For more information about how to set up the share, see “Creating a share from the NetBackup Appliance Shell Menu” in the NetBackup 52xx and 5330 Appliance Administrator's Guide.
Delegate management of RMAN catalogs	NetBackup lets you delegate permission and management of RMAN catalogs within NetBackup, which can simplify how you manage credentials and configuration.
Support for Replication Director	Replication Director can be used to create snapshots of the Oracle database. The snapshots can then be replicated to other NetApp disk arrays or backup the snapshot to a storage unit. To use Replication Director, the Oracle database must exist on a NetApp NAS disk array. Replication Director is not supported on SAN storage at this time. Oracle snapshot backups that use Replication Director are supported on UNIX and Linux platforms only.
NetBackup APIs	The same functionality in the web UI is also available with the NetBackup APIs.

Installing NetBackup for Oracle

This chapter includes the following topics:

- [Planning the installation of NetBackup for Oracle](#)
- [Verifying the operating system and platform compatibility](#)
- [License for NetBackup for Oracle](#)
- [NetBackup server and client requirements for NetBackup for Oracle](#)
- [Requirements for using NetBackup for Oracle in a NetBackup cluster](#)
- [About linking Oracle RMAN with NetBackup for UNIX](#)

Planning the installation of NetBackup for Oracle

To use the new features that are included in NetBackup for Oracle in NetBackup 11.0, upgrade your NetBackup for Oracle clients to NetBackup 11.0. The NetBackup media server must use the same version as or a higher version than the NetBackup for Oracle client.

[Table 2-1](#) shows the installation steps that are required to run NetBackup for Oracle.

Table 2-1 Installation steps for NetBackup for Oracle

Step	Action	Description
Step 1	Verify the operating system and platform compatibility.	See the NetBackup Compatibility Lists .

Table 2-1 Installation steps for NetBackup for Oracle *(continued)*

Step	Action	Description
Step 2	Verify that primary server has a valid license for NetBackup for Oracle and any NetBackup options or add-ons that you want to use.	See "License for NetBackup for Oracle" on page 20.
Step 3	Install the NetBackup client software on the computers that have the databases that you want to back up.	See "NetBackup server and client requirements for NetBackup for Oracle" on page 20. For an Oracle RAC cluster, you must install the client on all instances of a RAC cluster.
Step 4	To use NetBackup for Oracle in a NetBackup cluster, verify that your cluster environment is supported and that the NetBackup cluster is configured correctly.	See "Requirements for using NetBackup for Oracle in a NetBackup cluster " on page 21.
Step 5	Link the Oracle database server software with the NetBackup API library installed by NetBackup for Oracle.	See "About linking Oracle RMAN with NetBackup for UNIX" on page 21.

Verifying the operating system and platform compatibility

Verify that the NetBackup for Oracle agent is supported on your operating system or platform.

To verify operating system and compatibility

- 1 Go to the NetBackup compatibility list site.
<http://www.netbackup.com/compatibility>
- 2 Click on the following document:
Application/Database Agent Compatibility List
- 3 For information on support for Snapshot Client, see the Snapshot Client section in the following document:
Hardware and Cloud Storage Compatibility List (HCL).

License for NetBackup for Oracle

The NetBackup for Oracle agent is installed with the NetBackup client software. No separate installation is required. A valid license for the agent must exist on the primary server.

More information is available on how to add licenses.

See the [NetBackup Web UI Administrator's Guide](#).

For a NetBackup cluster, a valid license for NetBackup for Oracle must exist on each node where NetBackup server resides.

NetBackup server and client requirements for NetBackup for Oracle

Every NetBackup server includes the NetBackup client software by default. Therefore, you can use NetBackup for Oracle on a NetBackup server or client (if NetBackup for Oracle is supported on that platform).

Verify that the following requirements are met for the NetBackup server:

- The NetBackup server software is installed and operational on the NetBackup server. The NetBackup server platform can be any that NetBackup supports. See the [NetBackup Installation Guide](#).
- One or more Oracle database instances must exist.
- Make sure that you configure any backup media that the storage unit uses. The number of media volumes that are required depends on several things:
 - The devices used
 - The sizes of the databases that you want to back up
 - The amount of data that you want to archive
 - The size of your backups
 - The frequency of backups or archivesSee the [NetBackup Administrator's Guide, Volume I](#).
- Verify that the NetBackup client software is installed on the computer that has the databases you want to back up.

For a Oracle RAC cluster that is used in backups, all nodes must run the NetBackup client and the same NetBackup version. For Oracle RAC OIP, the clients must have NetBackup 8.3 or later.

If the database is clustered, you must use the same version of NetBackup on each node in the cluster.

See [“License for NetBackup for Oracle”](#) on page 20.

Requirements for using NetBackup for Oracle in a NetBackup cluster

If you plan to use NetBackup for Oracle on a NetBackup server configured in a NetBackup cluster, verify the following requirements:

- NetBackup supports your cluster environment.
See the [Software Compatibility List \(SCL\)](#).
- The NetBackup server software is installed and configured to work in a NetBackup cluster.
See the [NetBackup Installation Guide](#).
See the [NetBackup Clustered Primary Server Administrator's Guide](#).
- The NetBackup client software is installed and operational on each node to which NetBackup can failover.
- A valid license for NetBackup for Oracle must exist on each node where NetBackup server resides.

About linking Oracle RMAN with NetBackup for UNIX

Before writing to a storage unit, link the Oracle database server software with the NetBackup API library installed by NetBackup for Oracle. Oracle uses this library when it needs to write to or read from the devices that NetBackup media manager supports.

To link your Oracle software with the NetBackup API library, use one of the following methods:

- (Recommended) Use the SBT_LIBRARY parameter in the PARMS section of the allocate channel in the RMAN script. In the run block of the RMAN script, modify the ALLOCATE statement so that the SBT_LIBRARY parameter points to the NetBackup API library. For example:

```
ALLOCATE CHANNEL CH00 TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=  
/usr/opensv/netbackup/bin/libobk.so64.1';
```

- Use the linking script that NetBackup provides.

- Create the links manually.

The linking process differs depending on your hardware platform, your Oracle database server release level, and your OS level. This topic does not address all the supported combinations, but it specifies OS level differences.

The default location for the NetBackup API library is `/usr/opensv/netbackup/bin`. The name of the NetBackup API library differs depending on your platform.

[Table 2-2](#) lists the library names for the supported platforms.

Table 2-2 NetBackup API libraries

Platform	Oracle	Library name
AIX	64-bit	libobk.a64
HP Itanium	64-bit	libobk.so
Linux x86	64-bit	libobk.so64
IBM pSeries	64-bit	libobk.so
IBM zSeries	64-bit	libobk.so
Solaris (SPARC)	64-bit	libobk.so.64.1
Solaris (x86)	64-bit	libobk.so.1

Verifying environment variables and shutting down Oracle

The following procedure describes how to correctly define your environment variables and how to shut down the Oracle database instances.

To verify environment variables and shutdown Oracle database instances

- 1 Make sure that your Oracle environment variables are defined.

Define the variables as follows:

`ORACLE_HOME` The directory path to the Oracle software location.

`ORACLE_SID` The name of the Oracle database instance.

- 2 Become the Oracle user.

```
su - oracle
```

- 3 Determine if you need to link or re-link the library with NetBackup.
See [“Linking Oracle RMAN with NetBackup on UNIX platforms”](#) on page 23.
- 4 If this installation is an upgrade and the SBT_LIBRARY parameter is not in use, restart the Oracle database instance.
See [“Bringing the database instance offline on Windows”](#) on page 122.
See [“Bringing the database instance online on Windows”](#) on page 123.

Linking Oracle RMAN with NetBackup on UNIX platforms

The procedures in this topic show how to link RMAN with NetBackup. The automatic method is preferred. Use the manual method only if the link script fails or if you receive Oracle messages to indicate that manual linking is required.

Link the library with NetBackup when you license NetBackup for Oracle for the first time.

For more information about what Oracle database release that NetBackup for Oracle supports, review the [Application/Database Agent Compatibility List](#).

To automatically link Oracle RMAN with NetBackup

- 1 Run the `oracle_link` script that is located in `/usr/opensv/netbackup/bin/`.

This script determines the Oracle version level and then links Oracle with the NetBackup API library. This script writes output to `/tmp/make_trace.<pid>`. To change the trace file location, change the `MAKE_TRACE` variable in the `oracle_link` script.

- 2 If this installation is an upgrade, restart the Oracle database instance.

To manually link Oracle RMAN with NetBackup, follow the instructions in one of the platform-specific sub-topics:

- See [“Manually linking AIX \(64-bit\)”](#) on page 24.
- See [“Manually linking HP Itanium \(64-bit\)”](#) on page 24.
- See [“Manually linking Linux x86 \(64-bit\)”](#) on page 25.
- See [“Manually linking IBM pSeries or zSeries”](#) on page 26.
- See [“Manually linking Solaris x86 \(64-bit\)”](#) on page 27.
- See [“Manually linking Solaris SPARC \(64-bit\)”](#) on page 27.

Manually linking AIX (64-bit)

To manually link AIX (64-bit)

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.*
```

- 3 (Conditional) Use the `mv` command to move the Oracle library to an alternate location.

Perform this step if the output from step 2 shows that `libobk.a` exists.

For example:

```
mv libobk.a libobk.a.orig
```

- 4 Type the following `ln` command to create a new link:

```
ln -s /usr/openv/netbackup/bin/libobk.a64 libobk.a
```

- 5 If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Manually linking HP Itanium (64-bit)

To manually link HP Itanium (64-bit)

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.so
```

- 3 (Conditional) Use the `mv` command to move the Oracle library to an alternate location.

Perform this step if the output from step 2 shows that `libobk.so`.

For example:

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

- 4 Type the following `ln` command to create new links:

```
ln -s /usr/opensv/netbackup/bin/libobk.so libobk.so
```

- 5 If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Manually linking Linux x86 (64-bit)

To manually Linux x86 (64-bit)

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.so
```

- 3 (Conditional) Use the `mv` command to move the Oracle library to an alternate location.

Perform this step if the output from step 2 shows that `libobk.so` is present.

For example:

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

- 4 Type the following `ln` command to create a new link:

```
ln -s /usr/opensv/netbackup/bin/libobk.so64 libobk.so
```

- 5 If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Manually linking IBM pSeries or zSeries

To link manually IBM pSeries or zSeries

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.so
```

- 3 (Conditional) Use the `mv` command to move the Oracle library to an alternate location.

Perform this step if the output from step 2 shows that `libobk.so` is present.

For example:

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

- 4 Type the following `ln` command to create a new link:

```
ln -s /usr/opensv/netbackup/bin/libobk.so libobk.so
```

- 5 If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Manually linking Solaris x86 (64-bit)

To manually link Solaris x86 (64-bit)

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.so
```

- 3 Use the `mv` command to move the Oracle library to an alternate location. Perform this step if the output from step 2 shows that `libobk.so` is present.

For example:

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

- 4 Type the following `ln` command to create a new link:

```
ln -s /usr/opensv/netbackup/bin/libobk.so.1 libobk.so
```

- 5 If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Manually linking Solaris SPARC (64-bit)

To manually link Solaris (64-bit)

- 1 Type the following `cd` command to change directories:

```
cd $ORACLE_HOME/lib
```

- 2 Type the following `ls` command to determine whether the Oracle library exists:

```
ls -l libobk.so
```

- 3** (Conditional) Use the `mv` command to move the Oracle library to an alternate location.

Perform this step if the output from step 2 shows that `libobk.so` is present.

For example:

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

- 4** Type the following `ln` command to create a new link:

```
ln -s /usr/opensv/netbackup/bin/libobk.so64.1 libobk.so
```

- 5** If you run into problems and cannot re-link Oracle with the NetBackup API library, you can rollback what you have done. Enter the following:

```
cd $ORACLE_HOME/lib64
mv libobk.so.orig libobk.so
```

Configuring RBAC for the Oracle administrator

This chapter includes the following topics:

- [RBAC roles for the Oracle administrator](#)

RBAC roles for the Oracle administrator

To protect and restore Oracle assets, workload administrators must have permissions to access those assets, and to assign credentials to those assets. The RBAC role named Default Oracle Administrator provides the permissions that an Oracle administrator needs. These permissions include all Oracle assets, credentials, storage, and other applicable permissions. Or, that user must have the equivalent permissions in a custom RBAC role.

In addition, you may need other custom roles to give additional access to your Oracle administrators. For example, you may need a role that gives specific privileges to clone a source Oracle database or a pluggable database. This way the user doesn't have access to all the available Oracle backup and recovery operations.

For some environments, you may want to limit the permissions that some RBAC roles have for specific Oracle assets. You can apply specific permissions to the following types of Oracle assets: databases, RMAN catalogs, and Oracle credentials.

See [“Create a custom RBAC role for Oracle cloning”](#) on page 128.

Note the following requirements for creating roles and credentials:

- To create an RBAC role, you must have the RBAC Administrator role or the permissions to create roles.

- To create a credential, you must have the **Default Oracle Administrator** role or a similar role that has permissions to create credentials. The **Default Oracle Administrator** can manage all credentials.
- Contact your NetBackup administrator for assistance with creating roles and credentials.

For information on the RBAC permissions and default roles, see the NetBackup API documentation at <http://sort.veritas.com/>.

Managing Oracle discovery and databases

This chapter includes the following topics:

- [About Oracle discovery](#)
- [Managing Oracle databases](#)
- [Managing RMAN catalogs](#)
- [Manage a Data Guard](#)

About Oracle discovery

In the NetBackup web UI, the **Databases** tab in the **Workloads > Oracle** node contains the databases that NetBackup discovers and any databases that are added manually.

The NetBackup Discovery Service (`nbdisco`) discovers Oracle database instances throughout the NetBackup environment and collects the databases in a database repository. The user can access this repository with the NetBackup web UI, the `nboracmd` command, or the NetBackup APIs.

The Discovery Service reports to the primary server when it finds databases. It polls the clients upon NetBackup installation and periodically after installation (every 8 hours). By default, this service is enabled to report databases. You can also configure discovery of clients. The polling interval can be configured with the setting `ORACLE_ASSET_INTERVAL` on the primary server. On a periodic basis as a part of standard image cleanup activities, NetBackup removes any unregistered databases that were discovered more than 30 days ago.

See [“Configure discovery of Oracle clients”](#) on page 33.

See the [NetBackup Administrator's Guide, Volume I](#).

The NetBackup Discovery Service searches for databases in different areas where Oracle is installed. The following areas are where the Discovery Service searches:

- Non-RAC databases are discovered by searching the `oratab` file on UNIX and from the registry on Windows.
- NetBackup looks for the Oracle health check files that are found in the Oracle home. These are not cleaned up when a database is deleted. You may need to delete them manually otherwise NetBackup can continue to find the databases that are deleted.
- Oracle RAC databases are discovered when NetBackup queries the Oracle Cluster Ready Services (CRS) using the Oracle Clusterware high availability API.

Oracle RAC in the web UI does not support upgrades from legacy script-based policies. Also, there is no web UI support for the configurations that are created using *Appendix A: Deduplication best practices* or *Appendix B: Snapshot Client support of SFRAC*.

To allow the NetBackup web UI to discover a RAC instance or cluster:

- Remove the Oracle RAC from any configuration that is setup using *Appendix A: Deduplication best practices* or *Appendix B: Snapshot Client support of SFRAC*.
- Remove any Oracle RAC from any existing OIP policies.

Note: When an Oracle RAC database is discovered, that database does not have a **Database ID**. A **Database ID** is required to manually add additional RAC instances to the database. You must register the RAC database and provide a **Database ID** before adding additional instances.

See [“Register credentials for an Oracle database or RMAN catalog”](#) on page 50.

See [“Manually add a database”](#) on page 38.

Viewing the Oracle repository

You can view a complete list of all the Oracle databases that have been discovered in the NetBackup environment. In the web UI, click **Workloads > Oracle**. The following information is available:

- Databases
See [“View and edit Oracle database details”](#) on page 35.
- Pluggable databases

On this tab, the `PDB$SEED` is not visible. You can search the pluggable database list to view any of the `PDB$SEED` pluggable databases.

After an upgrade, be sure to run “Complete discovery” to make the PDB databases visible. Or, run an OIP backup. See the following tech note for details.

<http://www.veritas.com/docs/100061222>

See “[View and edit Oracle database details](#)” on page 35.

- RMAN catalogs
- Data Guard
- Instant access databases

See “[View the livemount details of an instant access mount](#)” on page 166.

Configure discovery of Oracle clients

By default, this service is enabled to report instances. However, you can use the `REPORT_CLIENT_DISCOVERIES` client configuration entry to shut down or restart the service on a particular client. By default, `REPORT_CLIENT_DISCOVERIES` is not present in the Windows registry or the UNIX `bp.conf` file.

To change the default setting, use `bpsetconfig` to add or change the entry:

- In the Windows registry.
- In the `/usr/opensv/netbackup/bp.conf` file on UNIX.

Use the following format: `REPORT_CLIENT_DISCOVERIES = TRUE | FALSE`

Set `REPORT_CLIENT_DISCOVERIES` to `FALSE` to shut down the discovery service.

The service shuts down within 10 minutes and remains down on the client. To turn on the discovery service on that client, set `REPORT_CLIENT_DISCOVERIES` to `TRUE` or remove the entire entry. Then run `bp.start_all` on the client to restart the service.

To set this value on a client remotely, run the following command from the primary server:

```
echo REPORT_CLIENT_DISCOVERIES=FALSE | bpsetconfig -h clientname
```

Manually discover databases

You can manually start the NetBackup discovery process if you want to immediately discover databases in the environment.

To manually discover databases

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Click **Discover**.
- 3 Indicate the hosts for which you want to perform the discovery.
 - In the **Hosts** box type one or more host names, separated by commas.
 - Select one or more host names from the list.
 - For NetBackup 8.1 and earlier clients, the host name does not automatically appear in the list. However, you can enter the host name and click **Discover** to discover that back-level host.
- 4 Click **Discover**.

Manually discover pluggable databases

You can manually start the NetBackup discovery process if you want to immediately discover pluggable databases for a specific Oracle host.

To manually discover pluggable databases

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Select the database that is associated with pluggable database.
- 3 Click **Actions > Discover pluggable database**.
- 4 After discovery completes, click **Close**.

Managing Oracle databases

All databases that you want backed up as part of an Oracle Intelligent Policy must be registered with credentials. You can apply a credential to one or more databases. The Oracle database user is required to have a certain level of credentials. It is recommended that the Oracle database user have `SYSDBA` privileges to limit NetBackup's exposure to your database. The Oracle database user can also have `SYSDBA` privileges (based on the version of Oracle).

You can view and manage Oracle databases in NetBackup as follows:

- See [“View and edit Oracle database details”](#) on page 35.
- See [“Manually discover databases”](#) on page 33.
- See [“Activate or deactivate a database”](#) on page 37.
- See [“Remove a database or pluggable database”](#) on page 37.

View and edit Oracle database details

NetBackup provides the capability to view the details for Oracle databases and pluggable databases.

NetBackup discovers pluggable databases when you validate the credentials for the database. Or, to manually discover the databases, click **Discover pluggable databases**. NetBackup clients at version 10.3 or later are automatically discovered the next time a backup runs. If you do not click “Complete discovery” following a primary server upgrade, you are prevented from browsing for pluggable databases to add to them to existing or to new policies. This behavior does not affect the databases in any existing policies or any current backups.

Oracle and NetBackup administrators can select a database and review the following information about a database or pluggable database:

- Protection. List the policies that protect the database.
- (Databases tab) Instances. Lists any instances that are associated with the database.
- (Databases tab) Pluggable databases. Lists any pluggable databases that are associated with the database. The PDB seed name appears in this table.
- Recovery points. Lists the recovery points that are available for the database. From a recovery point you can clone a database and configure an instant access mount.
- Restore activity. Displays any restore jobs that are currently running.
- Instant access databases. The instant access databases that are configured for this database.
- Permissions. Any RBAC roles that are applied to the database.

View the details of an Oracle database

To view the details of an Oracle database

- 1 On the left, click **Workloads > Oracle**. Then click the **Databases** tab.
- 2 Click on the database that you want to view.

View the details of an Oracle pluggable database

To view the details of a pluggable database

- 1 On the left, click **Workloads > Oracle**. Then click the **Pluggable databases** tab.
- 2 Click on the database that you want to view.

On the **Pluggable databases** tab, the `PDB$SEED` is not visible. You can search the pluggable database list to view any of the `PDB$SEED` pluggable databases. Alternatively, select **Workloads > Oracle > Database** and then select a container database to locate the pluggable database in the list.

Edit the details of an Oracle database

You can edit the details for an Oracle database. The web UI only displays the properties that apply to the particular database type.

To edit the details of an Oracle database

- 1 On the left, click **Workloads > Oracle**. Then click the **Databases** tab.
- 2 Locate the database that you want to edit and click **Actions > Edit**.
- 3 For RAC and Single Node RAC, enter the load balance and Oracle wallet information. For a Single instance database, enter the TNS name and Oracle wallet information.
- 4 When you are done with the changes, click **Save**.

Associate an RMAN catalog with a database

After you add an RMAN catalog to NetBackup, you can then associate databases with that RMAN catalog.

To associate an RMAN catalog with a database

- 1 On the left, click **Workloads > Oracle**.
- 2 Select one or more databases.
- 3 Click **Actions > Manage RMAN catalog**.
- 4 Click **Select RMAN catalog**.
- 5 Select the RMAN catalog that you want to associate with the selected databases.
- 6 Click **Confirm**.

Remove an RMAN catalog from a database

You can remove the associated RMAN catalog from a database. Note that if you remove this association, backups are no longer recorded in the RMAN catalog.

To remove an RMAN catalog from a database

- 1 On the left, click **Workloads > Oracle**.
- 2 Select one or more databases.
- 3 Click **Actions > Manage RMAN catalog**.
- 4 Click **Remove RMAN catalog from databases**.
- 5 Click **Confirm**.

Activate or deactivate a database

You can make a database inactive in NetBackup so it is excluded from a backup. For example, if the database is under maintenance.

To deactivate a database

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Locate the database that you want to deactivate.
- 3 Click **Actions > Deactivate > Deactivate**.

To activate a database

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Locate the database that you want to deactivate.
- 3 Click **Actions > Activate > Activate**.

Remove a database or pluggable database

You can remove an Oracle database from NetBackup. This database may be one that you added manually to NetBackup. Or, you can immediately remove a database without waiting for the automatic cleanup process to run. For example, a database that was uninstalled in the Oracle environment.

To remove a database

- 1 If necessary remove the Oracle database from any policies or protection plans.
- 2 On the left, click **Workloads > Oracle** and then click the **Databases** tab.

- 3 Select the database that you want to remove.
- 4 Click **Remove > Remove**.

To remove an Oracle pluggable database

- 1 On the left, click **Workloads > Oracle**. Then click the **Pluggable databases** tab.
- 2 Select the database.
- 3 Click **Remove > Remove**.

Manually add a database

Newly discovered Oracle databases on clients are automatically added to the NetBackup database. However, you may not want to wait for the discovery service to discover a new database. In this case you can add a database.

To manually add a database

- 1 On the left, click **Workloads > Oracle** and then click **Databases**.
- 2 Click **Add**.
- 3 For the **Database type**, select from the following options.
 - To add a RAC database, select **RAC**.
 - To add a RAC One Node database, select **RAC One Node**.
 - To add a single instance database, select **Single Instance**.
- 4 To register Oracle Wallet credentials with a database, additional fields are required when you add the database.

RAC. The configuration must include the **Oracle Wallet path**.

RAC One Node. The configuration must include the **Oracle Wallet path**.

Single instance database. The configuration must include the **TNS name** (the Net service name or TNS alias) and the **Oracle Wallet path**.

See [“Configure an Oracle Wallet with RAC within NetBackup”](#) on page 51.

- 5 Enter all the other required information for the database.
- 6 Click **Next**.
- 7 Enter all the required information for the Oracle instance and click **Next**.
- 8 (Conditional) For the database types **RAC** or **RAC One Node**, add any additional instances.

- 9 Add any RBAC roles that you want to have access to the database. Then click **Next**.
- 10 To associate the database with an RMAN catalog, select **Select the RMAN catalog to be associated with the database**. Then select the RMAN catalog from the list.
- 11 Choose from the following options.
 - To add the database without credentials, click **Finish**. The credentials can be added at a later time.
 - To register the database with a credential, continue with step 12.
- 12 Click **Manage credentials** and follow the prompts.

More information is available on the credential types for Oracle.

See [“Oracle credential types in the NetBackup Credential management”](#) on page 46.

Manually add an Oracle instance

Newly discovered Oracle instances on clients are automatically added to the NetBackup database. However, you may not want to wait for the discovery service to discover a new instance. In this case you can add an instance manually.

The following procedure applies to both single instance databases that don't have an instance and to RAC databases.

To manually add an instance

- 1 On the left, click **Workloads > Oracle**.
- 2 Click the **Databases** tab.
- 3 Click the database that is associated with the instance.
- 4 Click **Add**.

This action is available for any single instance databases that don't have an instance and for RAC databases. For RAC databases you can add multiple instances.

- 5 Enter the required information for the instance.

See [“Edit Oracle instance details”](#) on page 40.
- 6 Click **Save** to add the instance.

Edit Oracle instance details

For an Oracle instance that was discovered or added to NetBackup, you can edit the instance details.

To edit the instance details

- 1 On the left, click **Workloads** > **Oracle** and then click the **Databases** tab.
- 2 Click the database that is associated with the instance that you want to edit.
- 3 Select the instance that you want to edit and click **Edit**.
- 4 Enter the required information for the instance.
- 5 (Optional) Enter the **Override TNS admin** path if you need to override the default network administration directory on the client system. Enter the fully qualified path for the network administration directory on this host.
- 6 Click **Save**.

Copy a RAC instance

When you copy the details of a RAC instance, NetBackup pre-populates the instance information, which you can then edit to create a new instance.

To copy a RAC instance

- 1 On the left, click **Workloads** > **Oracle** and then click the **Databases** tab.
- 2 Click the database that is associated with the instance that you want to copy.
- 3 Select the instance and click **Copy**.
- 4 Edit the information for the new instance that you want to create.
- 5 (Optional) Enter the **Override TNS admin** path if you need to override the default network administration directory on the client system. Enter the fully qualified path for the network administration directory on this host.
- 6 Click **Finish**.

Remove an Oracle instance

You can remove an Oracle instance from NetBackup. This instance may be one that you added manually to NetBackup. Or, you can immediately remove an instance without waiting for the automatic cleanup process to run. For example, an instance that was uninstalled in the Oracle environment. Note that if a database still exists in the Oracle environment, then NetBackup discovers the instance again during the next discovery.

To remove an Oracle instance

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Click the database that is associated with the instance.
- 3 Select the instance that you want to remove.
- 4 Click **Remove**.

Load balance Oracle RAC instances

NetBackup can be set up to load balance the instances that makeup the Oracle RAC. Use this feature to distribute the backup load across all of the instances and to exclude any Oracle RAC instances from the backup.

To load balance Oracle RAC instances

- 1 On the left, click **Workloads > Oracle**. Then click the **Databases** tab.
- 2 You can load balance instances from the **Databases** tab or from the **Instances** tab in the database details.
 - Locate and select the database, then click **Actions > Load balance**.
 - Click on the database name, then click the **Instances** tab. Click **Load balance**.

- 3 In the list **Select number of instances to load balance**, select the number of instances to include for load balancing.

If you select **All**, all instances in the Oracle RAC are available for load balancing.

- 4 In the table, select the instance or instances that you want to move up or down in priority.

Click **Move up** or **Move down** to move the instances.

Click **Move to top** to move the instance or instances to the top of the list.

Click **Move to bottom** to move the instance or instances to the bottom of the list.

(Optional) If you select **Exclude from backup**, that instance moves to the table **The RAC instances excluded from backup** table. NetBackup does not use this instance when backup operations are performed.

- 5 Click **Save**.

Managing RMAN catalogs

NetBackup lets you associate an RMAN catalog with one or more databases in your Oracle environment. Adding an RMAN catalog then provides the benefit of recording backups and backup metadata in the RMAN Catalog. For Data Guard environments, Oracle requires that you use an RMAN catalog so you must add the catalog for Data Guard databases.

Add an RMAN catalog

After you add an RMAN catalog to NetBackup, you can then associate databases with that RMAN catalog.

See [“Register credentials for an Oracle database or RMAN catalog”](#) on page 50.

To add an RMAN catalog

- 1 On the left, click **Workloads** > **Oracle** and then click the **RMAN catalogs** tab.
- 2 Click **Add**.
- 3 Enter the required details for the RMAN catalog.
- 4 Click **Next**.
- 5 Add any RBAC roles that you want to have access to the RMAN catalog.
- 6 Click **Next** and follow the prompts.
- 7 Click **Finish**.

Edit an RMAN catalog

You can edit the TNS name for an RMAN catalog. The catalog name cannot be edited. If the catalog name has changed, you must remove and re-add the catalog.

To edit an RMAN catalog

- 1 On the left, click **Workloads** > **Oracle** and then click the **RMAN catalogs** tab.
- 2 Select the RMAN catalog.
- 3 Click **Edit**.
- 4 Review and edit the **TNS name**, if necessary.
- 5 Click **Save**.

Remove an RMAN catalog

To remove an RMAN catalog

- 1 On the left, click **Workloads** > **Oracle** and then click the **RMAN catalogs** tab.
- 2 Select the RMAN catalog.
- 3 Click **Remove** > **Confirm**.

Manage a Data Guard

A Data Guard configuration in NetBackup consists of a Data Guard object with the attributes of a Data Guard name, a database name, and a database ID.

View Data Guard details

You can view the list of RAC databases or single instance databases that are associated with a Data Guard configuration.

To view Data Guard details

- 1 On the left, click **Workloads** > **Oracle** and then click the **Data Guard** tab.
- 2 Click on the Data Guard name.

The **Databases** tab contains the details of the databases that are associated with the Data Guard. The **Permissions** tab indicates any RBAC roles or permissions that are associated with the Data Guard.

Add a Data Guard

To add a Data Guard

- 1 On the left, click **Workloads** > **Oracle** and then click the **Data Guard** tab.
- 2 Click **Add Data Guard**.
- 3 Select the databases to associate with the Data Guard.

You can only add the databases that have a matching **Database name** and **Database ID** to the same Data Guard configuration. When you select a database, non-matching databases are filtered out of the list.

- 4 Click **Next**.
- 5 Add any RBAC roles that you want to have access to the Data Guard.
- 6 Click **Finish**.

Add a database to a Data Guard configuration

If a database has not been backed up with NetBackup and needs future access to the database backups that NetBackup stores, it must be added to the Data Guard configuration.

To add a database to a Data Guard configuration

- 1 On the left, click **Workloads > Oracle** and then click the **Data Guard** tab.
- 2 Select the Data Guard and click **Edit**.
- 3 Click **Add**.
- 4 Select the database that you want to add.
- 5 Click **Save**.

Remove a Data Guard

You can remove a Data Guard from NetBackup. For example, you may have decommissioned the Data Guard or renamed the database.

If you remove a Data Guard, backups that other members of the Data Guard produce are no longer accessible to the removed node.

To remove a Data Guard

- 1 On the left, click **Workloads > Oracle** and then click the **Data Guard** tab.
- 2 Select the Data Guard you want to remove and click **Remove > Confirm**.

Managing Oracle credentials

This chapter includes the following topics:

- [About Oracle credentials](#)
- [Oracle credential types in the NetBackup Credential management](#)
- [Add a credential for an Oracle database or RMAN catalog](#)
- [View the credential name that is applied to an asset](#)
- [Edit or delete a named credential](#)
- [Register credentials for an Oracle database or RMAN catalog](#)
- [Manually validate database credentials](#)
- [Configure an Oracle Wallet with RAC within NetBackup](#)

About Oracle credentials

The Credential Management System (CMS) in NetBackup lets you add and manage the credentials for Oracle databases. Multiple databases can be associated with a single credential and an RMAN catalog, which saves time when you configure for the Oracle environment in NetBackup. The RBAC role **Default Oracle Administrator** has full control over credentials in NetBackup. Or, you can apply custom roles and permissions to Oracle credentials to limit who can manage specific Oracle credentials.

For more information about NetBackup credentials, see the following topic:

[Overview of credential management in NetBackup](#)

Oracle credentials can be managed in the following ways:

- From the NetBackup Credential Management System (CMS).
 A stored credential can be used for backup and for cloning operations.
 See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.
- From the **Databases** tab in **Workloads > Oracle**. For each database, you can add the Oracle credentials that are used to manage that database.
- From the **RMAN catalogs** tab in **Workloads > Oracle**. For each catalog, you can add the Oracle credentials that are used to manage that catalog.
 See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.

Oracle credential types in the NetBackup Credential management

The NetBackup web UI supports Wallet-based authentication, Oracle authentication, and OS authentication.

The NetBackup Credential Management System (CMS) in NetBackup lets you store credentials in NetBackup to use for specific NetBackup operations.

The following limitations apply for a stored credential that is created in Credential management:

- An Oracle clone operation only supports the Oracle credential types **Use operating system credential** and **Use Oracle and operating system credentials**. You cannot use **Use Oracle Wallet** or **Use Oracle credentials** for cloning operations.
- The credential option **Use Oracle and operating system credentials** can be used for cloning operations, but during the clone NetBackup ignores the Oracle credentials.

Refer to the following Oracle information for recommendations for the authentication of users.

[Operating System Authentication of Users](#)

[Authentication of Database Administrators by Using the Operating System](#)

[Table 5-1](#) describes the credential types that are available for performing Oracle operations.

Table 5-1 Credential types in Credential management

Option to register credentials	Configuration requirements and limitations
Use Oracle Wallet	Note: Cohesity and Oracle recommend that you use Oracle Wallet for maximum security. The wallet path must be provided to NetBackup. Edit the database and add the path as per your requirements. ■ The Oracle Wallet folder location must be on a file system. ■ Each instance must have its own entry in a shared wallet. ■ For Oracle RAC: - The same path must be used for each node of the cluster. - You must put a specific connection identifier in the wallet. - For more information about the connect identifier: See “Configure an Oracle Wallet with RAC within NetBackup” on page 51. ■ The configuration for a single instance database must include the TNS name (the Net service name or TNS alias) and the Oracle Wallet path. Limitations This option is not available for cloning operations.
Use Oracle credentials	If using a TNS connection, make sure to edit the Oracle database in NetBackup to include the TNS name. It is recommended that the Oracle database user have <code>SYSPRIVILEGES</code> privileges to limit NetBackup's exposure to your database. The Oracle database user can also have <code>SYSDBA</code> privileges (based on the version of Oracle). Limitations This option is not available for cloning operations.
Use operating system credentials	It is recommended that the Oracle database user have <code>SYSPRIVILEGES</code> privileges to limit NetBackup's exposure to your database. The Oracle database user can also have <code>SYSDBA</code> privileges (based on the version of Oracle). For cloning operations, the following requirements apply for Windows: ■ The account must have OS Administrator privileges and have Oracle DBA privileges. ■ For a domain account, also specify the Operating system domain. For a local account, leave the Operating system domain empty or use a “.” ■ See the following information for requirements for the Oracle Home User. See the section called “Oracle Home User requirements” on page 48. For cloning operations on UNIX/Linux, the account must have Oracle DBA privileges.
Use Oracle and operating system credentials	When this credential type is used for cloning, NetBackup ignores the Oracle credentials. See the configuration requirements for Use operating system credential.

Oracle Home User requirements

When you create a stored credential, you must also specify the Oracle Home User credentials when the target Oracle installation was installed using an OS account. (These credentials are not a virtual account or a Windows Built-in account like Local System.)

The Oracle Home User can be any of the following: a Windows Built-in account or a Virtual account or a standard Windows user account that is not an Administrator account.

Add a credential for an Oracle database or RMAN catalog

You can use a stored credential to register Oracle databases and for RMAN catalogs. This type of credential allows NetBackup to access an Oracle database or RMAN catalog.

For client versions 8.2 and earlier, NetBackup may require the database unique name to validate a credential. This information can be found by contacting your Oracle administrator or querying the database with the following command:

```
select DB_UNIQUE_NAME from v$database;
```

To add a credential to use an Oracle database or RMAN catalog

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, click **Add** and provide the following properties:
 - Credential name (for example: server_credential1)
 - Tag (for example: Workload name)
 - Description (for example: This credential is used to access workload name)
- 3 Click **Next**.
- 4 Select **Oracle**.
- 5 Provide the authentication details that are needed to connect to the Oracle database.

See [“Oracle credential types in the NetBackup Credential management”](#) on page 46.
- 6 Click **Next**.
- 7 In **Permissions**, add one or more RBAC roles that you want to have access to the credential.

- Click **Add**.
- Select the role name.
- Select the credential permissions that you want the role to have.
- Click **Save**.

Click **Next**.

- 8 On the **Review** page, verify that all information is correct and click **Finish**.

View the credential name that is applied to an asset

You can view the named credential that is configured for an asset type. If the credentials are not configured for a particular asset, this field is blank.

To view credentials for Oracle

- 1 On the left, select **Workloads > Oracle**.
- 2 On the **Instances** tab, locate the **Credentials** column.

Edit or delete a named credential

You can edit the properties for a named credential or delete a named credential NetBackup from the **Credential management**.

Edit a named credential

You can edit a named credential to change the following: credential tag, description, category, authentication details, or permissions. You cannot change the credential name.

To edit a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to edit.
- 3 Select **Edit** and update the credential as needed.
- 4 Review the changes and select **Finish**.

Delete a named credential

You can delete a named credential that you no longer need to use with NetBackup. Be sure to apply another credential to any assets that use the credential you want to delete. Otherwise, backups and restores may fail for those assets.

To delete a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to delete.
- 3 Select **Delete > Delete**.

Register credentials for an Oracle database or RMAN catalog

To protect Oracle you must register credentials with the Oracle databases. Use the **NetBackup web UI**, the `nboracmd` command, or the NetBackup APIs to access the repository for database registration.

After the discovery service adds new databases to the repository, you can register credentials with those databases. When you register credentials, NetBackup queries the Oracle database for metadata. The database must be in an OPEN state while NetBackup validates credentials. The NetBackup Credential Management System (CMS) stores the credentials for your databases.

You can register or update credentials for a database or an RMAN catalog at any time. When you manually add a database, you can choose not to include the credentials at that time.

When an Oracle RAC database is discovered, that database does not have a **Database ID**. A **Database ID** is required to manually add additional RAC instances to the database. You must register the RAC database and provide a **Database ID** before adding additional instances.

For client versions 8.2 and earlier, NetBackup may require the database unique name to validate a credential. This information can be found by contacting your Oracle administrator or querying the database with the following command:

```
select DB_UNIQUE_NAME from v$database;
```

Refer to the following Oracle information for recommendations for the authentication of users.

[Operating System Authentication of Users](#)

[Authentication of Database Administrators by Using the Operating System](#)

To register credentials for an Oracle database or RMAN catalog

- 1 On the left, click **Workloads > Oracle**.
- 2 Select one of the following:
 - To register credentials for a database, click the **Databases** tab.
 - To register credentials for an RMAN catalog, click the **RMAN catalogs** tab.
- 3 Select one or more databases or a catalog. Then click **Manage credentials**.
- 4 Click **Select from existing credentials**. Then click **Next**.
 Alternatively, you can add a new credential.
 See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.
- 5 Select the credential. Then click **Next**.
- 6 The credential is validated. NetBackup saves the credential even if the validation fails.
- 7 Review the **Registered** column to see that the database or catalog is now registered.

Manually validate database credentials

You can manually validate the credentials that are registered for a database.

To manually validate database credentials

- 1 On the left, click **Workloads > Oracle** and then click the **Databases** tab.
- 2 Select the databases for which you want to validate the credentials.
- 3 Click **Actions > Validate current credentials**.

Configure an Oracle Wallet with RAC within NetBackup

The configuration and setup of the Oracle Wallet in NetBackup is a two-step process. You add descriptors first, then you register the wallet. In the cases of Oracle RAC, your descriptors must enumerate the list of RAC instances that comprise your RAC cluster.

NetBackup Oracle Wallet prerequisites:

- The Oracle wallet location must be accessible from all nodes of the RAC cluster.

- Using a shared location is encouraged for maintainability.
An example storage location can be: An Oracle ACFS file system that is mounted on each node or an NFS share accessible to each node. The mount point of the shared location must be the same on each node.
- If the wallet is not in a shared location, it must be in an identically duplicate location on each node of the RAC cluster. The full contents of the wallet must also be duplicated on each node of the RAC cluster.

To configure Oracle Wallet with RAC in NetBackup

- 1 Retrieve the RAC connect descriptors for all instances in the RAC database. Place the list of connect descriptors in a text file for easy access for step 2. Use one of the following methods:

Get RAC connect descriptors from the NetBackup web UI:

- On the left, click **Workloads > Oracle** and then click **Databases**.
- Click on the RAC database.
- On the **Instances** tab, click **Show Oracle Wallet connect descriptors**.
- Click **Copy** to copy a descriptor.

Manually create the RAC connect descriptors:

- If you don't have this information, use the NetBackup web UI to retrieve it. You need to obtain the scan name, service name, and port number for the given RAC database. If the RAC instances for the RAC database are not known, collect the list of instance names.
- For each instance of a RAC, you must insert this information to create a connect descriptor. Insert the scan name, service name, and port number (from RAC database), as well as the instance name (from RAC instance) into the following example:

```
(DESCRIPTION= (ADDRESS= (PROTOCOL=TCP)
(HOST=<INSERT SCAN NAME>) (PORT=<INSERT PORT NUMBER>))
```

```
(CONNECT_DATA=(SERVER=DEDICATED) (SERVICE_NAME=<INSERT SERVICE NAME>)
(INSTANCE_NAME=<INSERT INSTANCE NAME>))
```

- 2** Add the connect descriptors with the Oracle `MKSTORE` utility. The descriptors are case-sensitive and must match exactly to what is in NetBackup.

```
mkstore -wrl /db/orac183/wallet/ -CreateCredential
' (DESCRIPTION= (ADDRESS= (PROTOCOL=TCP) (HOST=exampleScanName.domain.com)
(PORT=1521)) (CONNECT_DATA= (SERVER=DEDICATED)
(SERVICE_NAME=orac183.domain.com) (INSTANCE_NAME=orac1831))) '
testUser testPassword
```

```
mkstore -wrl /db/orac183/wallet/ -CreateCredential
' (DESCRIPTION= (ADDRESS= (PROTOCOL=TCP) (HOST=exampleScanName.domain.com)
(PORT=1521)) (CONNECT_DATA= (SERVER=DEDICATED)
(SERVICE_NAME=orac183.domain.com) (INSTANCE_NAME=orac1832))) '
testUser testPassword
```

- 3** In the NetBackup web UI, register the RAC with the credential that is configured for Oracle Wallet.

See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.

- 4** (Optional) If you get an error when you attempt to register the RAC, review the error message. Compare the descriptors in the error message with what you generated in step 1 and what you inserted into your Oracle Wallet.

Configuring Oracle policies

This chapter includes the following topics:

- [Preparing for NetBackup for Oracle configuration](#)
- [About Oracle Intelligent Policies \(OIP\)](#)
- [About script-based Oracle policies](#)
- [Configuring the logon account for the NetBackup Client Service for NetBackup for Oracle](#)
- [Testing configuration settings for NetBackup for Oracle](#)

Preparing for NetBackup for Oracle configuration

The major part of configuring NetBackup for Oracle is to create and configure the Oracle policies. The following topics prepare you to configure NetBackup for Oracle policies:

- See [“About Oracle policy configuration”](#) on page 54.
- See [“Oracle backup policy types”](#) on page 56.
- See [“Configuring the logon account for the NetBackup Client Service for NetBackup for Oracle”](#) on page 97.
- See [“Configuring the Maximum jobs per client”](#) on page 58.

About Oracle policy configuration

The NetBackup web UI offers two ways to configure an Oracle policy.

- **Oracle Intelligent Policies.** This method lets you create a single policy to protect multiple databases or pluggable databases that are spread over multiple clients. You select the Oracle databases that are automatically discovered in the

NetBackup environment. Among the features that these policies provide is the ability to schedule frequent backups of archived redo logs. These backups are accomplished in minutes instead of hours or days.

- Script-based policies. This method lets you create an Oracle backup policy by using a script that is based on a list of clients.

A backup policy for a database defines the backup criteria for a specific group of databases (Oracle Intelligent Policy) or clients (script-based policy).

The Intelligent Oracle Policy includes the following criteria:

- Storage unit and media to use
- Policy attributes
- Backup schedules. Automatic schedule and archive log schedule.
- Databases to back up
- Backup selections: Whole database, tablespaces, data files, FRA

The script-based policy includes the following criteria:

- Storage unit and media to use
- Policy attributes
- Backup schedules: Automatic schedule and application schedule.
- Clients to back up
- Backup script files to run on the clients

To back up the database environment, define at least one script-based Oracle policy with the appropriate schedules and clients. Or, you can configure a single Oracle Intelligent Policy that includes all databases.

Note: The Backup, Archive, and Restore GUI cannot be used for performing Oracle backups and restores.

Most requirements for database policies are the same as for file system backups. In addition to the policy attributes for Oracle, other attributes are available that you should consider.

See the [NetBackup Administrator's Guide, Volume I](#).

See [“About Oracle Intelligent Policies \(OIP\)”](#) on page 59.

See [“About policy attributes”](#) on page 79.

Oracle Home User permissions when NetBackup SAN Client is used

To use the NetBackup SAN Client to protect Oracle on Windows, the Oracle user must have administrator privileges. Starting with Oracle Database 12c Release 1 (12.1), Oracle Database on Windows supports the use of Oracle Home User. The Oracle Home User is specified at the time of Oracle Database installation and is used to run the Windows services for the Oracle home. The Oracle Home User that is used to run Windows services is similar to the Oracle user for Oracle Database on Linux.

For more information, refer to the Oracle document "Supporting Oracle Home User on Windows" at the following location:

http://docs.oracle.com/cd/E16655_01/win.121/e10714/oh_usr.htm

To use NetBackup SAN Client, make sure to select **Use Windows Built-in Account** during Oracle Database installation. Making this selection enables the Windows services for the Oracle home to run as `LocalSystem` or `LocalService`.

Oracle backup policy types

Table 6-1 shows the Oracle backup policy types that you can specify.

Table 6-1 Oracle backup types

Backup type	Description
Application backup – Script-based policy using streamed data only	The Application backup schedule enables user-controlled NetBackup operations from the client. These operations include those initiated from the client and those initiated by an automatic schedule on the primary server. NetBackup uses the Application backup schedule when the user starts a backup manually. Configure at least one Application backup schedule for each database policy. The Default-Application-Backup schedule is configured automatically as an Application backup schedule.
Full backup – Script-based policy	Stream-based backup: The specified script in the Backup selections tab is run. If the script is set up properly, RMAN initiates a full stream-based backup (full or incremental level 0). Note: The Application backup schedule properties (For example: storage and retention) are used. RMAN proxy backup: The specified script in the Backup selections tab is run. If the script is set up properly, RMAN initiates a proxy backup. Note: The Full backup schedule properties (For example: storage and retention) are used for the proxy portion of the backup. The Application backup schedule properties (For example: storage and retention) are used for the streamed portion of the backup.

Table 6-1 Oracle backup types (*continued*)

Backup type	Description
Differential incremental backup – Script-based policy	Stream-based backup: The specified script in the Backup selections tab is run. If the script is set up properly, RMAN initiates a stream-based incremental level 1 backup. Note: The Application backup schedule properties (storage, retention, etc.) are used. RMAN proxy backup: This backup type should only be used for BLI backups. If you do not use a proxy backup for a BLI backup, then a Full backup schedule should be used. The specified script in the Backup selections tab is run. If the script is set up properly, RMAN initiates a proxy backup. Note: The Differential incremental backup schedule properties (storage, retention, etc.) are used for the proxy portion of the backup. The Application backup schedule properties (storage, retention, etc.) are used for the streamed portion of the backup.
Cumulative incremental backup – Script-based policy	Stream-based backup: The specified script in the Backup Selections tab is run. If the script is set up properly, RMAN initiates a stream-based incremental level 1 cumulative backup. Note: The Application backup schedule properties (storage, retention, etc.) are used. RMAN proxy backup: This backup type should only be used for BLI backups. If you do not use a proxy backup for a BLI backup, then a Full backup schedule should be used. The specified script in the Backup selections tab is run. If the script is set up properly, RMAN initiates a proxy backup. Note: The Cumulative incremental backup schedule properties (storage, retention, etc.) are used for the proxy portion of the backup. The Application backup schedule properties (storage, retention, etc.) are used for the streamed portion of the backup.
Full backup – OIP policy	Stream-based backup: Dynamically generates an RMAN script on each client for the databases that are defined in the Instances and databases tab. The script initiates an Incremental Full (INCREMENTAL LEVEL 0) backup. RMAN proxy backup (Policy is defined to perform a snapshot): Dynamically generates an RMAN script on each client for the databases that are defined in the Instances and Databases tab to initiate a proxy backup. Note: The Full backup schedule properties (storage, retention, etc.) are used for both the streamed and the proxy data.

Table 6-1 Oracle backup types (*continued*)

Backup type	Description
Differential incremental backup – OIP policy	Stream-based backup: Dynamically generates an RMAN script on each client for the databases that are defined in the Instances and databases tab. The script initiates a Differential Incremental (INCREMENTAL LEVEL 1) backup. RMAN proxy backup (policy is defined to perform a snapshot): ■ The policy has Perform block level incremental backups selected. An RMAN script is dynamically generated on each client for the databases that are defined in the Instances and databases tab to initiate a proxy backup. ■ The policy does not have Perform block level incremental backups selected. An RMAN script is dynamically generated on each client for the databases that are defined in the Instances and databases tab. A Differential Incremental (INCREMENTAL LEVEL 1) stream-based backup is initiated. Note: The Differential incremental backup schedule properties (storage, retention, etc.) are used for both the streamed and the proxy data.
Cumulative Incremental backup – OIP policy	Stream-based backup: Dynamically generates an RMAN script on each client for the databases that are defined in the Instances and Databases tab. The script initiates a Cumulative Incremental (INCREMENTAL LEVEL 1 CUMULATIVE) backup. RMAN proxy backup (Policy is defined to perform a snapshot): ■ The policy has Perform block-level incremental backups selected. An RMAN script is dynamically generated on each client for the databases that are defined in the Instances and Databases tab to initiate a proxy backup. ■ The policy does not have Perform-block level incremental backups selected. An RMAN script is dynamically generated on each client for the databases that are defined in the Instances and databases tab. A Cumulative incremental (INCREMENTAL LEVEL 1 CUMULATIVE) stream-based backup is initiated. Note: The Cumulative incremental backup schedule properties (storage, retention, etc.) are used for both the streamed and the proxy data.
Archived redo log backup – OIP policy only	The policy dynamically generates an RMAN script on each client for the databases that are defined in the Instances and databases tab. The policy initiates a stream-based archive redo log backup. Note: The frequency is granular down to intervals of minutes.

Configuring the Maximum jobs per client

The **Maximum jobs per client** specifies the maximum number of concurrent backups that are allowed per client.

Note: Enter a large enough value for the **Maximum jobs per client** attribute to meet the number of jobs that Oracle runs. You may need to experiment with different values at your site.

To configure the maximum jobs per client

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary, select **Connect**. Then select **Edit primary server**.
- 4 Select **Global attributes**.
- 5 Select the appropriate value for **Maximum jobs per client**.

The default is 1.

About Oracle Intelligent Policies (OIP)

The Oracle Intelligent Policy (OIP) feature is a method of Oracle policy backup based on Oracle databases or RAC databases. This method precludes the need to create a script for your Oracle policies. The OIP feature has the following elements:

- You can create a single policy to protect multiple Oracle databases or RAC databases that are spread over multiple clients.
- All databases that you want to back up must be registered with credentials. If multiple database share the same credentials, you can select the same credential for those databases.
- The database administrator can use Credential management (CMS) to manage credentials and the registration of databases and RMAN catalogs, through the NetBackup web UI or the APIs. The use of the CMS provides improved security throughout the system.
- You are not required to know RMAN or to write RMAN scripts. Instead, this feature automatically generates the scripts at run-time.
- The job details in the Activity monitor let you view the backup summary, database state, RMAN input, and RMAN output for the OIP. The **Client or display name** column indicates the database that was backed up by the associated policy.
- Enhanced error codes enable faster identification, troubleshooting, and correction of problems. You can easily restart a failed job.
- OIP only requires automatic backup schedules for the data movement, which simplifies how retention works on the backup pieces. Application backup

schedules are not needed. An archived log schedule backs up the archived redo logs within intervals of minutes.

- The OIP automatically selects parameter settings at run-time that enable optimal deduplication.
- An OIP can protect an Oracle database when the Oracle DBA places database backups in the share on a NetBackup appliance.
- The OIP can create and maintain a full set of data file copies in the share on a NetBackup appliance. The Accelerator option is used to update the data file copies using only the changed blocks since the last full backup.
- Container databases (CDB) and pluggable databases (PDB) and can be protected.

Create an Oracle Intelligent Policy (OIP)

This topic guides you through the steps for setting up an Oracle Intelligent Policy (OIP). An OIP is used with Oracle CDB and PDB databases, Oracle Copilot, and regular Oracle database instance backups. This method precludes the need to create scripts for your Oracle policies.

Table 6-2 Steps for creating an OIP

Steps	Task	Instructions
Step 1	Register Oracle instances.	NetBackup automatically discovers Oracle instances and displays them in the instance repository. An instance must be registered in order for that instance to be included in an OIP. See “Register credentials for an Oracle database or RMAN catalog” on page 50.
Step 2	Add a policy.	In the NetBackup web UI on the left click Protection > Policies. Click Add. See “Oracle backup policy types” on page 56.
Step 3	Configure the Attributes tab.	In the Policy type list, select Oracle. This action causes the tabs along the top of the display to change to a unique Oracle tab set. The Use Accelerator option has a different function when used with an OIP. This option is automatically selected when certain options in the Backup Selections tab are set during Oracle Copilot configuration. See “About Oracle Copilot” on page 135. For information on additional policy attributes, see the NetBackup Administrator’s Guide, Volume I.

Table 6-2 Steps for creating an OIP (*continued*)

Steps	Task	Instructions
Step 4	Configure the Schedules tab.	The schedules that are defined on the Schedules tab determine when backups occur for an OIP. See “About NetBackup for Oracle schedule properties using Oracle Intelligent Policy” on page 62.
Step 5	Configure the Instances and Databases tab.	Select the databases that the OIP will back up. An OIP must use either the Protect non-RAC databases and pluggable databases option or the Protect RAC databases and pluggable databases option. See “Instances and databases tab” on page 66.
Step 6	Configure the Backup Selections tab.	You can backup the Whole database, Partial database - Tablespaces, Partial database - Datafiles, Fast Recovery Area (FRA), Database backup shares , or the Whole Database - Datafile Copy Share . See “Backup selections tab” on page 67. See “Configuring an OIP using a share on the NetBackup appliance (Oracle Copilot)” on page 137.
Step 7	Configure the Oracle tab.	The tab contains setup options for databases, tablespaces, data files, archived redo logs, file name formats, and database backup shares. See “Oracle tab” on page 68.
Step 8	(Optional) Create <code>bpstart_notify</code> and <code>bpend_notify</code> scripts	See “Create <code>bpstart_notify</code> and <code>bpend_notify</code> scripts for an Oracle Intelligent Policy (OIP)” on page 75.

Oracle database upgrade effect on Oracle Intelligent Policies

Upgrade of an Oracle database causes instance information for the upgraded database to become invalid. If this instance is associated with one or more current NetBackup for Oracle Intelligent Policies, run-time failures can occur. The issue occurs when an Oracle database is upgraded to a new version. The new version is likely to have a different ORACLE_HOME, ORACLE_SID, or Oracle User. If any of these values have changed, the existing instance information in the NetBackup instance repository and in the current Oracle Intelligent Policies becomes invalid. When the discovery service (`nbdisco`) polls the clients again, it discovers the database as a new instance. Consequently, there is no way to associate the new instance to the old instance.

This issue is not version-specific and can affect any valid Oracle upgrade patch, such as:

- Oracle 10 to version 11
- Oracle 10 to version 12
- Oracle 11 to version 12

For more information on valid Oracle upgrade paths, review the following documentation on the Oracle Support website:

<http://www.oracle.com/technetwork/database/upgrade/upgrading-oracle-database-wp-12c-1896123.pdf>

Therefore, when an existing Oracle database is upgraded and the ORACLE_HOME, ORACLE_SID, or Oracle User are modified, remove the existing instance in the instance repository. After the existing instance is removed, update the instance repository with the new instance information. Make sure to update any policies with the newly-discovered instances.

Configuring NetBackup for Oracle automatic backup schedules

Each policy has an automatic backup schedule. These schedules initiate automatic backups and specify when a user can initiate operations.

To configure an automatic backup schedule

- 1 On the **Policy** dialog box, click the **Schedules** tab.
- 2 Click **New**.
- 3 Specify a unique name for the schedule.
- 4 Select the **Type of backup**.
- 5 Specify the other properties for the schedule.

See [“Schedule properties for Oracle script-based policies ”](#) on page 80.

- 6 Click **OK**.

About NetBackup for Oracle schedule properties using Oracle Intelligent Policy

This topic describes the schedule properties that have a different meaning for Oracle Intelligent Policy backups than for file system backups. Other schedule properties vary according to your specific backup strategy and system configuration. Additional information about other schedule properties is available.

See the [NetBackup Administrator's Guide, Volume I](#).

Table 6-3 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure. See “Oracle backup policy types” on page 56.
Schedule type	You can schedule a backup in one of the following ways: ■ Frequency This setting is used only for scheduled backups. It is not used for user-directed backups. The Frequency specifies the period of time that can elapse until the next backup or archive operation begins on this schedule. For example, assume that the frequency is 7 days and a successful backup occurs on Wednesday. The next full backup does not occur until the following Wednesday. Typically, incremental backups have a shorter frequency than full backups. ■ Calendar This setting is used only for scheduled backups. It is not used for user-directed backups. The Calendar option lets you schedule the backup operations that are based on specific dates, recurring week days, or recurring days of the month.
Retention	Specifies a retention period to keep backup copies of files before they are deleted. The retention period for an automatic schedule controls how long NetBackup keeps records of when scheduled backups occurred. Set the time period to retain at least two full backups of your database. In this way, if one full backup is lost, you have another full backup to restore. The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be at least 2 weeks. The NetBackup scheduler compares the latest record of the automatic backup schedule to the frequency of that automatic backup schedule. This comparison is done to determine whether a backup is due. So if you set the retention period to expire the record too early, the scheduled backup frequency is unpredictable. However, if you set the retention period to be longer than necessary, the NetBackup catalog accumulates unnecessary records. Oracle is not notified when NetBackup expires a backup image. Use Oracle RMAN repository maintenance commands to periodically delete expired backup sets from the Oracle RMAN repository. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling.
Multiple copies	If you want to specify multiple copies of a backup for the policy, configure Multiple copies on the application backup schedule.

Table 6-3 Description of schedule properties (*continued*)

Property	Description
Accelerator forced rescan	This option instructs NetBackup to re-copy all the data files to the share. This option is only available when Whole Database - Datafile Copy Share is selected in the Backup selections tab and the Use Accelerator option is selected in the Attributes tab. This option forces the creation of a new set of database data file copies. When this option is not selected, the data file copies in the share are updated using an incremental backup. The incremental backup contains only the changed blocks since the last full backup.

Oracle Intelligent Policy - Storage and Retention

This topic describes storage and retention properties of the Oracle Intelligent Policy.

See the [NetBackup Administrator's Guide, Volume I](#).

Table 6-4 Storage and retention behavior

Property	Description
Policy is a snapshot type	If the policy is a snapshot type, the following are the possible scenarios of the retention behavior: ■ If the schedule does not override the policy storage unit, and the policy storage unit is a non-snapshot SLP, the SLP determines the retention period and the policy uses the policy storage unit. ■ If the schedule does not override the policy storage unit and the policy storage unit is not an SLP, the schedule determines the retention period, and the policy uses the policy storage unit. ■ If the schedule does override the policy storage unit with an SLP, and it is not a snapshot SLP, the override storage unit takes precedence over the policy storage unit, and the SLP determines the retention period. ■ If the schedule overrides the policy storage unit with a snapshot SLP, the policy storage unit must be a non-snapshot SLP. The SLP on the policy storage unit determines the retention period for the streamed data. Also, the SLP on the schedule determines the retention for the snapshot data.

Table 6-4 Storage and retention behavior (*continued*)

Property	Description
Policy is not a snapshot type	If the policy is not a snapshot type, the following are the possible scenarios of the retention behavior: ■ If the schedule does not override the policy storage unit and the policy storage unit is not an SLP, the schedule determines the retention period. ■ If the schedule does not override the policy storage unit and the policy storage unit is an SLP, the SLP determines the retention period. ■ If the schedule overrides the policy storage unit, and the schedule storage unit is not an SLP, the schedule determines the retention period. ■ If the schedule overrides the policy storage unit and the schedule storage unit is an SLP, the SLP determines the retention period.

The following are examples of the Oracle Intelligent Policy storage and retention behavior for snapshot-based policy types.

Policy storage	Schedule storage	Streamed data retention is derived from:	Snapshot data retention is derived from:
AdvancedDisk	-	Schedule	Schedule
AdvancedDisk	SLP	SLP	SLP
SLP	-	SLP	SLP
Tape library	-	Schedule	Schedule
Non-Snapshot SLP	Snapshot SLP	Non-Snapshot SLP	Snapshot SLP
AdvancedDisk	Snapshot SLP	Invalid configuration	Invalid configuration

The following are examples of the Oracle Intelligent Policy storage and retention behavior for stream-based policy types.

Policy storage	Schedule storage	Streamed data retention is derived from:
AdvancedDisk	-	Schedule
SLP	AdvancedDisk	Schedule
AdvancedDisk	SLP	SLP
SLP	-	SLP

About Oracle Intelligent Policy primary server behavior

By default for an Oracle Intelligent Policy, the client uses the first server in the server list to start the Oracle backup or restore operation. However, you may want the operation to recognize the primary server name that is passed down from the primary server. If so, do one of the following:

- On Windows, enter the `USE_REQUESTED_MASTER = TRUE` statement into a text file (for example, `new_config.txt`). Then use the following command on the primary or the media server to send this newly created configuration file to the client host:

```
# bpsetconfig -h myoracleclient new_config.txt
```

- On UNIX, add `USE_REQUESTED_MASTER = TRUE` to the `bp.conf` file, which enables more than one primary server to back up the client.

Instances and databases tab

If you add a new Oracle policy or change an existing Oracle policy, the **Instances and databases** tab displays as part of the policy configuration.

The following options are available.

- **Protect non-RAC databases and pluggable databases** (OIP option). This panel displays all databases that you have selected to back up for this policy. Select the databases that you want to add to the list. If a database does not appear in this panel because it is unregistered, you can register it and add it to the policy later.
- **Protect RAC databases and pluggable databases**. This option displays all RAC databases that you have selected to backup for this policy. If no RAC databases are displayed, you must add these databases. On the left select **Workloads > Oracle**. Then click the **Databases** tab.
- **Clients for use with scripts** (Non-OIP option). This option is not for use with OIP. Select this option if you want to use the client with scripts method of configuring an Oracle policy. If you select this option, the existing backup selections are erased. Also, the **Oracle** tab and the **Instances and databases** tab are removed, because those options must now be set in the RMAN script that the user supplies.

See [“Backup selections tab”](#) on page 67.

Backup selections tab

The **Backup selections** tab lets you change the type of Oracle backup. The following options are available.

Table 6-5 Oracle backup options

Backup option	Description
Whole database	Backs up the whole database. By default, the Whole database option is selected and the backup selections contain the directive <code>WHOLE_DATABASE</code>.
Partial database - Tablespaces	Backs up only tablespaces. Click Add to select the tablespaces for the policy to back up. This selection applies across all the databases and pluggable databases (PDBs) that are selected in the policy. If a tablespace is selected for one database or PDB, that same tablespace is backed up for all databases and PDBs in the policy.
Partial database - Datafiles	Backs up only datafiles. Click Add to select the data files for the policy to back up. This selection applies across all the databases and PDBs that are selected in the policy.
Fast Recovery Area - (FRA)	This option backs up the contents of the FRA. For the Oracle database to be restored and recovered, make sure that the FRA contains a recoverable image set when it is backed up.
Database Backup Shares	Note: This feature requires a NetBackup appliance that runs software version 2.7.1 or later. This option is used when the Oracle DBA places database backups in the share on a NetBackup appliance (Oracle Copilot). The directive <code>ALL_DATABASE_BACKUP_SHARES</code> is automatically added to the selection list, which backs up all the shares that are used on all appliances per database. Optionally, you can click Browse to select the appliance shares that Oracle DBAs have created for backups of the databases that are configured in the policy. Or click Add to add an appliance share to the policy manually. When you back up appliance shares for multiple databases, this option applies across all the databases that are selected in the policy. If a share is selected for one database, the data in that share is backed up for all the databases in that policy.

Table 6-5 Oracle backup options (*continued*)

Backup option	Description
Whole Database - Datafile Copy Share	Note: This feature requires a NetBackup appliance that runs software version 2.7.1 or later. This option is used to create and maintain a full set of data file copies in the share on a NetBackup appliance (Oracle Copilot). See the section called “Additional information for the Whole Database - Datafile Copy Share option” on page 68.

Backups with OIP that contains CDB with PDBs

If you set up an OIP and that policy contains a CDB with PDBs, the `CDB$ROOT` is automatically included in the backup. If the policy contains a PDB that is not found when a backup is performed, an error 5421 or 5422 appears in the Activity monitor.

Additional information for the Whole Database - Datafile Copy Share option

The **Whole Database - Datafile Copy Share** option allows the NetBackup Administrator to choose an appliance share as the destination for the first backup copy. When the policy runs the first time, an RMAN script is generated that creates a full set of Oracle data file copies. The copies reside in the appliance share. The next time that the full schedule runs, the backup is accelerated if the **Use Accelerator** option is selected. The RMAN script that is generated performs an incremental backup and the changed blocks are merged into the data files. This incremental backup creates an updated full set of Oracle data file copies. After the new full copy is created in the appliance share, an SLP is used to make additional copies of the full backup. The first copy is always a `remote_vxfs` snapshot.

The **Use Accelerator** feature is automatically selected when you configure an OIP with the **Whole Database - Datafile Copy Share** option. The first time that the full schedule runs it creates a full set of data file copies. After the first full schedule, only the changes are backed up as a backup set and merged with the existing full backup. Basically, an incremental merge is performed and Oracle's Block Change Tracking feature should be enabled for faster incremental backups. Only one share can be set up so if you have two or more databases, all databases reside in the same share.

See [“About Oracle Copilot”](#) on page 135.

Oracle tab

This tab includes the additional options that are available for Oracle backups.

Table 6-6 Oracle tab fields

Field	Description
Tablespace or datafile options	Number of parallel streams is the number of parallel backup streams that can be used in a backup operation. When an Oracle RAC is included in the policy, setting the Number of parallel streams option sets the streams per node. For example, if the Oracle RAC has two nodes that are active and the Number of parallel streams is set to 2, NetBackup uses four streams. NetBackup uses two parallel streams per active node. Select Specify read-only tablespace options to enable read-only tablespace options. SKIP means to skip the read-only tablespace during backup. FORCE means that RMAN backs up all files. Select Offline (cold) database backup to shut down the Oracle database and put it in the mount state. When this option is used with a PDB, the PDB is put in a mounted state for the backup. Once the backup has completed, the PDB is returned to the state it was in before the backup. Select Skip offline datafiles to direct the backup operation to not access offline data files.

Table 6-6 Oracle tab fields (*continued*)

Field	Description
Tablespace or datafile options (continued)	The Data Guard backup options let you specify a policy to always back up the primary or a standby database: ■ None - NetBackup does not look at the Data Guard role. At the policy run time each database or RAC database in the policy is backed up. ■ Require primary - This option backs up the primary Data Guard database in the list of databases (or RAC databases) in the policy. If the primary is not available, the backup fails. A message appears in the job details that indicates NetBackup cannot find the primary database. ■ Require standby - This option backs up a standby Data Guard database in the list of databases (or RAC databases) in the policy. The first standby database NetBackup finds, is backed up. Currently there is no way to set a priority order for the standby databases. Oracle Intelligent Policies do not support RMAN proxy type backups of Data Guard standby databases. If the standby is not available, the backup fails. A message appears in the job details that indicates NetBackup cannot find the standby database. ■ Prefer standby - This option backs up a standby Data Guard database in the list of databases (or RAC databases) in the policy. If a standby database is not found and the primary is available, then the primary database is backed up. Oracle Intelligent Policies do not support RMAN proxy type backups of Data Guard standby databases. Note: If you choose Require primary, Require standby, or Prefer standby, then the following applies: Only one database is backed up at policy run time. All databases or RAC databases that are listed in a policy must be part of the same Data Guard configuration. Each Data Guard configuration needs a separate policy. The policy cannot have the backup type Database Backup Shares or Whole Database – Datafile Copy Share selected. These selections are incompatible with Data Guard aware policies.

Table 6-6 Oracle tab fields (*continued*)

Field	Description
MSDP storage	The option Use multiple MSDP nodes allows backups of Oracle data files in parallel, which accelerates the backup and restore jobs. You can use multiple MSDP storage units, a storage unit of an MSDP cluster, or a storage unit of an MVG (MSDP volume group) volume. The dropdown list contains the available storage items, including MSDP storage units and SLPs that are configured in the system. Select the storage to which you want to distribute the data: ■ MSDP storage units or SLPs. Select each storage unit that you want to use. If replication is needed, all SLP targets must be in the same target domain. (These targets include the policy storage SLP and the selected storage SLPs.) This configuration ensures that the Oracle data files, archived redo log, and control file are in the same domain after replication. ■ A storage unit of an MSDP cluster. Note: Select only this one storage unit. Do not select any other storage units with this selection. This option distributes Oracle data files to multiple nodes in the cluster. The recommended value for Number of parallel streams is a multiple of the number of cluster nodes. If there are 4 nodes in the cluster, the value is 4n. For example: 4, 8, or 12. ■ A storage unit of an MVG (MSDP volume group) volume. Note: Select only this one storage unit. Do not select any other storage units with this selection. This option distributes Oracle data files to the MSDP disk volumes in the MVG volume. The recommended value for Number of parallel streams is a multiple of the number of MSDP disk volumes. If there are 4 MSDP disk volumes in the MVG volume, the value is 4n. For example: 4, 8, or 12. By default, <code>FILESPPERSET</code> is set to 1 for the best deduplication ratio. You can modify this value manually to tune the restore performance for the Oracle RAC environment. On each Oracle RAC node, in the <code>/usr/openv/netbackup/bp.conf</code>, add a configuration parameter <code>ORACLE_MULTIPLE_STORAGE_OVERRIDE_FILESPPERSET</code>. For example, <code>ORACLE_MULTIPLE_STORAGE_OVERRIDE_FILESPPERSET = 8</code>.

Table 6-6 Oracle tab fields (*continued*)

Field	Description
Specify maximum limits	Select Specify maximum limits to access several I/O and backup set limits. The following parameters should only have to be modified on rare occasions. If these values are not changed, the backup uses the default values that are defined in RMAN. RMAN default values usually provide the best performance. Maximum I/O Limits parameters: ■ Read rate (KB/sec) (RATE) specifies the maximum number of kilobytes (KB) that RMAN reads each second on this channel. This parameter sets an upper limit for bytes read so that RMAN does not consume too much disk bandwidth and degrade performance. ■ Size of backup piece (KB) (MAXPIECESIZE) specifies the maximum size of each backup piece that is created on this channel. ■ Number of open files (MAXOPENFILES) controls the maximum number of input files that the backup operation can have open at any given time. ■ Section size (SECTION SIZE) enables RMAN's multisection backup during an RMAN backup. When a multisection backup does run, an RMAN channel can process each file section independently (serially or in parallel). The multisection backup uses multiple channels to back up a single file. This option lets you set the Section size to either MB or GB. Size of backup piece (KB) and Section size are mutually exclusive. NetBackup automatically sets the Size of backup piece (KB) to 0 if this option is enabled. Maximum backup set limits parameters: ■ Number of files per backup set (FILESPERSET) specifies the maximum number of input files to include in each output backup set. ■ Size of the backup set (KB) (MAXSETSIZE) specifies a maximum size for a backup set in kilobytes.
Backup identifier options	Backup set identifier: (TAG) specifies a user-specified tag name for a backup set, proxy copy, data file copy, or control file copy. The tag is applied to the output files that the backup generates. Datafile copy tag: This option specifies a user-specified tag name when the Whole Database - Datafile Copy Share option is used. The tag is associated with the data files that are located on the appliance and is used during the incremental merge process.

Table 6-6 Oracle tab fields (*continued*)

Field	Description
Archived redo log options	Select Include archived redo logs in full and incremental schedules to include the archived redo logs on the full and the incremental schedule backups. Select Delete after making copies to delete the archived redo logs after the selected number of backups are successful. Clear the box or set to 0 to skip the delete operation and retain the logs after backup. Select Delete archived redo logs after (Hours) so that only the redo logs that are older than the selected value are deleted after a backup. Number of parallel streams controls the degree of parallelism within a backup. This number specifies the maximum number of connections between RMAN and a database. Each connection initiates an Oracle database server session on the target database. This server session performs the work of backing up backup sets. When an Oracle RAC is included in the policy, setting the Number of parallel streams option sets the streams per node. For example, if the Oracle RAC has two nodes that are active and the Number of parallel streams is set to 2, NetBackup uses four streams. NetBackup uses two parallel streams per active node. Specify Specify maximum limits to set custom limits for the archive redo logs. ■ Number of files per backup set specifies the maximum number of archived redo log files to include in each output backup set. ■ Size of backup set (KB) (MAXSETSIZE) specifies a maximum size for a backup set of archived redo logs on kilobytes.
User specified backup file name formats	Select Specify backup file name formats to set up formats for various backup file names for data files, archived redo logs, the control file, and Fast Recovery Area (FRA). Ensure that the format that is specified for all RMAN backup piece names (except for auto-backups of the control file) uses the <code>_%u</code> and ends with <code>_%t</code>. NetBackup uses this timestamp as part of its search criteria for catalog images. Without this timestamp, performance might degrade as the NetBackup catalog grows. These recommendations help to ensure proper backup, restore, and crosscheck functionality.

Table 6-6 Oracle tab fields (*continued*)

Field	Description
Database backup share options	These options let you set a time when backup sets and backup copies (data and control file copies) are automatically deleted from the appliance share. However, the files are only deleted if they have been successfully backed up from the share. The two options are Delete protected backup sets from share after and Delete protected backup copies from share after. The deletion is based on the age of the file in the share. The original dump time or the update time (if an incremental merge is done on the file) determines the age of the file. These options are only available when the Database backup shares option is selected in the Backup selections tab. Use the dropdowns to set the minutes, hours, days, or weeks. Note: The delete options in the Database backup share options can only be configured when the media server is a NetBackup appliance. This option does not work with any other type of media server. This feature requires a NetBackup appliance with software version 2.7.1 or later.

Create a Data Guard policy

Note: If you choose **Require primary**, **Require standby**, or **Prefer standby**, then the following applies:

Only one database is backed up at policy run time.

To create a Data Guard policy

- 1 Open the NetBackup web UI.
- 2 On the left click **Protection > Policies**. Then click **Add**.
- 3 Configure the **Attributes** tab.
- 4 Configure the **Instances and databases** tab.

All databases or RAC databases that are listed in a policy must be part of the same Data Guard configuration. Each Data Guard configuration needs a separate policy.

- 5 On the **Backup selections** tab, select the wanted option.

The policy cannot have the backup type **Database Backup Shares** or **Whole Database – Datafile Copy Share** selected. These selections are incompatible with Data Guard-aware policies.

- 6 Click the **Oracle** tab.

See [“Oracle tab”](#) on page 68.

7 Locate the **Data Guard backup options**.

The **Data Guard backup options** let you specify a policy to always back up the primary or a standby database:

- **None** - NetBackup does not look at the Data Guard role. At the policy run time each database or RAC database in the policy is backed up.
- **Require primary** - This option backs up the primary Data Guard database in the list of databases (or RAC databases) in the policy.
If the primary is not available, the backup fails. A message appears in the job details that indicates NetBackup cannot find the primary database.
- **Require standby** - This option backs up a standby Data Guard database in the list of databases (or RAC databases) in the policy. The first standby database NetBackup finds, is backed up. Currently there is no way to set a priority order for the standby databases.
Oracle Intelligent Policies do not support RMAN proxy type backups of Data Guard standby databases.
If the standby is not available, the backup fails. A message appears in the job details that indicates NetBackup cannot find the standby database.
- **Prefer standby** - This option backs up a standby Data Guard database in the list of databases (or RAC databases) in the policy. If a standby database is not found and the primary is available, then the primary database is backed up.
Oracle Intelligent Policies do not support RMAN proxy type backups of Data Guard standby databases.

8 Click **Create**.

Create `bpstart_notify` and `bpend_notify` scripts for an Oracle Intelligent Policy (OIP)

Optionally for an Oracle Intelligent Policy (OIP) you can create custom `bpstart_notify` and `bpend_notify` scripts. An OIP ignores the generic `bpstart_notify` and `bpend_notify` scripts. The scripts only work with the options **Protect non-RAC databases and pluggable databases** and **Protect RAC databases and pluggable databases**.

For more information about the `bpstart_notify` and `bpend_notify` scripts, refer to the [NetBackup Administrator's Guide, Volume II](#).

To create the scripts

1 Locate the generic scripts.

The generic `bpstart_notify` or `bpend_notify` scripts reside on the server. The location varies by platform:

- UNIX: `/usr/opensv/netbackup/bin/goodies/`
- Windows: `Install_path/NetBackup/bin/goodies/`

2 Copy and edit the scripts.

For an OIP, the scripts must include the `.policyname` or `.policyname.schedule` suffix or the script does not run at the start or end of the policy.

3 Rename the scripts.

Examples of renamed scripts:

UNIX:

```
bpstart_notify.oip_instance1  
bpend_notify.oip_instance2.full
```

Windows:

```
bpstart_notify.oip_instance1.bat  
bpend_notify.oip_instance2.full.bat
```

4 Place the scripts in the NetBackup bin directory on the client:

UNIX: `/usr/opensv/netbackup/bin/`

Windows: `install_path\NetBackup\bin\`

For RAC OIP backups the `bpstart_notify` and the `bpend_notify` scripts must reside on all Oracle RAC nodes.

Logging the RMAN input and output on a client

NetBackup has the ability to log the RMAN input and output that is logged locally on the client and also sent to the Activity Monitor. The `RMAN_OUTPUT_DIR` entry specifies which directory to place the RMAN input and output locally on the client for Oracle Intelligent Policy backups. The log is only created when a backup is run using an Oracle Intelligent Policy (OIP) and the file is continuously updated during the RMAN backup. Only one `RMAN_OUTPUT_DIR` entry per client is allowed in a Windows environment. In a UNIX environment, each user can place the output in a different location by adding the `RMAN_OUTPUT_DIR` entry to `$HOME/bp.conf` file. The value in the `$HOME/bp.conf` file takes precedence if it exists. NetBackup does not clean up the log files so the Oracle user has to clean up the log files manually.

You must use the `nbgetconfig` and the `nbsetconfig` commands to view, add, or change the option. The directory that is specified must exist and the Oracle user needs to have permission to create files within the directory.

Use the following format:

```
RMAN_OUTPUT_DIR = directory_name
```

The *directory_name* is a directory to which the Oracle user has permission to create files within the directory.

For information about these commands, see the [NetBackup Commands Reference Guide](#).

The file name has a specific format that includes the client name, policy name (OIP), schedule type, date stamp (yyyymmdd), and timestamp (hhmmss). The following is an example of how the file name looks in the directory:

```
orac121_backuppolycname_full_20160201_184157_GMT.log
```

The following are examples of `RMAN_OUTPUT_DIR` entires:

Windows: `install_path\oracle\oracle_logs\RMAN`

UNIX: `/oracle/oracle_logs/rman`

About script-based Oracle policies

NetBackup users or automatic schedules can start database backups by specifying a shell script in the file list of the Oracle policy. The shell script specifies the backup commands that RMAN performs on the client.

Note: All scripts must be stored and run locally. One recommendation is that scripts should not be world-writable. Scripts are not allowed to be run from network or remote locations. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

For more information about registering authorized locations and scripts, review the following article:

[Registering authorized locations used by a NetBackup database script-based policy](#)

On UNIX, NetBackup for Oracle includes a library of functions that enable RMAN to use NetBackup. You can link to this library.

See [“About linking Oracle RMAN with NetBackup for UNIX”](#) on page 21.

On Windows, NetBackup for Oracle includes a library of functions that enable RMAN to use NetBackup. This library is in `c:\Windows\system32`.

When you use the RMAN `backup` command, each resulting backup set contains at least one backup piece (data file, data file copy, control file, or archive log) from the target database. You must give each backup piece a unique name using the `format` operand. Several substitution variables are available to aid in generating unique names. You can specify the `format` operand in the `backup` command. NetBackup considers the backup piece name as the file being backed up, so this name must be unique in the catalog.

For a backup, the following items apply:

- The `rman` command starts the requested operation on the databases.
- When the process requires media to store backup data, RMAN starts a user-directed backup by issuing a backup request.
- The NetBackup media server connects to NetBackup for Oracle on the client. NetBackup for Oracle on the client sends the database data to the NetBackup media server which saves the data to secondary storage. A restore works in essentially the same manner except that RMAN issues a restore request. This request causes NetBackup to retrieve the data from secondary storage and send it to NetBackup for Oracle on the client.
- RMAN supports parallel operations, so a single `rman` command can start more than one backup, or restore on the NetBackup system.
- The status for an RMAN operation is stored in the RMAN catalog or in the database control file. This same status appears in the output of the RMAN command that is used to run the backup or restore. This status is the only status that a database administrator must check to verify that a backup or restore has been successful.
- NetBackup also logs status, but only for its own part of the operation. The database administrator cannot use the NetBackup status to determine whether `rman` was successful. Errors can occur in `rman` that do not affect NetBackup and are not recorded in its logs.

See [“About the environment variables set by NetBackup for Oracle”](#) on page 88.

See [“About the RMAN SEND command variables”](#) on page 89.

See [“About the bp.conf file on UNIX systems”](#) on page 92.

See [“About creating shell scripts”](#) on page 93.

Adding a new script-based Oracle policy

This topic describes how to add a new backup policy for a database.

To add a new script-based Oracle policy

- 1 Log on to the primary server as administrator (Windows) or root (UNIX), and open the web UI.
- 2 On the left, click **Protection > Policies**. Then click **Add**.
- 3 For the **Policy name**, type a unique name for the new policy.
- 4 For the **Policy type** list, select **Oracle**.
- 5 Click the **Instances and databases** tab and select **Clients for use with scripts**.

The tabs along the top of the dialog change to include a **Clients** tab.

- 6 Complete the entries on the **Attributes** tab.

See [“About policy attributes”](#) on page 79.

- 7 Add other policy information as follows:

- Add schedules.
 See [“Configuring NetBackup for Oracle automatic backup schedules”](#) on page 62.
 See [“Schedule properties for Oracle script-based policies ”](#) on page 80.

Note: Be aware of what may happen if an automatic schedule invokes a script that a user authored. NetBackup does not provide safeguards to prevent an automatic backup schedule from running a restore or a recovery script.

- Add clients.
 See [“Add clients to a policy”](#) on page 84.
- Add scripts to the backup selections list.
 See [“About adding backup selections to an Oracle policy”](#) on page 84.
 See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 319.

- 8 When you have added all the schedules, clients, and backup selections you need, click **Save**.

About policy attributes

With a few exceptions, NetBackup manages the policy attributes set for a database backup like a file system backup. Other policy attributes vary according to your specific backup strategy and system configuration.

For more information on policy attributes, see the [NetBackup Administrator's Guide, Volume I](#).

Table 6-7 Policy attributes for NetBackup for Oracle policies

Attribute	Description
Policy type	Determines the types of clients that can be backed up with the policy. For Oracle databases, select the policy type Oracle .
Limit jobs per policy	Sets the maximum number of databases that NetBackup can back up concurrently with this policy.
Follow NFS	This option is available for Oracle policies on UNIX. Select this attribute to back up the files from NFS-mounted file systems. If this option is not selected, NetBackup cannot perform a backup of NFS-mounted files. Also see the NetBackup Administrator's Guide, Volume I. Note: This option is not available for snapshot backups.
Keyword phrase	For NetBackup for Oracle, the Keyword phrase entry is ignored.
Snapshot Client and Replication Director	This group contains the options that enable backups with Snapshot Client and Replication Director.

Schedule properties for Oracle script-based policies

This topic describes the schedule properties that have a different meaning for database backups than for file system backups. Other schedule properties vary according to your specific backup strategy and system configuration. Additional information about other schedule properties is available. See the [NetBackup Administrator's Guide, Volume I](#).

Table 6-8 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure.

Table 6-8 Description of schedule properties (*continued*)

Property	Description
Retention	Clients for use with scripts: The retention period for an application backup type schedule refers to the length of time that NetBackup keeps backup images for stream-based backups. The retention period for an automatic backup type schedule refers to the length of time that NetBackup keeps backup images for proxy based backups. Oracle Intelligent Policy (OIP): The retention period for an automatic backup type schedule refers to the length of time that NetBackup keeps backup images for all backup types (stream-based, proxy, etc.) The Oracle database also has retention settings for backup images in the RMAN catalog. As a general recommendation, the NetBackup retention of a backup image should be longer than the database retention of the same backup image. The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be more than one week. The NetBackup scheduler compares the latest record of the automatic backup schedule to the frequency of that automatic backup schedule. This comparison is done to determine whether a backup is due. So if you set the retention period to expire the record too early, the scheduled backup frequency is unpredictable. However, if you set the retention period to be longer than necessary, the NetBackup catalog accumulates unnecessary records. Oracle is not notified when NetBackup expires a backup image. Use Oracle RMAN repository maintenance commands to periodically delete expired backup sets from the Oracle RMAN repository. Oracle XML export operations create archives for long-term storage and recovery. Set the retention level to a period of years or to infinity. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling.

Script-based policy - Storage and Retention

This topic describes storage and retention properties of the script-based policies. See the [NetBackup Administrator's Guide, Volume I](#).

Table 6-9 Storage and retention behavior

Property	Description
Policy is a snapshot type	If the policy is a snapshot type, the following are the possible scenarios of the retention behavior: ■ If a schedule has overridden the policy storage, the override storage on the schedule takes precedence over the policy storage. ■ If the policy storage is a snapshot SLP, the application schedule must override the policy storage. The storage that is specified on the application schedule may not be a snapshot SLP. ■ If the storage being used is not an SLP, the schedule determines the retention for the snapshot data. ■ If you use the policy storage unit as an SLP, the SLP determines the retention for the snapshot data. Streamed data is processed by using the application schedule. Snapshot data is processed by using the automatic schedule.
Policy is not a snapshot type	If the policy is not a snapshot type, the following are the possible scenarios of the retention behavior: ■ If the application schedule has overridden the policy storage, the override storage on the schedule takes precedence over the policy storage. ■ If the storage being used is not an SLP, the retention is derived from the schedule. ■ If the storage being used is an SLP, the retention is derived from the SLP. Since all data is streamed, the data is processed using the application schedule.

The following are examples of the script-based policy storage and retention behavior for snapshot-based policy types:

Policy storage	Application schedule storage	Full/Incremental schedule storage	Streamed data retention is derived from:	Snapshot data retention is derived from:
AdvancedDisk	-	-	Application Schedule	Full/Incremental Schedule
AdvancedDisk	-	Non-Snapshot SLP	Application Schedule	Non-Snapshot SLP
Non-Snapshot SLP	AdvancedDisk	-	Application Schedule	Non-Snapshot SLP
Tape library	Non-Snapshot SLP	-	Non-Snapshot SLP	Full/Incremental Schedule

Policy storage	Application schedule storage	Full/Incremental schedule storage	Streamed data retention is derived from:	Snapshot data retention is derived from:
Snapshot SLP	AdvancedDisk (must be specified)	Snapshot SLP	Application Schedule	Snapshot SLP on Full/Incremental Schedule
AdvancedDisk	-	Snapshot SLP	Application Schedule	Snapshot SLP
Non-Snapshot SLP	-	-	Non-Snapshot SLP	Non-Snapshot SLP
AdvancedDisk	Non-Snapshot SLP	Snapshot SLP	Non-Snapshot on Application Schedule	Snapshot SLP on Full/Incremental Schedule
Snapshot SLP	Non-Snapshot SLP (must be specified)	Snapshot SLP	Non-Snapshot on Application Schedule	Snapshot SLP on Full/Incremental Schedule

The following are examples of the script-based policy storage and retention behavior for stream-based policy types:

Policy storage	Schedule storage	Application schedule storage	Streamed data retention is derived from:
AdvancedDisk	-	N/A	Application Schedule
Non-Snapshot SLP	AdvancedDisk	N/A	Application Schedule
AdvancedDisk	Non-Snapshot SLP	N/A	Non-Snapshot SLP
Non-Snapshot SLP	-	N/A	Non-Snapshot SLP
AdvancedDisk	-	Non-Snapshot SLP	Non-Snapshot SLP
Snapshot SLP	-	AdvancedDisk	Application Schedule

See [“About NetBackup for Oracle schedule properties using Oracle Intelligent Policy”](#) on page 62.

Add clients to a policy

The client list contains a list of the clients on which your scripts are run during an automatic backup. This list determines the clients that can send backup requests to the application schedule. A NetBackup client must be in at least one policy but can be in more than one.

NetBackup attempts to run each script in the backup selections list for each client in the client list. If a script is not valid on a particular client, the script is skipped. (For example, if the Oracle home that is specified in the script does not exist on that client.) A policy can contain multiple clients and multiple scripts. Only a subset of the scripts needs to be valid on each client. If the valid scripts are successful, the entire backup is successful.

To add clients to a policy

- 1 Open the policy and select the **Clients** tab.
- 2 Select **Add**.
- 3 Type the name of the client and select the hardware and operating system of the client.
- 4 Select **Add**.

About adding backup selections to an Oracle policy

The backup selections list in a database policy has a different meaning than for non-database policies. For example, in a Standard or MS-Windows policy, the list contains files and directories to be backed up.

In a database policy, you specify scripts to be run.

Observe the following rules when you use scripts:

- Make sure that the scripts reside on each client in the client list.
- NetBackup installs sample scripts when you install the software; you can modify these scripts for your own use.
- All scripts must be in an authorized location.
See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 319.
- If you use NetBackup for Oracle in a NetBackup server cluster, make sure that the scripts reside in a location that is available after a failover.

Note: All scripts must be stored and run locally. One recommendation is that scripts should not be world-writable. Scripts are not allowed to be run from network or remote locations. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

For more information about registering authorized locations and scripts, review the knowledge base article:

[Registering authorized locations used by a NetBackup database script-based policy](#)

Add scripts to the backup selections list only if you want to set up a policy for automatic backups. These scripts are run for manual backups and for automatic schedules as specified under the **Schedules** tab. NetBackup runs the scripts in the order that they appear in the backup selections list.

See [“Adding a script to the backup selections list”](#) on page 85.

Adding a script to the backup selections list

The following procedure describes how to add a script to the backup selections list.

Note: Be sure to specify the correct script name in the backup selections list to prevent an error or a wrong operation.

To add a script to the backup selections list

- 1 Open the policy.
- 2 Select the **Backup selections** tab.
- 3 Select **Add**.
- 4 In the **Script** box, type the full path name of a script on the client.

For example:

```
/backup_scripts/db/cold_backup.sh  
C:\backup_scripts\db\cold_backup.cmd
```

See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 319.

- 5 Select **Add** to add the script to the list.
- 6 Select **Save**.

Note: Be aware of what may happen if an automatic schedule invokes a script that a user authored. NetBackup does not provide safeguards to prevent an automatic backup schedule from running a restore or a recovery script.

See [“Using the NetBackup for Oracle sample scripts”](#) on page 96.

About configuring the run-time environment

When the Oracle Intelligent Policy is not used, there are many user configurable variables that can affect the operation of NetBackup for Oracle. Most can be set only in an RMAN script. A few can be set in multiple places, most notably those that specify the primary server, client name, policy, and schedule.

When these variables are configured in multiple places, the following order of precedence is used (the list is ranked highest to lowest):

- RMAN SEND command variables, if specified in the backup script.
- RMAN ENV parameter variables, if specified in the backup script.
 See [“About the RMAN SEND command variables”](#) on page 89.
- Environment variables that are inherited from the Oracle listener process startup environment, if RMAN connects to the database using TNS SQL*Net.
- Environment variables that are set in the backup script before RMAN is started.
- The environment variables that the login or shell inherits.
 See [“About the Oracle RMAN environment in NetBackup for Oracle”](#) on page 87.
- The environment variables that the primary server initiation of an automatic schedule sets.
 See [“About the environment variables set by NetBackup for Oracle”](#) on page 88.
- On UNIX, the Oracle user’s configuration file; \$HOME/bp.conf.
- The NetBackup configuration:
 - Windows: The
 HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config
 registry keys.
 - UNIX: The /usr/opensv/netbackup/bp.conf file.
 - Both: The NetBackup GUI settings for server, client name, optional default policy, and optional default schedule.
 See [“About the bp.conf file on UNIX systems”](#) on page 92.
- The following defaults apply:
 - A server must be specified, there is no default.

- The client name defaults to the host name.
- The primary server selects the first policy of type Oracle for the client name.
- The primary server selects the first schedule of type Application Backup (stream-based) or Automatic Full Backup (proxy) from the policy.

See [“Preparing for NetBackup for Oracle configuration”](#) on page 54.

About the Oracle RMAN environment in NetBackup for Oracle

The Oracle RMAN program inherits the environment of the program or shell from which it was started. The environment may come from a number of places:

- The global environment or profile for the host
- The profile of the user
- The NetBackup primary server
- A non-NetBackup scheduler
- A backup script
- An interactive terminal session

Additionally, once RMAN is started it connects to the database and starts the Oracle database server processes that perform the backup. If the connection is by local logon and password (without a TNS alias), the Oracle database server process is a child of the RMAN program. The Oracle database server process inherits the environment from RMAN. Because the NetBackup for Oracle agent is a shared library loaded into the Oracle database server process it too inherits that environment.

However, if RMAN connects to the database by SQL*Net (logon and password@TNSalias) the Oracle database server process is a child of the SQL*Net listener service. This SQL*Net listener service was started previously and independently of RMAN. As a result, the NetBackup for Oracle agent does not inherit the environment from RMAN. Instead, the agent inherits the environment from which the listener service was started

To avoid unexpected results, it is recommended to configure RMAN to always use the `send` command to pass the desired variables and values to NetBackup explicitly. Alternatively the RMAN ENV parameter can be used to make the variables and values available to NetBackup.

Example 1. Use the `send` command to specify the policy and server to use for a database backup. As this example shows, specify the variables in the string in the

RMAN script after all channels have been allocated and before the `backup` command.

```
run {
    allocate channel t1 type 'SBT_TAPE';
    allocate channel t2 type 'SBT_TAPE';
    send 'NB_ORA_POLICY=your_policy,NB_ORA_SERV=your_server';
    backup (database format 'bk_%U_%t');
    release channel t1;
    release channel t2;
}
```

Example 2. Use the `parms` operand to specify the policy and server to use for a database backup. The `parms` operand is set with each `allocate channel` command in the shell script.

```
run {
    allocate channel t1 DEVICE TYPE 'SBT_TAPE'
        PARMS "SBT_LIBRARY=/usr/opensv/netbackup/bin/libobk.so,
        ENV=(NB_ORA_POLICY=your_policy,NB_ORA_SERV=your_server)";
    allocate channel t2 DEVICE TYPE 'SBT_TAPE'
        PARMS "SBT_LIBRARY=/usr/opensv/netbackup/bin/libobk.so,
        ENV=(NB_ORA_POLICY=your_policy,NB_ORA_SERV=your_server)";
    backup (database format 'bk_%s_%p_%t');
    release channel t1;
    release channel t2;
}
```

About the environment variables set by NetBackup for Oracle

When an automatic schedule runs, NetBackup sets environment variables for shell scripts to use. These variables are set only if the backup is started from the server, either automatically by the NetBackup scheduler or manually through the administrator interface.

On UNIX and Windows, these variables can be used to perform conditional operations within the backup script.

[Table 6-10](#) shows the variables.

Table 6-10 Variables that NetBackup for Oracle sets

Environment variable	Purpose
NB_ORA_SERV	Name of the NetBackup server that initiated the automatic schedule.
NB_ORA_POLICY	Name of the Oracle policy that contained the automatic schedule.
NB_ORA_CLIENT	Name of the NetBackup client in the policy.
NB_ORA_FULL	Set to 1 for a Full schedule.
NB_ORA_INCR	Set to 1 for a Differential incremental schedule.
NB_ORA_CINC	Set to 1 for a Cumulative incremental schedule.
NB_ORA_PC_SCHED	Name of the automatic schedule.

About the RMAN SEND command variables

The Oracle `SEND` command and the `ENV` parameter support several options that are used with NetBackup for Oracle. The variables that the `SEND` command specifies supersede those specified by the `ENV` parameter. Also, spaces are not permitted when the variables and values are specified.

[Table 6-11](#) describes the options you can set for the RMAN SEND command.

Table 6-11 Options for the SEND command

Option	Purpose
BKUP_IMAGE_PERM	Lets you set the permissions on a backup image at backup time. Possible values are the following: USER - set the permissions to 600. Only the original user who backed up the data has access to the backup images. GROUP - set the permissions to 660. Anyone from the same group as the original user who backed up the data has access to the backup images. ANY - set the permissions to 664. Anyone has access to the backup images. If this keyword is not specified, the permissions default to 660. To specify this keyword, use the send command to set the variable. For example: <pre>SEND 'BKUP_IMAGE_PERM=ANY';</pre> Note: The BKUP_IMAGE_PERM option does not affect the permissions for the physical files that are included in an RMAN Proxy copy backup. Ensure the physical file owner, group, and permissions are set correctly before the backup. For more information, review the following: Oracle RMAN PROXY restore from cataloged snapshot fails to match the existing piece name
NB_ORA_CLIENT	Specifies the name of the Oracle client.
NB_ORA_COPY_NUMBER	Specifies which copy of the backup image to use for the restore.
NB_ORA_PARENT_JOBID	Enables the parent ID of the job ID to be displayed in the Activity Monitor (only valid if it is a scheduled job).
NB_ORA_PC_RESTORE	Specifies a snapshot rollback restore using a script or RMAN command.
NB_ORA_PC_SCHED	Specifies the NetBackup for the Oracle schedule that NetBackup uses for a proxy copy file-based backup. This schedule can be Full, Differential Incremental, or Cumulative Incremental backup type. For scheduled backups, this variable is passed from the scheduler.

Table 6-11 Options for the SEND command (*continued*)

Option	Purpose
NB_ORA_PC_STREAMS	Specifies the number of backup streams that NetBackup starts simultaneously in each proxy copy session. When a backup is started, NetBackup groups all data files into a specified number of backup streams that are based on the file sizes. NetBackup tries to create streams of equal size. The default value for NB_ORA_PC_STREAMS is 1. This option also can be used to specify the number of restore streams that start simultaneously. For more information about restores, refer to: See “About Oracle multistream restore for proxy backup” on page 114.
NB_ORA_POLICY	Specifies the name of the policy to use for the Oracle backup.
NB_ORA_RESTORE_PRIORITY	Specifies the restore priority in NetBackup.
NB_ORA_SCHED	Specifies the name of the Application Backup schedule to use for the Oracle backup.
NB_ORA_SERV	Specifies the name of the NetBackup primary server.
NB_ORA_SERVER_READ_TIMEOUT	Configured to instruct the <code>dbclient</code> to lengthen or shorten the time-out on the media server. The media server uses this time out when it waits for a progress status update from the client during transfer of the backup image. Typically, this setting should not be adjusted. To review setting information and delay examples, refer to the following article: http://www.veritas.com/docs/TECH227741
NB_ORA_DISK_MEDIA_SERVER	Specifies which media server to use when more than one has access to the image to be restored. Supersedes any <code>FORCE_RESTORE_MEDIA_SERVER</code> setting on the primary server.
CPF1_POLICY	Policy to be used for duplex copy number 1.
CPF1_SCHED	Application backup schedule for duplex copy number 1.
CPF2_POLICY	Policy to be used for duplex copy number 2.
CPF2_SCHED	Application backup schedule for duplex copy number 2.

Table 6-11 Options for the SEND command (*continued*)

Option	Purpose
CPF3_POLICY	Policy to be used for duplex copy number 3.
CPF3_SCHED	Application backup schedule for duplex copy number 3.
CPF4_POLICY	Policy to be used for duplex copy number 4.
CPF4_SCHED	Application backup schedule for duplex copy number 4.

For more information, see the [NetBackup System Administrator's Guide, Volume I](#).

See [“About configuring the run-time environment”](#) on page 86.

See [“About the bp.conf file on UNIX systems”](#) on page 92.

See [“About the environment variables set by a user in the XML export parameter file”](#) on page 294.

See [“About the environment variables set by NetBackup for Oracle”](#) on page 88.

About the bp.conf file on UNIX systems

A NetBackup for Oracle user can create a `bp.conf` file in the Oracle user's home directory on the NetBackup for Oracle client host. When a NetBackup for Oracle operation is started, the user's `bp.conf` file is searched before the primary configuration file (`/usr/openv/netbackup/bp.conf`). Any option that is found at the user level overrides the same option's setting at the primary level.

[Table 6-12](#) shows the options that you can set in the user's `bp.conf` file.

Table 6-12 Options for the user bp.conf file

Option	Purpose
BPBACKUP_POLICY	This option specifies the name of the policy to use for the backup.
BPBACKUP_SCHED	This option specifies the name of the Application Backup type of schedule to use for the backup.
CLIENT_NAME	This option specifies the name of the Oracle client. This name is especially useful for a redirected restore operation.

Table 6-12 Options for the user `bp.conf` file (*continued*)

Option	Purpose
<code>CLIENT_READ_TIMEOUT</code>	Use this option to increase the number of seconds that the Oracle client initially waits for a response from the NetBackup server. The default is the greater of 900 or <code>CLIENT_READ_TIMEOUT</code> .
<code>SERVER</code>	This option specifies the name of the NetBackup primary server. There can only be one <code>SERVER</code> option in the user <code>bp.conf</code> file.
<code>VERBOSE</code>	This option causes NetBackup to include more information in its debug logs.

For more information, see the [NetBackup System Administrator's Guide, Volume I](#).

The following shows example `bp.conf` entries for an Oracle user:

```
SERVER=jupiter
CLIENT_READ_TIMEOUT=900
VERBOSE=1
```

See [“About configuring the run-time environment”](#) on page 86.

See [“About the environment variables set by a user in the XML export parameter file”](#) on page 294.

See [“About the environment variables set by NetBackup for Oracle”](#) on page 88.

See [“About the RMAN SEND command variables”](#) on page 89.

About creating shell scripts

RMAN scripts contain the commands that run NetBackup RMAN backup and recovery jobs. The scripts must be created before NetBackup can perform scheduled backups. The shell scripts are specified in policy configuration on the NetBackup server.

Note: The Backup, Archive, and Restore GUI cannot be used for performing Oracle backups and restores.

See [“About creating RMAN scripts manually”](#) on page 94.

RMAN shell scripts

You can use RMAN shell scripts with the NetBackup for Oracle agent. Shell scripts that the user writes must conform to RMAN and operating system shell syntax. Sample backup and recovery shell scripts are installed on the client with the NetBackup for Oracle agent. Modify these scripts to meet your individual requirements.

See [“About creating RMAN scripts manually”](#) on page 94.

About creating RMAN scripts manually

You can create RMAN scripts manually and when you create a script, you need to specify the type of backup and assign a name to the output file. Keep in mind the following considerations:

Backup type	RMAN supports the following different types of backups (In the examples, <i>n</i> must be 1 or higher): ■ BACKUP FULL ■ BACKUP INCREMENTAL LEVEL 0 (Full backup base for incremental backups) ■ BACKUP INCREMENTAL LEVEL <i>n</i> (Differential incremental backup) ■ BACKUP INCREMENTAL LEVEL <i>n</i> CUMULATIVE (Cumulative incremental backup)
-------------	--

When generating a data file backup set, you can make either an incremental backup or a full backup. Both a full backup and an incremental level 0 perform a complete backup of the data file. However, an incremental level 0 backup can be used as the base for incremental level *n* and or incremental level *n* cumulative backups.

File names	Observe the following with regard to file names: ■ Each output file must have a unique name. Use the %U format specifier to satisfy this restriction. %U is equivalent to %u_%p_%c, and it guarantees the uniqueness of the backup set name in all circumstances. ■ Put %t at the end of the backup file name format. NetBackup uses the timestamp as part of its search criteria for catalog images. Without this timestamp, performance might degrade as the NetBackup catalog grows. ■ Ensure that the format that is specified for all RMAN backup piece names does not contain any space characters.
------------	--

See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 319.

See [“About the NetBackup for Oracle sample scripts”](#) on page 95.

See [“Using the NetBackup for Oracle sample scripts”](#) on page 96.

See [“About the set duplex command”](#) on page 96.

About the NetBackup for Oracle sample scripts

When you install NetBackup for Oracle, there are some sample scripts that can be used as examples. You must modify any sample script you use to work within your environment. The NetBackup installation writes example scripts to the following directory:

Windows:

```
install_path\NetBackup\dbext\Oracle\samples\rman
```

UNIX:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman
```

The Oracle example scripts are as follows:

Windows:

```
cold_duplex_cluster_database_backup_full.cmd  
cold_cluster_database_backup.cmd  
cold_database_backup.cmd  
cold_pdb_backup.cmd  
cluster_database_restore.cmd  
complete_database_restore.cmd  
complete_pdb_restore.cmd  
hot_database_tablespace_backup_proxy.cmd  
hot_pdb_backup.cmd  
hot_database_backup.cmd  
hot_pdb_tablespace_backup_proxy.cmd  
pit_database_restore.cmd  
pit_cluster_database_restore.cmd  
pit_pdb_restore.cmd
```

UNIX:

```
cold_database_backup.sh  
cold_pdb_backup.sh  
hot_database_backup.sh  
hot_database_tablespace_backup_proxy.sh  
hot_pdb_backup.sh  
hot_pdb_tablespace_backup_proxy.sh
```

```
complete_database_restore.sh  
complete_pdb_restore.sh  
pit_database_restore.sh  
pit_pdb_restore.sh
```

See [“Using the NetBackup for Oracle sample scripts”](#) on page 96.

Using the NetBackup for Oracle sample scripts

The following procedure describes how to use the sample scripts to manually create your own script.

To use sample scripts provided by NetBackup for Oracle

- 1 Copy the example scripts to a different directory on your client. Oracle scripts can be located anywhere on the client.
- 2 Modify each script for your environment.
- 3 On UNIX, make sure the `su` command logs into the correct user.

If you do not include an `su - user` (`user` is Oracle administrator account) in your Oracle scripts, they do not run with the proper permissions and environment variables. Problems with your database backups and restores can arise.

See [“Adding a script to the backup selections list”](#) on page 85.

See [“About creating RMAN scripts manually”](#) on page 94.

About the set duplex command

RMAN provides an API that lets you make up to four backup sets simultaneously, each an exact duplicate of the others. Using NetBackup, for example, you can back up each copy to a different tape to protect against disaster, media damage, or human error. Use the `set duplex` and the `send` commands to take advantage of this feature.

The `set duplex` command specifies the number of copies of each backup piece to create. The `set duplex` command affects all channels that are allocated after you issue the command. It remains in effect until explicitly disabled or changed during the session. You cannot issue the `set duplex` command after allocating a channel.

The command syntax is:

```
set duplex = {ON | OFF | 1 | 2 | 3 | 4}
```

By default, duplex is `OFF` (a single backup set is produced). If you specify `ON`, it produces two identical backup sets.

Note that you must enable the `BACKUP_TAPE_IO_SLAVES` initialization parameter to perform duplexed backups. RMAN configures all media as needed for the number of backup copies you request. For more information on `BACKUP_TAPE_IO_SLAVES`, see your Oracle documentation.

Use the `send` command to specify the policy and schedule to use with each backup. Because NetBackup uses the policy or schedule to determine what media to use, this information is required for each copy, or an error occurs.

The command syntax is as follows:

```
send 'keyword=value [, keyword=value,...]';
```

The keywords that are used to specify a policy are `CPF1_POLICY`, `CPF2_POLICY`, `CPF3_POLICY`, and `CPF4_POLICY`, which specify the backup policy for duplexed file 1 through duplexed file 4.

The keywords that are used to specify a schedule are `CPF1_SCHED`, `CPF2_SCHED`, `CPF3_SCHED`, and `CPF4_SCHED`, which specify the Application Backup schedule for duplexed file 1 through duplexed file 4.

About storing shell scripts

Shell scripts must reside on the NetBackup client. Backup shell scripts are associated with a policy by specifying the file name (including path) in the policy backup selections list. For server-directed or scheduled backups, each client in the policy's client list must have a copy of the script with the same name in the same location.

See [“About adding backup selections to an Oracle policy”](#) on page 84.

The backup and the recovery process sometimes require passwords for Oracle database access or system user accounts. Because a shell interprets the shell scripts, store the passwords in clear text.

See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 319.

Configuring the logon account for the NetBackup Client Service for NetBackup for Oracle

This topic applies to those that are running NetBackup for Oracle on a Windows platform.

Because the NetBackup Client Service is started by default under the `SYSTEM` account, you must also give special attention to database user authentication. The

`SYSTEM` account does not have permission to connect to the target database if you use OS authentication instead of passwords.

If you use OS authentication, run the NetBackup client service under an account that has SYSDBA privileges.

For more information on OS authentication, see your Oracle documentation.

Note: In a cluster environment, perform the steps on each database node in the cluster. For an off-host backup, perform the steps on the alternate client.

To configure the logon account for the NetBackup Client Service for NetBackup for Oracle

- 1 Open the Windows Services application.
- 2 Double-click the **NetBackup Client Service** entry.
- 3 Click the **Log On** tab.
- 4 Type the account name with SYSDBA privileges.
- 5 Type the password.
- 6 Click **OK**.
- 7 Stop and start the NetBackup Client Service.
- 8 Close the Services control panel application.

Testing configuration settings for NetBackup for Oracle

After you configure the servers and clients in your environment, test the configuration settings. Perform a manual backup (or backups) with the automatic backup schedules you created. A description of status codes and other troubleshooting information is available.

See the [NetBackup Status Codes Reference Guide](#).

See the [NetBackup Troubleshooting Guide](#).

To test the configuration settings

- 1 Log on to the primary server as administrator (Windows) or root (UNIX).
- 2 Sign in to the NetBackup web UI.
- 3 In the left pane, click **Protection > Policies**.
- 4 Select the policy that you want to test.

- 5** Click **Manual backup**.
- 6** Select the schedule that you want to use and the clients or instances that you want to back up.
- 7** Click **Backup**.
- 8** To check the status of the backup, click **Activity monitor**.

Performing backups and restores of Oracle

This chapter includes the following topics:

- [Overview of using NetBackup for Oracle](#)
- [Maintaining the RMAN repository](#)
- [Querying the RMAN repository](#)
- [About NetBackup for Oracle backups](#)
- [Browsing backups using the bplist command](#)
- [Managing expired backup images](#)
- [About NetBackup for Oracle restores](#)
- [Using NetBackup for Oracle in a Windows Server Failover Cluster \(WSFC\)](#)

Overview of using NetBackup for Oracle

The NetBackup graphical user interfaces and command line interfaces let you perform Oracle backup and recovery operations using Oracle RMAN utilities. You can also use the Oracle Enterprise Manager to perform Oracle backup and recovery operations. The Oracle RMAN command line interface is also used to maintain and query the RMAN repository.

See [“About NetBackup for Oracle backups”](#) on page 106.

See [“Maintaining the RMAN repository”](#) on page 101.

See [“Querying the RMAN repository”](#) on page 105.

See [“About creating RMAN scripts manually”](#) on page 94.

Maintaining the RMAN repository

The RMAN repository is the collection of metadata about your target databases that RMAN uses to conduct its backup, recovery, and maintenance operations. You can either create a recovery catalog in which to store this information or let RMAN store it exclusively in the target database control file. Although RMAN can conduct all major backup and recovery operations using only the control file, some RMAN commands function only when you use a recovery catalog.

[Table 7-1](#) shows the tasks that are required to maintain the RMAN repository and a subset of the repository maintenance commands that perform the tasks. Some of these commands might not be available with all versions of RMAN.

Table 7-1 Tasks and commands

Task	Commands that perform the task
Register a database with the recovery catalog	Before using RMAN with a recovery catalog, register the target database in the recovery catalog. To register, start and mount the target database but do not open it. At the RMAN prompt, issue a <code>register database</code> command.
Reset the incarnation in the recovery catalog	The <code>reset database</code> command directs RMAN to create a new database incarnation record in the recovery catalog.

Table 7-1 Tasks and commands (*continued*)

Task	Commands that perform the task
Crosscheck the information in the RMAN repository	Because NetBackup can expire images independently from Oracle, the RMAN repository can contain outdated information. Run an RMAN crosscheck to ensure that data in the recovery catalog or control file is in sync with data in the backup image catalog. The crosscheck queries NetBackup for the existence of each backup piece and then marks it as available or expired in the RMAN repository. Use one of the following commands to check the specified files. You need to run separate commands to delete images or repository records. ■ The <code>change...crosscheck</code> command queries NetBackup to determine if a backup piece is available. If not, RMAN marks the backup piece as expired. If it was expired but is now available, RMAN marks the backup piece as available. The command syntax is as follows: <pre>change backuppiece {primary_keylist filename_list tag} crosscheck; change backupset {primary_keylist} crosscheck;</pre> ■ The <code>crosscheck backupset</code> command operates on available and expired backup pieces. RMAN updates their status with the result (available or expired). ■ To crosscheck a database, start RMAN and connect to the target database and to the recovery catalog (if used). At the <code>rman</code> command prompt, enter the following: <pre>allocate channel for maintenance type 'SBT_TAPE'; crosscheck backupset of database;</pre> The length of time to perform an RMAN crosscheck depends on several factors: ■ Number of RMAN backup pieces being crosschecked. ■ Number of RMAN backup pieces past their NetBackup retention period when NetBackup expires them, not RMAN. ■ Format of the RMAN backup piece name and if the Cohesity recommended <code>_%t</code> appears at the end of the format statement. ■ Number of Oracle clients. ■ Number of NetBackup policies of any kind. ■ Length of time NetBackup retains backups and the number of backup images for the client in the NetBackup catalog. ■ Scheduling time and the length of time between RMAN catalog maintenance operations. ■ Speed and accuracy of host name and reverse host name resolution on the NetBackup primary server. ■ Number and complexity of the operations that the NetBackup primary server performs during each crosscheck request. ■ Normal performance.

Table 7-1 Tasks and commands (*continued*)

Task	Commands that perform the task
Crosscheck using the Oracle Copilot share	If files on an Oracle Copilot share are deleted outside of RMAN, the subsequent incremental merge backups that are done to the share fail. An RMAN crosscheck of the share must be done before the next backup to prevent more failures. This version of the RMAN crosscheck is slightly different from the other crosscheck examples because of the need to specify <code>type disk</code> instead of <code>type SBT_TAPE</code>. When running the RMAN crosscheck, the default is the <i>NetBackup_policyname</i>. However, if the Datafile copy tag is changed in the Oracle tab, then that tag name must be used in place of <i>NetBackup_policyname</i>. An example of the command syntax follows (using the default <i>NetBackup_policyname</i>): <pre>Run { Allocate channel ch00 type 'disk'; crosscheck backup tag <Netbackup_policyname>; delete noprompt expired backup; crosscheck copy <Netbackup_policyname>; delete noprompt expired copy; release channel ch00; }</pre>
Delete obsolete backups	The <code>DELETE OBSOLETE</code> command deletes the backups that are no longer needed to satisfy specified recoverability requirements. You can delete obsolete pieces according to the configured default retention policy, or another retention policy that a <code>DELETE OBSOLETE</code> option specifies. As with other forms of the <code>DELETE</code> command, the deleted files are removed from the backup media (expired from NetBackup). Then they are deleted from the recovery catalog, and marked as <code>DELETED</code> in the control file. If you specify the <code>DELETE OBSOLETE</code> command with no arguments, then RMAN deletes all the obsolete backups that the currently configured retention policy defines. For example: <pre>Allocate channel for maintenance type 'SBT_TAPE'; DELETE OBSOLETE;</pre> You can also use the <code>REDUNDANCY</code> or <code>RECOVERY WINDOW</code> clauses with <code>DELETE</code> to delete the backups that are obsolete under a specific retention policy instead of the configured default: <pre>DELETE OBSOLETE REDUNDANCY = 3; DELETE OBSOLETE RECOVERY WINDOW OR 7 DAYS;</pre>

Table 7-1 Tasks and commands (*continued*)

Task	Commands that perform the task
Delete expired backups	The <code>delete expired backupset</code> command operates only on the expired backup pieces that are found in the recovery catalog. RMAN removes them from the recovery catalog and also from the backup media (expires them from NetBackup). To delete expired backup sets of a database from the recovery catalog, start RMAN and connect to the target and the recovery catalog databases. At the RMAN command prompt, type the following commands: <pre>allocate channel for maintenance type 'SBT_TAPE'; delete expired backupset of database;</pre> The <code>crosscheck</code> and <code>delete backupset</code> commands restrict the list of objects to only those that are operated on. The restrictions are placed on the specified Oracle device type (disk or SBT tape), object type (archived logs or database files), and date range.
Resynchronize the recovery catalog	RMAN compares the recovery catalog to either the current control file of the target database or a backup control file. It subsequently updates the catalog with the missing information or changed information. If you are running in <code>ARCHIVELOG</code> mode, do the following: Resynchronize the recovery catalog regularly because the recovery catalog is not updated automatically when a log switch occurs or when a redo log is archived. You must also resynchronize the recovery catalog after making any change to the physical structure of the target database. As with log archive operations, the recovery catalog is not automatically updated when a physical schema change is made. The RMAN <code>backup</code>, <code>copy</code>, <code>restore</code>, and <code>switch</code> commands update the recovery catalog automatically when the target database control file is available. The recovery catalog database is available when one of these commands is executed. If the recovery catalog is unavailable when you issue <code>backup</code> or <code>copy</code> commands, you should resynchronize it manually. To resynchronize the recovery catalog, start RMAN and issue the <code>resync catalog</code> command.

Table 7-1 Tasks and commands (*continued*)

Task	Commands that perform the task
Change the availability of a backup set or file copy	Periodically, you might need to notify RMAN that the status of a backup set, backup piece, data file copy, or archived redo log has changed. The RMAN <code>change</code> command enables you to make a variety of useful record changes. The <code>change ... uncatalog</code> command removes references to a backup piece, data file copy, or archive log from the recovery catalog. This command works only with a recovery catalog. The <code>change ... delete</code> command removes references to a backup piece, data file copy, or archive log from the control file and recovery catalog. It physically deletes the file. This command works with or without a recovery catalog. The <code>change ... crosscheck</code> command removes references to a backup piece, data file copy, or archive log from the control file and recovery catalog. The references are removed when that file no longer exists. This command works with or without a recovery catalog. The <code>change ... unavailable</code> command marks a backup piece, data file copy, or archive log as unavailable. This command works only with a recovery catalog.
Validate the restore of backups	A restore validation retrieves the backup pieces from storage (NetBackup) and checks that the retrieved pieces are intact. But the restore validation discards the backup pieces without saving the contents into the database. Use <code>restore ... validate</code> when you want RMAN to choose the backups to test. Use <code>validate backupset</code> when you want to specify the backup sets to test.

See [“Overview of using NetBackup for Oracle”](#) on page 100.

See [“Querying the RMAN repository”](#) on page 105.

See [“About the Oracle RMAN environment in NetBackup for Oracle”](#) on page 87.

Querying the RMAN repository

RMAN lets you generate a number of reports relevant for backup and recovery using the `report` and `list` commands. The `list` command lists the contents of the recovery catalog or control file, and the `report` command performs a more detailed analysis.

Use the `report` and `list` commands to determine what you have backed up and what you need to back up. The information is available whether or not you use a recovery catalog.

You can use the `report` command to answer many different questions.

Some examples are as follows:

- Which files need a backup?
- Which files have not had been backed up in awhile?
- Which files are not recoverable due to unrecoverable operations?
- Which backup files can be deleted?
- What was the physical schema of the database at some previous point in time?

The `list` command queries the recovery catalog and control file and produces a listing of its contents. The primary purpose of the `list` command is to determine the backups that are available.

You can list the following information:

- Backup sets containing a backup of a specified list of data files.
- Backup sets containing a backup of any data file that is a member of a specified list of tablespaces.
- All backup sets or copies of all data files in the database.
- Backup sets containing a backup of any archive logs with a specified name or within a specified range.
- Incarnations of a specified database or of all databases that are known to the recovery catalog.

For more information on querying the RMAN repository, see your Oracle documentation.

About NetBackup for Oracle backups

You can perform different types of backups using NetBackup. Backups can be run automatically by using the schedules that you determine. Or, you can run a backup manually. The following table describes these methods of running a backup.

Automatic backups with OIP policies	The Oracle Intelligent Policy (OIP) feature is a method of Oracle policy backup based on Oracle databases or RAC databases. This method precludes the need to create a script for your Oracle policies.
-------------------------------------	---

See [“About Oracle Intelligent Policies \(OIP\)”](#) on page 59.

Automatic backups with classic policies and scripts	When the NetBackup scheduler invokes a schedule for an automatic backup, the NetBackup for Oracle backup shell scripts run as follows: ■ In the same order as they appear in the file list ■ On all clients in the client list The NetBackup for Oracle backup shell scripts start the database backup by running the <code>rman</code> command. When the backup is started through NetBackup, RMAN performs error checking. The <code>rman</code> command generates an error if it considers a command invalid, but it allows any of the commands it typically considers valid to proceed. When you specify the wrong script file name, you can start an unintended operation.
Manual backups	You can use the NetBackup web UI to manually run an automatic backup schedule for the Oracle policy. See “Testing configuration settings for NetBackup for Oracle” on page 98.

Running the NetBackup for Oracle shell script

When you run a NetBackup for Oracle shell script on a client to initiate a backup from the command prompt, specify the full path name to the file that contains the script. For example:

Windows:

```
install_path\oracle\scripts\db_full_backup.cmd
```

UNIX:

```
/oracle/scripts/db_full_backup.sh
```

The shell starts the database backup by running the Oracle shell script. The Oracle shell script contains commands to run `rman`.

The NetBackup installation script installs sample scripts in the following location:

Windows:

```
install_path\NetBackup\dbext\oracle\samples\rman
```

UNIX:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman
```

Running RMAN

As an Oracle user, you can run the `rman` command from the command prompt with the RMAN command file as a parameter. This topic describes how to set the primary server to `hag` and the Oracle policy to `obk` before you start the backup.

On Windows, RMAN functionality runs as a service, so use the `send` operand to set up the run-time environment. To start a backup using the `rman` command from the command prompt, type the following:

```
# send " 'NB_ORA_POLICY=obk,NB_ORA_SERV=hag' " cmdfile \
"install_path\oracle\scripts\db_full_backup.rcv"
```

On UNIX, type the following at the command prompt:

```
# rman target 'internal/oracle@ORCL' rcvcat 'rman/rman@RCAT'
# send " 'NB_ORA_POLICY=obk,NB_ORA_SERV=hag' " cmdfile \
'/oracle/scripts/db_full_backup.rcv'
```

If you intend to connect to a database using a TNS alias, the RMAN `send` command specifies the environment variables. The example sets the primary server to `hag` and the Oracle policy to `obk` before you start the backup.

See [“About the bp.conf file on UNIX systems”](#) on page 92.

Note: To run script files for database operations other than backups or restores, it is recommended that you run the `rman` command directly rather than NetBackup.

For `rman` command script syntax and examples, see your Oracle documentation.

Browsing backups using the bplist command

You can use the `bplist` command to browse Oracle backups. The command returns a list of backup file names.

Before using this command, log as follows into either the primary server or to the client:

- On Windows, log on as administrator to the primary server and to the client with the appropriate *altnames* entry.
- On UNIX and Linux, log on as root to the primary server and to the client with the appropriate *altnames* entry.

The following example uses the command to search all Oracle backups for a client named `jupiter`:

```
# bplist -C jupiter -t 4 -R /  
  
/exb_n2bm5bco_1_1392342936  
/exb_mabm02ko_1_1392170136  
/exb_lqbltds6_1_1392083334
```

The `-t 4` on this command specifies the Oracle backups. The `-R` specifies the default number (999) of directory levels to search.

For more information on the `bplist` command, see the [NetBackup Commands Reference Guide](#).

You can also use the `RMAN report` and `list` commands to browse Oracle backups.

See [“Querying the RMAN repository”](#) on page 105.

See [“Browsing XML export archives using bplist”](#) on page 304.

Managing expired backup images

NetBackup and Oracle each maintain a repository of RMAN-initiated backup image information. The retention setting in the Application Backup schedule for RMAN stream-based backups determines the NetBackup image retention. But for RMAN proxy backups and OIP backups, the retention setting on the Automatic Backup schedule determines retention of the NetBackup image.

To manage expired backup images from the NetBackup repository, access the Retention setting of the Application backup schedule. Specify the length of time before NetBackup expires a backup image.

See [“Schedule properties for Oracle script-based policies”](#) on page 80.

You can also manage the expired backup images from the Oracle repository. This method sets the backup retention as an RMAN attribute, rather than a NetBackup attribute. RMAN deletes the obsolete but not the unexpired backups from NetBackup. The following items are also part of this process:

- Set the NetBackup backup retention for Oracle backups to be either infinite or significantly longer than the RMAN retention.
- Set the RMAN retention to the number or duration to keep the backup sets in the RMAN catalog. If no RMAN catalog exists, then use SQL to set an appropriate value for `"control_file_record_keep_time"`. The minimum appropriate time is the catalog backup retention time plus the maximum time between catalog maintenance operations.
- On a regular basis, run the `RMAN delete obsolete` command to expire obsolete images from the RMAN catalog, the control file, and from NetBackup.

- If a cross-check of the catalog is required, perform the cross-check after RMAN deletes the obsolete backups.
- Stagger the initiation of RMAN catalog maintenance functions. Staggering is done to limit the number of concurrent checks or deletion requests that RMAN makes of the NetBackup primary server.
- Perform the RMAN catalog maintenance functions on a more frequent basis to limit the number of NetBackup catalog requests in a single session.
- Ensure that the format that is specified for all RMAN backup piece names (except for autobackups of the control file) ends with `_%t`.
- Ensure that the format that is specified for all RMAN backup piece names does not contain any space characters.
- Avoid the creation of excessive, small backup pieces of database files or archive logs.

You can manually remove references to backup images from the Oracle RMAN repository. Use RMAN repository maintenance commands to remove references to backup files. You can use these commands to delete backup image information from both the Oracle RMAN repository and the NetBackup repository.

More information is available on the RMAN repository maintenance commands.

See [“Maintaining the RMAN repository”](#) on page 101.

When a request is issued to delete a backup file from the RMAN repository, RMAN sends the request to NetBackup. The request tells NetBackup to delete the corresponding image from the NetBackup repository, regardless of the retention level.

About preventing the direct expiration of backup images

Catalog maintenance operations on Oracle send requests into NetBackup to synchronize the database catalog with the NetBackup catalog. As part of the catalog synchronization, the database may initiate an image expiration (delete) request to the NetBackup catalog. These requests may also come from the DBA when command-line options are used. For compliance reasons, you may want to prevent the expiration of images in the NetBackup catalog from a database request by using a `bp.conf` entry on the primary server.

To prevent the expiration of backup images, use the following `bp.conf` entry on the primary server:

PREVENT_ORACLE_DIRECT_EXPIRE

YES: This setting prevents the image delete requests from the database. The delete request receives a status code of 1420 and the log message indicates that the image is on `LEGAL HOLD`.

NO: The default setting. All image delete requests are honored from the database.

Note: The normal image expiration (retention) and the `bpexpdate` command are unaffected by this setting.

In a clustered primary server environment, these settings should be set and match in all the primary server `bp.conf` files.

The following is an example of a log entry when a status code of 1420 is produced.

The `dbclient` log message for Oracle:

```
Image is on LEGAL HOLD - Failed to remove, %s, from image catalog.
```

About NetBackup for Oracle restores

Make sure that a backup has completed successfully before you attempt a restore.

A recovery can be a complex process so it might be necessary to perform manual steps as part of the operation. For more information, see your Oracle documentation.

About the Oracle recovery shell script on the client

You can initiate a database recovery from the command prompt by typing the full path to the shell script that performs an Oracle recovery. For example:

Windows:

```
install_path\oracle\scripts\database_restore.cmd
```

UNIX:

```
/oracle/scripts/database_restore.sh
```

The operating system shell starts the database restore by running the Oracle shell script file. The Oracle shell script file contains commands to run `RMAN`.

The NetBackup installation script writes sample scripts to the following location:

Windows:

```
install_path\Netbackup\dbext\oracle\samples\rman\
```

UNIX:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman
```

See [“Running RMAN on the client”](#) on page 113.

See [“Redirecting a restore to a different client”](#) on page 115.

See [“About creating shell scripts”](#) on page 93.

About catalog names with RAC and Data Guard aware intelligent policies

Oracle Data Guard and RAC aware intelligent policies catalog the backup images under a special client name within the NetBackup catalog. Data Guard enabled policies use the format `databaseName_databaseID` as the client name. Database assets that are associated with a Data Guard asset are cataloged under the Data Guard catalog name. RAC aware policies use the format `databaseUniqueName_databaseID`. If a RAC database is in a Data Guard enabled policy, the catalog name uses the Data Guard catalog name.

Access to RAC database backup images

A RAC database configuration within NetBackup consists of a RAC database entry along with the associated instances. Each instance has a single client (host) name that can access the backup images that NetBackup stores. RAC databases and the associated instances can be viewed and managed using the NetBackup web UI. They can also be viewed and managed with the NetBackup APIs and the `nboraclmd` command.

Access to Data Guard backup images

A Data Guard configuration in NetBackup consists of a Data Guard object with attributes of database name and database ID. The list of Data Guard configurations that are stored in the NetBackup relational database can be viewed using the NetBackup APIs and the `nboraclmd` command.

RAC databases and single instance databases can be associated with a single Data Guard configuration. For Data Guard enabled policies, the Data Guard configuration in NetBackup is added or updated automatically at policy run time. Only RAC databases and single instance databases are added to a Data Guard configuration after a back up has occurred. Any client that is part of a RAC database configuration or single instance databases has access to the associated Data Guard backup images. Username and user IDs along with group name and group IDs

should match across hosts in the Data Guard configuration to have access to the backed-up files.

See the [NetBackup Commands Reference Guide](#) for details of the `nboracmd` command.

Access to Auto Image Replication (A.I.R.) backup images

If restoring from an alternate NetBackup server domain (A.I.R. domain target), for access to images that relate to Data Guard or RAC, the primary server must have certain conditions set. That server must have the same Oracle repository configuration duplicated into it to allow access from alternate nodes in the RAC cluster. Or, the server must have an alternate standby database in the Data Guard. The Data Guard is automatically recreated on import. The RMAN catalogs must be manually recreated along with all the CMS credentials that any registered databases use.

About Oracle Clients with multiple network interfaces

Additional host names can be added as aliases to an instance within NetBackup. This process is useful if a client is part of a RAC cluster or Data Guard configuration and has multiple network interfaces used to connect to NetBackup servers. Aliases can be listed, added, or deleted with the NetBackup APIs.

Running RMAN on the client

You can run the `rman` command from a command prompt on the client. Use the appropriate RMAN command file as a parameter.

On UNIX, the following example assumes that you are logged on as an Oracle administrator.

To run the RMAN command on the client:

- ◆ At the command prompt, type the following:

```
Windows: rman target 'internal/oracle@ORCL' rcvcat 'rman\rman@RCAT'  
cmdfile 'install_path\oracle\scripts\database_restore.rcv'
```

```
UNIX: rman target 'internal/oracle@ORCL' rcvcat 'rman/rman@RCAT'  
cmdfile '/oracle/scripts/database_restore.rcv'
```

See [“About the Oracle recovery shell script on the client”](#) on page 111.

See [“Redirecting a restore to a different client”](#) on page 115.

See [“Running RMAN”](#) on page 108.

About Oracle multistream restore for proxy backup

NetBackup lets you specify the number of restore streams that start simultaneously when the RMAN command is used. You can use the SEND command variable `NB_ORA_PC_STREAMS` or the RMAN `ENV` parameter to specify the number of restore streams. When you send the request to NetBackup, there may not be the same number of streams or jobs running during restore. NetBackup adjusts the stream count based on the count that is specified when you use the `NB_ORA_PC_STREAMS` variable. Or, NetBackup uses the number of images the requested restore job needs if the restore job needs more than one image. NetBackup selects whichever is the minimum number needed to complete the restore job.

When the restore job needs only one image and it is a snapshot, the stream count is based on the count that you specify in `NB_ORA_PC_STREAMS`. Or, NetBackup uses the number of files the requested restore job needs to complete. NetBackup selects whichever is the minimum number needed to complete the restore job. Also, the files are evenly distributed across the streams based on the file size.

When the restore needs only one image and that image is not a snapshot, then NetBackup does not attempt to perform a multistream restore.

See [“About the RMAN SEND command variables”](#) on page 89.

When the multistream restore is started, a parent job is created that initiates a child job for each stream. If you cancel the parent job, all incomplete child jobs are canceled and the job exits with a status of 150. If one of the child jobs is successful before parent cancellation, then the parent job exits with a status of 1. If you cancel one of the running child jobs, the child exits with status 150 and the parent job exits with a status of 1.

Note: Multistream restore only works when using Oracle backup images and is only accessible using command-line inputs.

Multistream restore supports the following snapshot method images:

- `remote_vxfs`
- `VxFS_Checkpoint`
- `VxVM`

Multistream restore is not supported when using the following types of images:

- Block level incremental images
- Off-host supported snapshot method images

It is recommended to configure RMAN to always use the SEND command to pass the desired variables and values to NetBackup explicitly. Alternatively the RMAN

ENV parameter can be used to make the variables and values available to NetBackup. The following are examples of running the multistream restore:

Example 1. Use the SEND command to specify the NB_ORA_PC_STREAMS variable.

```
RUN {  
  ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';  
  SEND 'NB_ORA_PC_STREAMS=<number of restore streams>';  
  RESTORE DATABASE; RECOVER DATABASE;  
  RELEASE CHANNEL ch00;  
}
```

Example 2. Use the PARMS operand to specify the NB_ORA_PC_STREAMS variable.

```
RUN {  
  ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE'  
  PARMS "ENV=(NB_ORA_PC_STREAMS= <number of restore streams>)" ;  
  RESTORE DATABASE; RECOVER DATABASE;  
  RELEASE CHANNEL ch00;  
}
```


See [“Running RMAN”](#) on page 108.

See [“About NetBackup for Oracle restores”](#) on page 111.

Redirecting a restore to a different client

With NetBackup for Oracle you have the option to restore a database to a client other than the one that originally performed the backup. The process of restoring data to another client is called a redirected restore.

[Figure 7-1](#) shows a redirected restore.

Figure 7-1 Redirected restore

The user on client A cannot initiate a redirected restore to client B. Only the user on client B, which is the client receiving the backup image, can initiate the redirected restore. Any user who belongs to the database group that performed the backup can restore it, unless the `BKUP_IMAGE_PERM` variable was set to `USER`.

See [“About performing a redirected restore with RMAN”](#) on page 117.

See [“Example - Performing a redirected restore of Oracle”](#) on page 118.

Preparing the primary server for an alternate restore

The examples in the following procedure assume that the database instance `ORAC11` was backed up by `client2`, and you want to restore `ORAC11` to `client1`.

For more information on how to manage client restores, see the [NetBackup Administrator's Guide, Volume I](#).

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its group is `svgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

To prepare the NetBackup primary server for alternate restores

- 1 Log on to the NetBackup primary server that hosts the policy that backed up database instance `ORAC11`.
- 2 Create a `dest_client_name` file on the NetBackup primary server.
 - Windows: `install_path\NetBackup\db\altnames\dest_client_name`
 - UNIX: `/usr/opensv/netbackup/db/altnames/dest_client_name`

Where `dest_client_name` is the name of a client that is allowed to be a destination client for alternate restores. For example, `client1`.

- 3 After creating a `dest_client_name` file, add the name of the NetBackup for Oracle source client to the `dest_client_name` file. For example, add the following line to this file:

```
client2
```

For more information on managing a client restore, see the [NetBackup Administrator's Guide, Volume I](#).

About performing a redirected restore with RMAN

Perform the following procedure on the destination client host if you want to restore any RMAN backups that another client owns.

The user on client A cannot initiate a redirected restore to client B. Only the user on client B, which is the client receiving the backup image, can initiate the redirected restore. Any user who belongs to the database group that performed the backup can restore it, unless the `BKUP_IMAGE_PERM` variable is set to `USER`.

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its group is `svgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

Note: If the RMAN catalog database has been lost, restore the catalog database first before continuing with the redirected restore.

To perform a redirected restore

- 1 Enable a network connection to the RMAN catalog database that the source client used.
- 2 Do one of the following:
 - On Windows, use the `rman parms` option to set the `NB_ORA_CLIENT` environment variable to the source client.
 - On UNIX, set the `NB_ORA_CLIENT` environment variable to the source client.
- 3 On UNIX, check the `bp.conf` files on the source client. Make sure that the `CLIENT_NAME` variable either is not set or is set to the host name of the source client.
- 4 Make the `init.ora` file of the source client available to the destination client.
Copy the file to the destination client or modify the file on the destination client.
Change all location-specific parameters.
- 5 Create a folder or set the permissions for a directory to restore the data files:
 - On Windows, create and start an Oracle service for the previously set `ORACLE_SID`. Create the folder to which you want to restore the data files.
 - On UNIX, grant write permission to the directory to which you want to restore the data files.
- 6 Set up a password file for the destination client database.
- 7 Start the database in the `nomount` state.
- 8 Start RMAN, connecting to the catalog. On Windows, also connect to the target database.
- 9 On UNIX, set `dbid` to be the DBID of the source client database. Connect to the target database without using a user ID and password.
- 10 Run an RMAN restore script. On UNIX, you can alternatively type the RMAN commands for the restore.

See [“Redirecting a restore to a different client”](#) on page 115.

See [“Example - Performing a redirected restore of Oracle”](#) on page 118.

Example - Performing a redirected restore of Oracle

For example, assume the following:

- Source client is `camel`
- Destination client is `giraffe`

- Primary server is `lion`
- `ORACLE_SID` is `test`
- The user is connected to the Oracle database using a local connection, not SQL*Net
- UNIX user is `ora` on both `camel` and `giraffe`

To perform a redirected restore (example)

- 1** Create the following file on server `lion`:
 Windows: `install_path\NetBackup\db\altnames\giraffe`
 UNIX: `/usr/opensv/netbackup/db/altnames/giraffe`
 Edit `giraffe` to contain the name `camel`:
- 2** Do one of the following:
 - Windows: Use the BAR GUI to set `lion` as the primary server.
 - UNIX: Log onto `giraffe` as `ora`. Set `SERVER=lion` in `$ORACLE_HOME/bp.conf`. This server must be the first server that is listed in the `bp.conf` file.
- 3** Modify the network `tnsnames.ora` file to enable the RMAN catalog connection.
- 4** Create `inittest.ora` file.
- 5** Windows: Using Oracle administration, create and start `ORACLESERVICETEST`.
- 6** Set the environment variable `ORACLE_SID` to `test`. On UNIX, also set `NB_ORA_CLIENT` to `camel`.

- 7 Make sure that the destination database directory exists and has appropriate access permissions.

The data files are restored to the directory path with the same name they had when they were backed up.

- 8 Start the database in a `nomount` state.

On UNIX, the following is the output:

```
SQL> startup nomount pfile=$ORACLE_HOME/dbs/inittest.ora
%rman catalog rman/rman@rcat
RMAN> set dbid=<dbid of source database on camel>
RMAN> connect target/
RMAN> run {
RMAN>     ALLOCATE CHANNEL CH00 TYPE 'SBT_TAPE';
RMAN>     SEND 'NB_ORA_SERV=lion, NB_ORA_CLIENT=camel';
RMAN>     restore controlfile;
RMAN> }

SQL> alter database mount;
%orapwd file=$ORACLE_HOME/dbs/orapwtest password=<oracle>
%rman catalog rman/rman@RCVCAT

RMAN>set dbid=<Saved dbID of Source Target>
RMAN>connect target/
RMAN>run {
RMAN>     ALLOCATE CHANNEL CH00 TYPE 'SBT_TAPE';
RMAN>     ALLOCATE CHANNEL CH01 TYPE 'SBT_TAPE';
RMAN>     SEND 'NB_ORA_SERV=lion, NB_ORA_CLIENT=camel';
RMAN>     restore database;
RMAN>     restore archivelog all;
RMAN> }
```

```
SQL>recover database until cancel using backup controlfile;
```

Now apply the archived logs. Type `cancel` when you decide to stop recovery.

See [“Redirecting a restore to a different client”](#) on page 115.

See [“About performing a redirected restore with RMAN”](#) on page 117.

Using NetBackup for Oracle in a Windows Server Failover Cluster (WSFC)

To use NetBackup for Oracle in a Windows Server Failover Cluster (WSFC) environment, the following must be installed in the cluster nodes:

- NetBackup client or server
- NetBackup for Oracle on Windows
- Oracle Database version 10g or greater
- Oracle Failsafe 3.11 for Oracle 10g or greater
 - Review the Oracle compatibility list for complete information.

NetBackup for Oracle users in a Windows Server Failover Cluster environment must take some additional steps to prepare for server-directed backups, user-directed backups, and user-directed restores.

See [“About backups of an Oracle clustered database on Windows”](#) on page 121.

See [“User-directed backup or restore from the Windows client”](#) on page 123.

About backups of an Oracle clustered database on Windows

The most convenient way to back up your clustered databases is to set up schedules for automatic backups. NetBackup for Oracle comes with sample scripts for clustered Oracle databases. The NetBackup for Oracle installation process installs the sample scripts in the following location:

`install_path\NetBackup\dbext\oracle\samples\rman\`

Modify the scripts to give values to the following variables:

- Oracle SID
- Oracle Home
- Cluster Name, Domain
- Failsafe Home
- Failsafe user ID
- Failsafe Password
- Failsafe Database Resource Name
- Virtual Oracle Database Name

You can also manually back up an Oracle policy. Refer to the following procedure:

See [“Testing configuration settings for NetBackup for Oracle”](#) on page 98.

For more information on how to back up or restore Windows Server Failover Cluster using NetBackup, see the [NetBackup Clustered Primary Administrator's Guide](#).

See [“Bringing the database instance offline on Windows”](#) on page 122.

See [“User-directed backup or restore from the Windows client”](#) on page 123.

Bringing the database instance offline on Windows

Before you can perform a user-directed backup or restore from the client, you must take the database instance offline. You can use the Failsafe graphical user interface or the Failsafe command line (FSCMD).

To take the database instance offline with Failsafe graphical user interface

- 1 Select the Oracle database resource in the Failsafe graphical user interface.
- 2 Choose to bring it offline.

To take the database instance offline with Failsafe command line (FSCMD), type the following command:

```
■ fscmd offlineresource salesdb /cluster=curly /offline=immediate
   /domain=domainname /user=user /pwd=pwd
```

To bring the resource offline, the preceding command sets `offline=immediate`.

Alternately, based on your need you can specify one of the following as the argument:

<code>abort</code>	Shuts down the database instantaneously by aborting the database instance.
<code>immediate</code>	Shuts down the database immediately by terminating SQL statements in progress, rolling back uncommitted transactions and disconnecting users.
<code>normal</code>	Shuts down the database and doesn't allow new connections after the command was issued. This command waits for the connected users to disconnect before the database is shut down.
<code>transactional</code>	Shuts down the database only after all of the current transactions have completed.

Because the `offlineresource` operation shuts down the Oracle database service, enter the following command to start the Oracle database service:

```
net start OracleService
```

See [“User-directed backup or restore from the Windows client”](#) on page 123.

See [“Bringing the database instance online on Windows”](#) on page 123.

Bringing the database instance online on Windows

After you perform a user-directed backup or restore from the client, you must bring the database instance online. You can use the Failsafe graphical user interface or the Failsafe command line (FSCMD).

To bring the database instance online with Failsafe graphical user interface

- 1** Select the resource in the Failsafe graphical user interface.
- 2** Choose to bring it online.

To bring the database instance online with Failsafe command line (FSCMD), type the following command:

```
■ fscmd online resource salesdb /cluster=curly
   /offline=immediate /domain=domainname /user=user /pwd=pwd
```

See [“Bringing the database instance offline on Windows”](#) on page 122.

See [“User-directed backup or restore from the Windows client”](#) on page 123.

User-directed backup or restore from the Windows client

This section explains the process to prepare a Windows Server Failover Cluster environment for a user-directed backup or restore operation.

Note: When performing user-directed backups, make sure that you are on the node that owns the shared drive where the Oracle database is installed.

Note: When user-directed client restores are performed with different configuration options of NetBackup failover media servers and a Linux or Windows primary server, see the [NetBackup Administrator's Guide, Volume I](#).

To perform a user-directed backup or restore from the client

- 1** Take the clustered Oracle database instance offline.

See [“Bringing the database instance offline on Windows”](#) on page 122.

- 2** Shut down and then startup the database in `mount` state.

The sequence is necessary to perform administrative tasks like backup and recovery. Use the `svrmgrl` or `sqlplus` utility from Oracle. At the command line, type the following:

```
Shutdown option [normal, abort, immediate]
startup mount
```

- 3** Perform the backup or recovery.
- 4** Bring the Oracle database online with failsafe after the desired backup or restore is complete. The database is then enabled to fail over between the configured cluster of nodes.

See [“Bringing the database instance online on Windows”](#) on page 123.

Oracle cloning

This chapter includes the following topics:

- [About cloning an Oracle database or a pluggable database](#)
- [Create a custom RBAC role for Oracle cloning](#)
- [Clone an Oracle database](#)
- [Clone a pluggable database](#)
- [Cloning a database from an Oracle Intelligent Policy \(OIP\) that is Data Guard aware](#)
- [Additional functionality in NetBackup APIs](#)

About cloning an Oracle database or a pluggable database

The NetBackup web UI supports the following types of clone operations:

Table 8-1 Supported clone operations

Type of Oracle configuration backup	Clone to standalone Oracle installation	Clone pluggable database into existing container database
Standalone container database (CDB)	X	X
Standalone non-container database	X	
RAC database	X	X

Table 8-1 Supported clone operations (*continued*)

Type of Oracle configuration backup	Clone to standalone Oracle installation	Clone pluggable database into existing container database
Data Guard primary database	X	X
Data Guard standby database	X	See limitations regarding Data Guard.

The clone features in the NetBackup web UI offer a streamlined and a more efficient method to clone an Oracle database.

- A clone operation is performed from a good backup image and does not require access to the original database.
- Beginning with NetBackup 10.3, NetBackup collects additional information (metadata) at the time of backup that helps to create the RMAN cloning script.
- The validation process checks for any issues before NetBackup starts the clone operation.

Requirements for cloning an Oracle database or a pluggable database

Oracle cloning has the following requirements:

- The Oracle version must be the same on the source server and the target server.
- The operating system of the source file system and the target file system must be the same.
- The target operating system file system of the clone must be the same as the original operating system file system from which the backup was taken. For example, a clone of ext4 to XFS is not supported.
- An additional requirement may exist when the backup source is ASM and you are cloning to file system using OMF. Or, when the backup source is a file system and you are cloning to ASM using OMF. Before the clone you may need to manually create the target directories for the temp files and the redo log files. Otherwise, an error like the following may occur:

```
ORA-00344: unable to re-create online log
'E:\APP\ORACLE\PRODUCT\19.3\DB_HOME1\RAC_RESTORE2\DATA\CDB19R\
ONLINELOG\GROUP_1.264.1126884003'
```

See [“Error: Unable to re-create online log”](#) on page 237.

- To clone from an image that was created with an Oracle Intelligent Policy (OIP) that is Data Guard-aware, the Oracle Data Guard databases must be defined within NetBackup. You can configure Data Guards with the NetBackup web UI, the NetBackup APIs, or the `nboracmd` command.

See [“Cloning a database from an Oracle Intelligent Policy \(OIP\) that is Data Guard aware”](#) on page 133.

Limitations for cloning an Oracle database or a pluggable database

In NetBackup, Oracle cloning has certain limitations or is not supported in some situations, as follows:

- The backup is from an instant recovery snapshot.
- A clone is supported only from the primary backup copy. A clone from other backup copies is not supported.
- A clone is not supported from when you use the backup option **Database backup shares** (for Oracle Copilot).
- A clone from a hot full backup or an incremental backup that does not also include archived redo logs may fail due to RMAN limitations. Ensure that the option **Include archived redo logs in full and incremental schedules** is selected on the **Oracle** tab. Or configure a schedule for **Archived redo log backup** to perform a clone.
- Clone from a source database that has an offline data file.
- Clone from backup set where incarnation changes during the set of backups.
- When the following objects are encrypted: NetBackup displays an error in these cases.
 - Column encryption: Single PDB clone
 - Tablespace encryption:
 - Single PDB clone
 - Full CDB, non-CDB, and single PDB clone for encryption of Oracle default tablespace.
 - Oracle-managed tablespaces: SYSTEM, SYSAUX, UNDO, TEMP
- When the Oracle backup source has a source location in the control file for a volume that doesn't exist on the target.

On the target, you must create the source location that is named in the control file.

Review the following Oracle documentation: [Full RMAN Backup Set Duplication = Error in the Alert Log: ORA-7445\[skgfifi\(\)+4307\]](#) (Doc ID 2443991.1)

- If the source backup was a proxy backup, only a file system restore is supported. If you want to move the cloned database to ASM storage, additional steps are required. More information is available on how to restore from a data file copy to ASM storage using RMAN.

See [“About restoring from a data file copy to ASM storage using RMAN”](#) on page 160.

- NetBackup creates a clone control file and database server parameter file (SPFILE) based on the control file and SPFILE of the source database. NetBackup modifies the control file and SPFILE with parameters provided by the user during the clone operation. See the Oracle database documentation on supported characters for SPFILE or control file parameters. Different database versions may have different requirements.

The following examples may cause a clone to fail:

- If the SPFILE of the clone that is created contains a path with localized characters.
- If the control file paths or the DB_CREATE_FILE_DEST contains localized characters.
- If a source database is not OMF, you cannot clone the database to ASM storage. When you request a clone operation, NetBackup detects this case and prompts you to convert the cloned database to OMF. If you select “Yes”, NetBackup converts the cloned database during the clone operation.
 See the following Oracle support article that describes some of the limitations of the ASM file system naming conventions when used with OMF and RMAN duplication. These same limitations apply to NetBackup clone operations. Also see the section on “Uniqueness of an ASM file”.
<https://support.oracle.com/epmos/faces/DocContentDisplay?id=1910175.1>
 NetBackup appends a randomly generated suffix onto the OMF-created ASM files to try to mitigate RMAN's restore behavior.
- NetBackup supports a source OS (for example, Chinese OS) with non-ASCII and an Oracle installation with one of the Oracle-defined character sets (for example, ZHS16GBK). Users can clone to a target that has the same OS and Oracle character set.

Create a custom RBAC role for Oracle cloning

A custom role can allow an Oracle administrator to clone an Oracle database or a pluggable database. Optionally you can provide a user access to a stored OS credential.

Use this role if you do not want users to have the Default Oracle Administrator role. Or, you do not want to give users access to all the available Oracle backup and recovery operations.

To create a custom role for Oracle cloning

- 1 If you want the user to have access to the OS credential, this credential must already exist in NetBackup Credential Management.
 See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.
 Contact your NetBackup administrator for assistance.
- 2 On the left, select **Security > RBAC** and click **Add**.
- 3 Select **Custom role** and click **Next**.
- 4 Provide a **Role name** and a description.
 For example, include a description that the role allows users to perform a database cloning.
- 5 Under **Permissions**, click **Assign**.
- 6 On the **Global** tab, expand the **NetBackup management** section. Locate **Jobs** and select the **View** permission.
- 7 On the **Assets** tab, locate **Oracle assets**. Select the following permissions.

Permission	Description
View	View Oracle databases and pluggable databases.
Restore	List recovery points for Oracle backups. Clone a database or pluggable database from a recovery point to the original host. Create an instant access database from a recovery point.
Restore to alternate location	(Optional) Clone a database or pluggable database from a recovery point to an alternate host.

- 8 Click **Assign**.
- 9 Under **Workloads**, click **Assign**.
- 10 On the **Oracle** tab, choose from the following.
 - To apply the same permissions to all Oracle assets, select **Apply selected permissions to all existing and future Oracle assets**.
 - To make only specific Oracle assets available to users with this RBAC role, click **Add**. Then add each of the databases or pluggable databases that you want to add to the role.

- 11 Click **Assign**.
- 12 Under **Users**, click **Assign**. Then add the users that you want to have this RBAC role.
- 13 When you are done configuring the role, click **Add role**.
- 14 Configure the credential that users need to perform cloning and assign the role that you created in this procedure or another role to that credential.
See [“Add a credential for an Oracle database or RMAN catalog”](#) on page 48.

Clone an Oracle database

The following procedure describes how to clone an Oracle database. The user must have an RBAC role that gives them permissions to perform cloning operations.

See [“Create a custom RBAC role for Oracle cloning”](#) on page 128.

To clone an Oracle database

- 1 On the left, click **Workloads > Oracle**.
- 2 Click the **Databases** tab.
- 3 Click on the name of the database you want to clone.
- 4 Select the **Recovery points** tab.
- 5 Locate the recovery point from which you want to create the clone.
- 6 Click **Actions > Clone database**.
- 7 In **Recovery target**, enter a host name. Or, click **Select host** to select a host.
See [the section called “Target directory notes”](#) on page 131.
- 8 (Oracle 21 and later) The path for the **Oracle base config** is required. Otherwise a clone to an Oracle 21 or later database fails.
- 9 Enter all other required information.
- 10 Click **Next**.
- 11 In **Credentials**, specify the applicable credential options.
 - Linux or UNIX: You can select from existing credentials or use OS credentials.
 - Windows: You can select an existing credential or use OS credentials. If the target Oracle database requires them, also enter the Oracle Home User credentials.

After all required information is entered, click **Next**.

- 12 In **Recovery options**, enter the number of channels and the SCN value and then click **Next**.
- 13 On the **Review** page, NetBackup runs a pre-recovery check for any errors. You should correct any errors that are highlighted.
During the pre-recovery check, NetBackup looks for:
 - Credentials
 - Database name
 - Oracle base config path
 - Image access
 - Existence of Oracle home
 - SCN Value
- 14 Click **Start recovery** to begin the cloning operation.
You can start the cloning operation without fixing any of the errors that are shown on the **Review** page.
- 15 (Optional) You can monitor the progress of the cloning operation in the **Activity monitor**.

Target directory notes

Note the following when you specify the target directory for cloning:

- If using an ASM path for any target, the syntax can be as follows:
+ASM^{MDG}/path
- If the source backup used OMF, selecting an ASM directory for control files or redo files requires a path after the disk group.
- If the source backup used OMF and from ASM, selecting an ASM directory for any new location can only contain the disk group. Oracle does not honor any paths after the disk group.
- When you clone to a specified target directory, the best effort is made to create the target directories along with the source subdirectories.

Clone a pluggable database

The following procedure describes how to clone a pluggable database. The user must have an RBAC role that gives them permissions to perform cloning operations.

See [“Create a custom RBAC role for Oracle cloning”](#) on page 128.

To clone a pluggable database

- 1 On the left, click **Workloads > Oracle**.
- 2 Click the **Pluggable databases** tab.
- 3 Click on the name of the pluggable database you want to clone.
- 4 Select the **Recovery points** tab.
- 5 Locate the recovery point from which you want to create the clone.
- 6 Click **Actions > Clone pluggable database**.
- 7 In **Recovery target**, enter a host. Or, click **Select destination**. That option automatically fills in both the **Host** and the **Target instance**.

A container must exist and must be open (read-write mode) for the pluggable database to be plugged into. None of the existing PDBs in the container can have the same name as the new PDB.
- 8 (Oracle 21 and later) The path for the **Oracle base config** is required. Otherwise a clone to an Oracle 21 or later database fails.
- 9 Enter all other required information.
- 10 In **Credentials**, specify the applicable credential options.
 - Linux or UNIX: You can select an existing credential or use OS credentials.
 - Windows: You can select an existing credential or use OS credentials. If the target Oracle database requires them, also enter the Oracle Home User credentials.

After all required information is entered, click **Next**.

- 11 In **Recovery options**, enter the number of channels and the SCN value and then click **Next**.
- 12 On the **Review** page, NetBackup runs a pre-recovery check for any errors. You should correct any errors that are highlighted.

During the pre-recovery check, NetBackup looks for:

- Credentials
- Database name
- Oracle base config path
- Image access
- Existence of Oracle home
- SCN Value

- 13** Click **Start recovery** to begin the cloning operation.

You can start the cloning operation without fixing any of the errors that are shown on the **Review** page.

- 14** (Optional) You can monitor the progress of the cloning operation in the **Activity monitor**.

Cloning a database from an Oracle Intelligent Policy (OIP) that is Data Guard aware

If you want to clone a database that was backed up with an Oracle Intelligent Policy that is Data Guard aware, the databases in the Oracle Data Guard must be defined within NetBackup.

See [“Manage a Data Guard”](#) on page 43.

Note the following about cloning from Data Guard configurations:

- A clone from a primary node backup or from a logical standby is similar to a clone from a typical database.
- For a clone from a physical or a snapshot standby, the success of the clone depends on the state of the redo log sets that NetBackup has access to.

The following example RMAN steps may resolve an issue after a clone:

```
SQL> alter database drop standby logfile group 1;
Database altered.
SQL> alter database drop standby logfile group 2;
Database altered.
SQL> alter database drop standby logfile group 3;
Database altered.
SQL> ALTER DATABASE ADD LOGFILE group 1 ('c:\clone\REDO01.LOG') size 5M;
Database altered.
SQL> ALTER DATABASE ADD LOGFILE group 2 ('c:\clone\REDO02.LOG') size 5M;
Database altered.
SQL> alter database activate standby database;
Database altered.
SQL> alter database open;
```

Additional functionality in NetBackup APIs

The NetBackup APIs offer additional functionality for Oracle cloning:

- If you want to rename every file to a specific path or location, the Recovery API provides this ability for the data file name, redo log, temp file, and control file.
- The Recovery Point Device API provides additional information about recovery point attributes and the chain of backup images.
- For the following cloning operations:
 - From a file system to a file system
 - From ASM to a file system
 The Recovery API allows a change from non-OMF to OMF or from OMF to non-OMF.

Refer to the NetBackup API Reference documentation on the Cohesity SORT website:

<https://sort.veritas.com/documents>

NetBackup Copilot for Oracle

This chapter includes the following topics:

- [About Oracle Copilot](#)
- [Configuring an OIP using a share on the NetBackup appliance \(Oracle Copilot\)](#)
- [Creating an instant recovery point from an Oracle Copilot image \(NetBackup Appliance share\)](#)
- [Deleting an instant recovery point for Oracle Copilot instant recovery](#)
- [Configuring an OIP using universal shares \(Oracle Copilot\)](#)
- [Managing an instant access mount from an Oracle Copilot universal share image](#)
- [Cleaning up the Oracle Copilot share after point in time restore of database](#)
- [Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount](#)
- [About restoring from a data file copy to ASM storage using RMAN](#)

About Oracle Copilot

This feature enhances the Oracle Intelligent Policy by giving you two options for protecting an Oracle database using a share on a NetBackup appliance or universal shares. The first option gives you better control of backups when Oracle database backups are placed in shares (including NetBackup appliance share and universal shares) by the DBA. The second option lets you choose an appliance share or universal shares as the destination for the first backup copy. Now you do not have to rely on the DBA to create backups in the share.

The **Database Backup Shares** option provides a single NetBackup Appliance share or universal shares for the DBA and is protected using the Oracle Intelligent Policy. The backup occurs on the storage server as an off-host backup and all data movement occurs on the storage server and does not affect the Oracle client. Since the OIP protects the shares, the backups are visible when the DBA uses RMAN or Oracle Enterprise Manager.

The **Whole Database - Datafile Copy Share** option enhances the OIP to allow the NetBackup Administrator to choose an appliance share or universal shares as the destination for the first backup copy. When the policy runs the first time, an RMAN script is generated that creates a full set of Oracle data file copies. The data file copies reside in the shares. The next time that the full schedule runs, the backup is accelerated if the **Use Accelerator** option is selected. The RMAN script that is generated performs an incremental backup and the changed blocks are merged into the data files. This incremental backup creates an updated full set of Oracle data file copies. After the new full copy is created in the database backup shares, an SLP is used to make additional copies of the full backup. The first copy is always a `remote_vxfs` snapshot. For NetBackup appliance share, the `remote_vxfs` snapshot creates the `vxfs_checkpoint` snapshot. For universal shares, there are several `vpfs` snapshots created.

On the **Attributes** tab, the **Use Accelerator** feature is automatically selected when you configure an OIP with **Whole Database - Datafile Copy Share** selected in the **Backup Selections** tab. The first time that the full schedule runs it creates a full set of data file copies. After the first full schedule, only the changes are backed up as a backup set and merged with the existing full backup. Basically, an incremental merge is performed. Oracle's Block Change Tracking feature should be enabled for faster incremental backups.

When using Oracle Copilot to protect your database, NetBackup does not protect extended attributes, extent attributes, or Access Control Lists associated with the database's data files.

Oracle Copilot works with both NetBackup appliance share and MSDP universal shares that are created on BYO, Flex, Flex WORM, and Flex Scale. Oracle Copilot with universal shares supports instant access. See the chapter *Oracle Copilot with Instant Access*.

Note: Oracle Copilot for Flex WORM is supported only on Red Hat Enterprise Linux media server.

Note: Create a share on the appliance using the procedures in the Managing shares section of the [NetBackup 52xx and 5330 Appliance Administrator's Guide](#). If you enter a share path that is not located on an appliance, nothing is backed up.

Note: The NetBackup appliance media server can only be configured in a pure IPv4 or dual stack IPv4/IPv6 mode.

See [“Configuring an OIP using a share on the NetBackup appliance \(Oracle Copilot\)”](#) on page 137.

See [“Backup selections tab”](#) on page 67.

See [“Create an Oracle Intelligent Policy \(OIP\)”](#) on page 60.

See [“About restoring from a data file copy to ASM storage using RMAN”](#) on page 160.

Configuring an OIP using a share on the NetBackup appliance (Oracle Copilot)

The following prerequisites exist for using a NetBackup appliance share:

- You must provision a share on the appliance for these options using the NetBackup Appliance Shell Menu.
- These options are available to you but only configurable when you have a NetBackup appliance configured as the media server.
- To use copilot with universal shares, see the chapter *Oracle Copilot with instant access*.

Note: This feature requires a NetBackup appliance running software version 2.7.1 or later.

The **Database backup shares** option protects the database backups that an Oracle DBA creates on a share on the NetBackup appliance.

The **Whole Database - Datafile Copy Share** option enhances the OIP to allow the NetBackup Administrator to choose an appliance share as the destination for the first backup copy. The backup copy is a full set of data file copies that are maintained by updating only the changed blocks if **Use Accelerator** is selected.

If using Oracle Copilot with RAC then the following must be configured:

- Each node of the RAC cluster must have access to the Oracle Copilot share.
- The Oracle Copilot share has to be mounted on each node of the RAC cluster using the same mount point name on each node.

Oracle Copilot with NetBackup Appliance share workflow

Table 9-1 Appliance share workflow

Step	Description
Create Oracle Copilot Policy.	Create an Oracle Copilot policy. See “Configuring an OIP using a share on the NetBackup appliance (Oracle Copilot)” on page 137.
Creating an instant recovery point.	Creating an instant recovery point. See “Creating an instant recovery point from an Oracle Copilot image (NetBackup Appliance share)” on page 139.
Common tasks	See “Cleaning up the Oracle Copilot share after point in time restore of database” on page 149. See “Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount” on page 156. See “About restoring from a data file copy to ASM storage using RMAN” on page 160.

Use the following procedure to set up a backup policy that protects shares on the NetBackup appliance.

To configure an OIP using the Database backup shares or Whole Database - Datafile Copy Share options

- (Database Backup Shares option)** The Oracle DBA asks NetBackup administrator for the appliance share information.
- The NetBackup administrator uses the NetBackup Appliance Shell Menu to create a share on the appliance and then sets permissions for the share.
 - For more information about how to set up the share, see [Creating a share from the NetBackup Appliance Shell Menu in the NetBackup 52xx and 5330 Appliance Administrator’s Guide](#)
- The NetBackup administrator sends information about the appliance share to system administrator.
- The system administrator mounts an appliance share on the Oracle database server using the OS tools.

- 5 (Database backup shares option)** The Oracle DBA uses RMAN to create a database backup on the appliance share.
- 6** Configure an OIP.

Attributes tab	The Use Accelerator option is automatically selected when the Whole Database - Datafile Copy Share option is selected. If the Use Accelerator option is unchecked the full set of data files are copied again (including changed blocks). When the Use Accelerator option is used, the Oracle Change Block tracking should be enabled for better performance. For information on the Attributes tab, see the NetBackup Administrator's Guide, Volume I.
Schedules tab	Click Add and select Full backup. For information on the Schedules tab, see the NetBackup Administrator's Guide, Volume I.
Instance and databases tab	Select the Oracle database. See "Instances and databases tab" on page 66.
Backup selections tab	Select the Database backup shares option. See "Backup selections tab" on page 67.
Oracle tab	Configure the Database backup share options to delete backup sets and backup copies. See "Oracle tab" on page 68.

Note: The NetBackup appliance media server can only be configured in a pure IPv4 or dual stack IPv4/IPv6 mode.

Creating an instant recovery point from an Oracle Copilot image (NetBackup Appliance share)

The `nborair` command can determine if an image is available for Oracle Copilot instant recovery.

Note: The functionality for creating an instant recovery point is not in the NetBackup web UI. This feature is a command-line option only.

Refer to the [NetBackup Commands Reference Guide](#) for more usage options using the `nborair` command.

To create an instant recovery point

- 1 Determine if there are any images available for instant recovery by running the `nborair -list_images [-client name] [-server primary] command`.

The NetBackup administrator or the DBA can run this command from the NetBackup client or primary server.

Example output:

```
# nborair -list_images -client orachost1.demo.com -server primarysrv123
Time: 08/30/2016 15:51:17 ID: orachost1.demo.com_1472590277 Full Backup policy1
Time: 08/31/2016 11:20:17 ID: orachost1.demo.com_1472660417 Full Backup policy1
Time: 09/02/2016 10:42:45 ID: orachost1.demo.com_1472830965 Full Backup policy1
```

- 2 List the files that are included in the backup image by running the `nborair -list_files -backupid backup_id command`.

The NetBackup administrator or the DBA can run this command from the NetBackup client or primary server. The DBA sees only the files they can access when this command is run.

Example output:

```
# nborair -list_files -backupid orachost1.demo.com_1472590277
-rw-r----- oracl12 dba 807411712 Sep 02 10:42 /backup/data_D-ORAC112_I-3955369132_TS-SYSAUX...
-rw-r----- oracl12 dba 744497152 Sep 02 10:42 /backup/data_D-ORAC112_I-3955369132_TS-SYSTEM...
-rw-r----- oracl12 dba 52436992 Sep 02 10:42 /backup/data_D-ORAC112_I-3955369132_TS-UNDOTBS...
-rw-r----- oracl12 dba 5251072 Sep 02 10:42 /backup/data_D-ORAC112_I-3955369132_TS-USERS_FN...
-rw-r----- oracl12 dba 163328 Sep 02 10:42 /backup/arch_D-ORAC112_I-3955369132_SCN-3744354_5i...
-rw-r----- oracl12 dba 2560 Sep 02 10:42 /backup/arch_D-ORAC112_I-3955369132_SCN-3744354_5i...
-rw-r----- oracl12 dba 98304 Sep 02 10:42 /backup/spfile_D-ORAC112_I-3955369132_T-20160902...
-rw-r----- oracl12 dba 1425408 Sep 02 10:42 /backup/cf_D-ORAC112_I-3955369132_T-20160902_5k...
```

- 3 Create the instant recovery point by running the `nborair -create_recovery_point -backupid backup_id -dest_client name` command. The `backup_id` is the same `backup_id` found in step 1.

The NetBackup administrator must run this command from the NetBackup primary server.

For this example, the destination client is **oracdest**.

Example output:

```
# nborair -create_recovery_point -backupid orachost1.demo.com_1472590277
  -dest_client oracdest
Appliance: appl5330
Export path: /shares/share1_orachost1.demo.com_1472590277_rp1
Export options: oracdest(rw,no_root_squash,insecure)
```

- 4 Mount the recovery point on the destination client using the OS tools and with the required mount options per Oracle documentation.

Example:

```
mount -t nfs
appl5330:/shares/share1_orachost1.demo.com_1472590277_rp1 /mnt
```

For Windows, Oracle's DNFS needs to be configured. The recovery point has to be exported with the insecure option.

- 5 (Conditional) On the destination host, verify the mount point is from the backup ID that was requested by running the `nborair -validate -backupid backup_id -mount_path mount_path` command.

The NetBackup administrator or the DBA can run this command on the destination host.

Example output:

```
# nborair -validate -backupid orachost1.demo.com_1472590277 -mount_path /mnt
Validation successful - Recovery point mounted on /mnt was
created from backup ID orachost1.demo.com_1472590277
```

Note: The NetBackup appliance media server can only be configured in a pure IPv4 or dual stack IPv4/IPv6 mode.

See [“Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount”](#) on page 156.

Deleting an instant recovery point for Oracle Copilot instant recovery

The `nborair` command can delete an instant recovery point that is available for Oracle Copilot instant recovery.

Note: The functionality for deleting an instant recovery point is not in the GUI. This feature is command line option only.

Refer to the [NetBackup Commands Reference Guide](#) for more usage options using the `nborair` command.

To delete an instant recovery point

- 1 (Conditional) Verify the recovery point is unmounted from the destination client using the OS tools.

UNIX: `umount /mnt`

- 2 List the recovery point on the NetBackup appliance by running the `nborair -list_recovery_points -appliance appliance_name` command.

The NetBackup administrator must run this command from the NetBackup primary server.

Example output:

```
# nborair -list_recovery_points -appliance app15330
Total 1 recovery points found.
```

```
Export path: /shares/share1_orachost1.demo.com_1472590277_rp1
Share name: share1
Export options: oracdest(rw,no_root_squash,insecure)
```

- 3 Delete the recovery point on the NetBackup appliance by running the `nborair -delete_recovery_point -appliance appliance_name -export_path export_path` command.

The NetBackup administrator must run this command from the NetBackup primary server.

Example output:

```
# nborair -delete_recovery_point -appliance app15330
-export_path /shares/share1_orachost1.demo.com_1472590277_rp1
```

Configuring an OIP using universal shares (Oracle Copilot)

Table 9-2 Oracle Copilot with universal shares workflow

Step	Description
Create a universal share.	A universal share can be created on one or more MSDP storage servers, Flex WORM instance, and Flex Scale Cluster. If the universal share is located on multiple storage servers, ensure that the NetBackup media server can access all the storage servers. For more information, refer to the NetBackup Web UI Administrator's Guide. Section name: Create a universal share
Mount universal share to the Oracle server.	Refer to the documentation for NFS options. If the Oracle server has the Direct NFS option enabled, add the servers or shares into the Oracle DNFS configuration file. For more information, refer to the NetBackup Deduplication Guide. Section name: Mounting a universal share created from the NetBackup web UI
Create an Oracle Copilot policy.	Enter the mount share path during the selection. See "Configuring an OIP using universal shares (Oracle Copilot)" on page 143.
Start the Oracle backup and wait for the backup job to succeed.	See "Backing up an Oracle database using Oracle Copilot policy with a universal share" on page 165.
Select an Oracle database from the Oracle workload page (web UI only).	See "Managing an instant access mount from an Oracle Copilot universal share image" on page 149. See "Configure an instant access mount" on page 165.
Create an instant access mount based on the recovery point (web UI only).	When you create an instant access mount, if you select Direct NFS, ensure that the DNFS check box is selected. If you want to parallelly restore from RAC nodes, enter the host name of the RAC nodes as the target server. See "Configure an instant access mount" on page 165.

Table 9-2
Oracle Copilot with universal shares workflow (*continued*)

Step	Description
Mount instant access to target Oracle server.	Refer to Oracle documentation for NFS options. If you want to parallelly restore from RAC nodes, ensure that the instant access mount is mounted on all nodes with the same name. For example use this command to mount the instant access mount: <code>mount -t.</code>
Catalog data on Instant Access mount	For RMAN > catalog, start with '<a_mount_share>' See “Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount” on page 156.
Use the cataloged data file copy to do RMAN operations, such as restore data file or restore database from an instant access share. For example, perform disaster recover on an alternative Oracle server.	Instant access share provides a full set of data of the source database. An Oracle Database Administrator can decide how to use the data. See “Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount” on page 156.
Common tasks	See “Cleaning up the Oracle Copilot share after point in time restore of database” on page 149. See “Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount” on page 156. See “About restoring from a data file copy to ASM storage using RMAN” on page 160.

Supported platforms for Oracle Copilot backup with universal shares

Table 9-3
Supported platforms

Universal shares	Platform	Version
Multiple	BYO	NetBackup 10.2 and later
Single	BYO, Flex Appliance, Flex Scale	NetBackup 10.0 and later
Multiple	Flex Appliance	NetBackup Flex Appliance 3.0 and later
Multiple	Flex Scale	NetBackup Flex Scale 3.1 and later
Single	NetBackup Appliance	NetBackup Appliance 5.0 and later

Table 9-3
Supported platforms (*continued*)

Universal shares	Platform	Version
Multiple	Flex WORM	NetBackup 10.5 and later
Single	Flex WORM	NetBackup 10.5 and later

The **Database backup shares** option protects the database backups that an Oracle DBA creates on a universal share.

The **Whole Database - Datafile Copy Share** option enhances the OIP to allow the NetBackup Administrator to choose a single or multiple universal shares as the destination for the first backup copy. The backup copy is a full set of data file copies that are maintained by updating only the changed blocks if **Use Accelerator** is selected.

If using Oracle Copilot with RAC then the following must be configured:

- Each node of the RAC cluster must have access to the Oracle Copilot shares.
- The Oracle Copilot shares have to be mounted on each node of the RAC cluster using the same mount point name on each node.

To configure an OIP using universal shares (Oracle Copilot)

- 1 Create universal shares with the NFS protocol.
Refer to the information on creating a universal share in the *NetBackup Web UI Administrator's Guide*.
- 2 Mount all the universal shares on the Oracle clients.
- 3 Create an Oracle Copilot policy and then in **Backup selections**, select the mount points of the universal shares.

Note: If multiple storage servers are configured for Oracle Copilot with Universal Share policy, you must create an MSDP storage unit on each MSDP storage server where the universal shares are located. This is irrespective of whether the storage units are used as backup from snapshot with SLP destination storage or not.

For example, a setup has two storage servers and a universal share is created on server1 but backup from snapshot with SLP destination storage is on server2. You must create an MSDP storage unit on both server1 and server2.

Recommendations when you configure an Oracle Copilot with universal shares

Note the following recommendations when you configure an Oracle Copilot with universal shares:

- Ensure that there is at least one datafile on one share or backup may fail with an empty share error.
- Ensure the following:
 - Multiple shares are mounted with the correct NFS options on the Oracle server.
See [“About policy attributes”](#) on page 79.
 - Multiple mount paths are added into the selection list of an Oracle Copilot policy.
- Ensure that the backup datafile stream number is larger than number of shares in the Oracle Copilot policy.
If the number of streams are less, NetBackup automatically changes the number of parallel streams to match the shares.
- If the source Oracle is of the type RAC, ensure that multiple shares are mounted on each instance node with the same mount path.
For example, if there are four shares on four nodes, each node must have all shares mounted.
- If you want to replace an old share, ensure that the new share is mounted on the same path. Run RMAN crosscheck before you start any backup job.
- If you want to add a new share to the policy, ensure that the new share is mounted on all Oracle nodes.
If you have new datafiles added to the database, the datafiles are backed up to the new share. If there is no new datafile added, the new share is empty and the backup job partially succeeds. It is recommended that you create a new policy for all the shares including the new share.
- You can restore datafile copies on the new mount shares from multiple nodes of Oracle RAC. It helps to improve the restore performance. To improve the performance, start by configuring multiple instances in RMAN script.

Troubleshooting issues related to Oracle Copilot with universal shares

The following table describes the issues that are related to Oracle Copilot with universal shares:

Table 9-4 Troubleshoot issues

Job state	Error message	Issue	Action
Partial success	ORA-19504: failed to create file "XXXXXXXXXXXX" ORA-27040: file create error, unable to create file	Shares that failed from nodes.	1 Verify that the shares on all nodes have correct permissions associated. 2 Remount the share, run crosscheck manually, and then run the job again.
Partial success	ORA-19625: error identifying file"XXXXXXXXXXXX" ORA-27037: unable to obtain file status	Datafile on the share may be corrupted or missing	1 If the file is missing, verify that the shares are mounted on all nodes. 2 Try to remount the share, run crosscheck manually, and then run the job again. 3 If the file exists, use Oracle tools to verify the file and ensure that the file is not corrupted. 4 If the file is corrupted, keep all the logs and call technical support. 5 Move the corrupted file from a shared drive to a local drive, run crosscheck, and then run the job again.

Table 9-4 Troubleshoot issues (*continued*)

Job state	Error message	Issue	Action
Partial success	Snapshot failed	Snapshot rotation limitation exceeded.	1 Check if a previous backup from snapshot job failed. 2 Increase the maximum snapshot number and run crosscheck manually. 3 Modify datafile copy tag in the Oracle tab of policy, and run a new full backup. If the error persists, call technical support.
Partial success	RMAN Catalog failed	Shares that failed from nodes.	1 Verify that the shares on all nodes have correct permissions associated. 2 Remount the failed share, find the backup set on share, and try to manually catalog it.
Success	Backup job takes more time than expected.	Performance issue	1 Verify the operating system resource and ensure that there is no performance issue from Oracle. 2 Run the job again with the force rescan option.
Fail	SLP job keeps failing.	The ushare is created on a different backup target storage.	1 Verify and change the SLP backup snapshot storage target. 2 Run the job again.

Managing an instant access mount from an Oracle Copilot universal share image

This feature is only available on the web UI.

See the following topics.

See [“Configure an instant access mount”](#) on page 165.

See [“Delete an instant access mount”](#) on page 169.

Cleaning up the Oracle Copilot share after point in time restore of database

After a point in time restore of an Oracle database, RMAN can leave files from the previous database incarnations on an Oracle Copilot share. NetBackup does not automatically clean up the files from the previous database incarnation. This procedure describes how to manually clean up the share using RMAN.

Note: The functionality for cleaning up an Oracle Copilot share is not in the GUI. This feature is command line option only.

To clean up the Oracle Copilot share

- 1 Open a command prompt on the NetBackup client.
- 2 Set the `NLS_DATE_FORMAT` to display hours, minutes, and seconds.

UNIX:

```
NLS_DATE_FORMAT=DD-MON-YY_HH24:MI:SS
export NLS_DATE_FORMAT
```

Windows:

```
set NLS_DATE_FORMAT=DD-MON-YY_HH24:MI:SS
```

- 3 Log into RMAN and if NetBackup uses the RMAN catalog, it is required to log in to the catalog.

- 4** Use the `RMAN> list incarnation of database;` command to find the Reset Time for the current incarnation.

Example:

List of Database Incarnations

DB Key	Inc Key	DB Name	DB ID	STATUS	Reset SCN	Reset Time
10046	10054	ORACLEC2	3019371157	PARENT	1	11-SEP-14_08:40:48
10046	10047	ORACLEC2	3019371157	PARENT	2233668	27-APR-17_10:23:22
10046	11551	ORACLEC2	3019371157	CURRENT	2323198	28-APR-17_10:41:37

- 5** Use the `list backup summary` completed before "`to_date()`" device type disk; command to find all the backup pieces from the previous incarnation by using the reset time ("`to_date()`" must match `NLS_DATE_FORMAT`).

Example:

```
RMAN> list backup summary completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" device type disk;
```

List of Backups

```
=====
```

Key	TY	LV	S	Device	Type	Completion Time	#Pieces	#Copies	Compressed	Tag
10192	B	F	A	DISK		27-APR-17_10:42:59	1	1	NO	TAG20170427T104257
10193	B	F	A	DISK		27-APR-17_13:16:37	1	1	NO	TAG20170427T131636
10194	B	F	A	DISK		27-APR-17_13:16:55	1	1	NO	TAG20170427T131654
10195	B	F	A	DISK		27-APR-17_13:28:52	1	1	NO	TAG20170427T132851
10196	B	F	A	DISK		27-APR-17_13:29:08	1	1	NO	TAG20170427T132906
10197	B	F	A	DISK		27-APR-17_14:00:31	1	1	NO	TAG20170427T140031
10198	B	F	A	DISK		27-APR-17_14:00:43	1	1	NO	TAG20170427T140043
10199	B	F	A	DISK		27-APR-17_14:07:31	1	1	NO	TAG20170427T140730
10200	B	F	A	DISK		27-APR-17_14:07:48	1	1	NO	TAG20170427T140747
10759	B	A	A	DISK		28-APR-17_10:28:46	1	1	NO	DCS_CDB
10786	B	F	A	DISK		28-APR-17_10:28:56	1	1	NO	DCS_CDB
10814	B	F	A	DISK		28-APR-17_10:29:08	1	1	NO	DCS_CDB

- 6** Use the `list backup summary` completed before `"to_date()"` device type disk tag `' '`; command to find the backup pieces on the share by using the tag (by default, the tag is the NetBackup policy name).

Example:

```
RMAN> list backup summary completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" device type disk tag 'DCS_CDB';
```

List of Backups

=====

Key	TY	LV	S	Device	Type	Completion Time	#Pieces	#Copies	Compressed	Tag
-----	--	--	-	-----	-----	-----	-----	-----	-----	---
10759	B	A	A	DISK		28-APR-17_10:28:46	1	1	NO	DCS_CDB
10786	B	F	A	DISK		28-APR-17_10:28:56	1	1	NO	DCS_CDB
10814	B	F	A	DISK		28-APR-17_10:29:08	1	1	NO	DCS_CDB

- 7 Use the** list backup completed before "to_date()" device type disk tag ''; command to remove the summary option to see what files need deleting.

Example:

```
RMAN> list backup completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" device type disk tag 'DCS_CDB';
```

List of Backup Sets

=====

BS Key	Size	Device Type	Elapsed Time	Completion Time
10759	40.00K	DISK	00:02:55	28-APR-17_10:28:46
BP Key: 10762 Status: AVAILABLE Compressed: NO Tag: DCS_CDB				
Piece Name: C:\HA_NBA_SHARE\ARCH_D-ORACLEC2_I-3019371157_SCN-2323355_26S2QF5F_DCS_CDB				

List of Archived Logs in backup set 10759

Thrd	Seq	Low SCN	Low Time	Next SCN	Next Time
1	19	2322734	28-APR-17_10:16:54	2323527	28-APR-17_10:25:48
1	20	2323527	28-APR-17_10:25:48	2323546	28-APR-17_10:25:49

BS Key	Type	LV	Size	Device Type	Elapsed Time	Completion Time
10786	Full		80.00K	DISK	00:03:02	28-APR-17_10:28:56
BP Key: 10789 Status: AVAILABLE Compressed: NO Tag: DCS_CDB						
Piece Name: C:\HA_NBA_SHARE\SPFILE_D-ORACLEC2_I-3019371157_T-20170428_27S2QF5I_DCS_CDB						
SPFILE Included: Modification time: 27-APR-17_14:57:53						
SPFILE db_unique_name: ORACLEC2						

BS Key	Type	LV	Size	Device Type	Elapsed Time	Completion Time
10814	Full		17.17M	DISK	00:03:11	28-APR-17_10:29:08
BP Key: 10816 Status: AVAILABLE Compressed: NO Tag: DCS_CDB						
Piece Name: C:\HA_NBA_SHARE\CF_D-ORACLEC2_I-3019371157_T-20170428_28S2QF5L_DCS_CDB						
Control File Included: Ckp SCN: 2323603 Ckp time: 28-APR-17_10:25:57						

- 8 Use the delete backup completed before "to_date()" device type disk tag '' ; command to delete the unwanted backup pieces.**

Example:

```
RMAN> delete backup completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" device type disk tag 'DCS_CDB';
```

```
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=242 device type=DISK
```

List of Backup Pieces

BP Key	BS Key	Pc#	Cp#	Status	Device	Type	Piece Name
10762	10759	1	1	AVAILABLE	DISK		C:\HA_NBA_SHARE\ARCH_D-ORACLEC2_I-3019371157_SCN-2323355_26S2QF5F_DCS_CDB
10789	10786	1	1	AVAILABLE	DISK		C:\HA_NBA_SHARE\SPFILE_D-ORACLEC2_I-3019371157_T-20170428_27S2QF5I_DCS_CDB
10816	10814	1	1	AVAILABLE	DISK		C:\HA_NBA_SHARE\CF_D-ORACLEC2_I-3019371157_T-20170428_28S2QF5L_DCS_CDB

Do you really want to delete the above objects (enter YES or NO)? YES

deleted backup piece

backup piece handle=C:\HA_NBA_SHARE\ARCH_D-ORACLEC2_I-3019371157_SCN-2323355_26S2QF5F_DCS_CDB
 RECID=50 STAMP=942488751

deleted backup piece

backup piece handle=C:\HA_NBA_SHARE\SPFILE_D-ORACLEC2_I-3019371157_T-20170428_27S2QF5I_DCS_CDB
 RECID=51 STAMP=942488754

deleted backup piece

backup piece handle=C:\HA_NBA_SHARE\CF_D-ORACLEC2_I-3019371157_T-20170428_28S2QF5L_DCS_CDB
 RECID=52 STAMP=942488758

Deleted 3 objects

- 9** Use the `list copy completed before "to_date()" tag '';` command to find the data file copies on the share using the same reset time and tag.

Example:

```
RMAN> list copy completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" tag 'DCS_CDB';
```

specification does not match any control file copy in the repository

List of Datafile Copies

=====

Key	File S	Completion Time	Ckp SCN	Ckp Time
10649	1	A 28-APR-17_10:25:39	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-1_1GS2QE1J_S-48_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
10251	2	A 28-APR-17_10:15:32	2243146	27-APR-17_10:31:51
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-2_1LS2QEGQ_S-53_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
Container ID: 2, PDB Name: PDB\$SEED				
10648	3	A 28-APR-17_10:25:39	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-3_1IS2QE8G_S-50_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
10249	4	A 28-APR-17_10:13:19	2243146	27-APR-17_10:31:51
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-4_1JS2QEBG_S-51_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
Container ID: 2, PDB Name: PDB\$SEED				
10647	5	A 28-APR-17_10:25:38	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-UNDOTBS1_FNO-5_1HS2QE57_S-49_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
10646	6	A 28-APR-17_10:25:37	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-USERS_FNO-6_1NS2QEJV_S-55_I-3019371157_DCS_CDB				
Tag: DCS_CDB				

10 Use the delete copy completed before "to_date()" tag '' command to delete the data file copies on the selected share.

Example:

```
RMAN> delete copy completed before "to_date('28-APR-17_10:41:37',
'DD-MON-YY_HH24:MI:SS')" tag 'DCS_CDB';
```

```
released channel: ORA_DISK_1
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=242 device type=DISK
specification does not match any control file copy in the repository
List of Datafile Copies
=====
```

Key	File S	Completion Time	Ckp SCN	Ckp Time
10649	1	A 28-APR-17_10:25:39	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-1_1GS2QE1J_S-48_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
10251	2	A 28-APR-17_10:15:32	2243146	27-APR-17_10:31:51
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-2_1LS2QEGQ_S-53_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
Container ID: 2, PDB Name: PDB\$SEED				
10648	3	A 28-APR-17_10:25:39	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-3_1IS2QE8G_S-50_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
10249	4	A 28-APR-17_10:13:19	2243146	27-APR-17_10:31:51
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-4_1JS2QEBG_S-51_I-3019371157_DCS_CDB				
Tag: DCS_CDB				
Container ID: 2, PDB Name: PDB\$SEED				
10647	5	A 28-APR-17_10:25:38	2323417	28-APR-17_10:25:15
Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-UNDOTBS1_FNO-5_1HS2QE57_S-49_I-3019371157_DCS_CDB				
Tag: DCS_CDB				

Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount

```

10646      6      A 28-APR-17_10:25:37 2323417      28-APR-17_10:25:15
          Name: C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-USERS_FNO-6_1NS2QEJV_S-55_
I-3019371157_DCS_CDB
          Tag: DCS_CDB

Do you really want to delete the above objects (enter YES or NO)? YES
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-1_1GS2QE1J_
S-48_I-3019371157_DCS_CDB RECID=36 STAMP=942488739
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSTEM_FNO-2_1LS2QEGQ_
S-53_I-3019371157_DCS_CDB RECID=29 STAMP=942488132
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-3_1IS2QE8G_
S-50_I-3019371157_DCS_CDB RECID=35 STAMP=942488739
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-SYSAUX_FNO-4_1JS2QEBG_
S-51_I-3019371157_DCS_CDB RECID=27 STAMP=942487999
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-UNDOTBS1_FNO-5_1HS2QE57_
S-49_I-3019371157_DCS_CDB RECID=34 STAMP=942488738
deleted datafile copy
datafile copy file name=C:\HA_NBA_SHARE\DATA_D-ORACLEC2_I-3019371157_TS-USERS_FNO-6_1NS2QEJV_
S-55_I-3019371157_DCS_CDB RECID=33 STAMP=942488737
Deleted 6 objects

```

Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount

The following procedure shows how to use RMAN to restore from a recovery point. This procedure is only viable after the command `nborair -create_recovery_point` is run and the recovery point is mounted on a target client.

Note: The functionality for single-step restore to ASM storage is not in the GUI. This feature is run with RMAN only.

The procedure example assumes that a recovery point is already mounted and uses the mount point of `/db_mp` as the example. All RMAN commands must run from the target host.

Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount**To perform a single-step restore to ASM storage from a recovery point****1 Catalog the backups from the recovery point.**

```
RMAN> catalog start with '/db_mp/';
```

```
searching for all files that match the pattern /db_mp/
```

```
List of Files Unknown to the Database
```

```
=====
```

```
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-SYSAUX_FNO-2_8hrqu3qd_s-1297_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-SYSTEM_FNO-1_8irgu3qk_s-1298_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-UNDOTBS1_FNO-3_8jrgu3qr_s-1299_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-USERS_FNO-4_8krgu3qt_s-1300_I-3955369132
File Name: /db_mp/arch_D-ORAC112_I-3955369132_SCN-5248163_a8rh0s3b
File Name: /db_mp/spfile_D-ORAC112_I-3955369132_T-20160929_a9rh0s3c
File Name: /db_mp/cf_D-ORAC112_I-3955369132_T-20160929_aarh0s3d
```

```
Do you really want to catalog the above files (enter YES or NO)? YES
```

```
cataloging files...
```

```
cataloging done
```

```
List of Cataloged Files
```

```
=====
```

```
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-SYSAUX_FNO-2_8hrqu3qd_s-1297_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-SYSTEM_FNO-1_8irgu3qk_s-1298_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-UNDOTBS1_FNO-3_8jrgu3qr_s-1299_I-3955369132
File Name: /db_mp/data_D-ORAC112_I-3955369132_TS-USERS_FNO-4_8krgu3qt_s-1300_I-3955369132
File Name: /db_mp/arch_D-ORAC112_I-3955369132_SCN-5248163_a8rh0s3b
File Name: /db_mp/spfile_D-ORAC112_I-3955369132_T-20160929_a9rh0s3c
File Name: /db_mp/cf_D-ORAC112_I-3955369132_T-20160929_aarh0s3d
```

Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount**2 Restore the data files from the point in time of the recovery point.**

The following RMAN restore is from disk (DISK). Also, this example uses the `NLS_DATE_FORMAT="DD-MM-YYYY-HH24:MI:SS"` command that was set in the environment before RMAN was run. Use the date format for your environment.

```
RMAN> restore until time '2016-09-29-10:00:00' database;
```

```
Starting restore at 2016-10-12:15:51:22
```

```
allocated channel: ORA_DISK_1
```

```
channel ORA_DISK_1: SID=193 device type=DISK
```

```
channel ORA_DISK_1: restoring datafile 00001
```

```
input datafile copy RECID=461 STAMP=925055096
```

```
file name=/demo_2/data_D-ORAC112_I-3955369132_TS-SYSTEM_FNO-1_8irgu3qk_s-1298_I-3955369132
```

```
destination for restore of datafile 00001: /db/orac112/app/oradata/orac112/system01.dbf
```

```
channel ORA_DISK_1: copied datafile copy of datafile 00001
```

```
output file name=/db/orac112/app/oradata/orac112/system01.dbf RECID=0 STAMP=0
```

```
Finished restore at 2016-10-12:15:51:34
```

Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount**3 Recover the database.**

If the restore of archive logs is not available on disk, then the logs are restored from NetBackup (`sbt_tape`).

```

RMAN> run
{
allocate channel ch00 type sbt_Tape;
recover database;
release channel ch00;
}

released channel: ORA_DISK_1
allocated channel: ch00
channel ch00: SID=193 device type=SBT_TAPE
channel ch00: Veritas NetBackup for Oracle - Release 8.0 (2016091418)

Starting recover at 2016-10-12:15:54:13

starting media recovery

archived log for thread 1 with sequence 508 is already on disk as file
/db/orac112/app/fast_recovery_area/ORAC112/archivelog/2016_09_29/o1_mf_1_508_cytbkv22_.arc
archived log for thread 1 with sequence 509 is already on disk as file
/db/orac112/app/fast_recovery_area/ORAC112/archivelog/2016_09_29/o1_mf_1_509_cytbkv36_.arc
....
archived log file name=
/db/orac112/app/fast_recovery_area/ORAC112/archivelog/2016_09_29/o1_mf_1_508_cytbkv22_.arc
thread=1 sequence=508
archived log file name=
/db/orac112/app/fast_recovery_area/ORAC112/archivelog/2016_09_29/o1_mf_1_509_cytbkv36_.arc
thread=1 sequence=509
....
media recovery complete, elapsed time: 00:00:55
Finished recover at 2016-10-12:15:55:09

released channel: ch00

RMAN>

```

See [“Creating an instant recovery point from an Oracle Copilot image \(NetBackup Appliance share\)”](#) on page 139.

See [“About Oracle Copilot”](#) on page 135.

See “[Configuring an OIP using a share on the NetBackup appliance \(Oracle Copilot\)](#)” on page 137.

About restoring from a data file copy to ASM storage using RMAN

When you use a proxy method for data file copies, NetBackup cannot place the file directly back in ASM storage. You need to do a two-step restore for the data file copies.

If the backups are stream-based then restore directly from NetBackup.

When you restore back to the appliance share, make sure that the share on the appliance is configured with the `no_root_squash` NFS export option enabled.

For more information, refer to the Managing shares chapter in the [NetBackup Appliance Administrator's Guide](#).

The first step is to stage the files to a file system. The second step is to use RMAN to restore the files into ASM storage.

The following is an example RMAN script to stage the files to a file system:

```
RUN {
  ALLOCATE CHANNEL ch00
    TYPE 'SBT_TAPE';
  SEND 'NB_ORA_CLIENT=clientname,NB_ORA_SERV=servername';
  SET NEWNAME FOR TABLESPACE USERS TO '/dump/%U';
  RESTORE TABLESPACE USERS;
  RELEASE CHANNEL ch00;
}
```

Once the file is on a file system, then you can restore to ASM storage by running the following:

```
RUN {
  ALLOCATE CHANNEL dc00 DEVICE TYPE DISK;
  RESTORE TABLESPACE USERS;
  RECOVER DATABASE;
  RELEASE CHANNEL dc00;
}
```

Note: Any proxy backup of an Oracle database on ASM cannot be cloned to another ASM environment. Only a File System Restore or the 2-step restore process is supported.

See [“Single-step restore to ASM storage from an Oracle Copilot recovery point or instant access mount”](#) on page 156.

See [“About NetBackup for Oracle restores”](#) on page 111.

Oracle Copilot with instant access

This chapter includes the following topics:

- [Prerequisites when you configure an instant access Oracle database](#)
- [Things to consider before you configure an instant access mount point](#)
- [Backing up an Oracle database using Oracle Copilot policy with a universal share](#)
- [Configure an instant access mount](#)
- [View the livemount details of an instant access mount](#)
- [Configuring Auto Image Replication for Oracle instant access backups](#)
- [Delete an instant access mount](#)
- [NetBackup for Oracle terms](#)
- [Frequently asked questions](#)

Prerequisites when you configure an instant access Oracle database

The following prerequisites apply when you configure Oracle instant access databases:

- Instant access is only supported for Oracle backup images when the following conditions are met:

- The backup image is created using an Oracle Copilot Universal Share policy. The policy has two options, that you can select:
 - **Whole Database-Datafile Copy Share** is supported by the web UI.
 - **Database backup shares** is only supported from the REST API.
For more information, see the following topic:
See [“Backup selections tab”](#) on page 67.
- The backup is a full database backup.
- Supported platforms for Oracle Copilot Backup with Instant Access:

Primary	Platform	Version
Backup From Single NFS Server	BYO, Flex Appliance, Flex Scale, MSDP AKS (Azure Kubernetes Services), EKS (Amazon Elastic Kubernetes Service)	NetBackup 10.0 and later
Backup From Single NFS Server	Flex WORM	NetBackup 10.5 and later
Backup From Single NFS Server	NetBackup Appliance	NetBackup Appliance 4.0 and later
Backup From Single NFS Server	Cloud LSU (logical storage unit)	NetBackup 10.0.1 and later
Backup From Multiple NFS Server	Flex Scale	NetBackup Flex Scale 3.1 and later
Backup From Multiple NFS Server	BYO	NetBackup 10.2 and later
Backup From Multiple NFS Server	Flex Appliance	NetBackup Flex Appliance 3.0 and later

- For Build Your Own (BYO) server, the operating system version must be same as the latest appliance operating system version that is RHEL 7.6 and later.
- For BYO server, NGINX is installed on the storage server.
 - The NGINX version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (EPEL).
 - If SE Linux is configured, ensure that the `policycoreutils` and `policycoreutils-python` (for RHEL 7) or `policycoreutils-python-utils`

(for RHEL 8) packages are installed from the same RHEL yum source (RHEL server). Then run the following commands:

- semanage port -a -t http_port_t -p tcp 10087
- setsebool -P httpd_can_network_connect 1
- For BYO server, the /mnt folder on the storage server cannot be mounted by any mount points directly. User mount points must be mounted to its subfolders.
- For BYO server, enable the logrotate permission in SELinux using the following command:
semanage permissive -a logrotate_t

Hardware configuration requirement of instant access

Table 10-1 Hardware configuration

CPU	Memory	Disk
■ Minimum 2.2-GHz clock rate. ■ 64-bit processor. ■ Minimum 4 cores; 8 cores recommended. For 64 TB storage, the Intel x86-64 architecture requires eight cores.	■ 16 GB (For 8 TB to 32 TB of storage) ■ 1 GB RAM for 1 TB storage. ■ 32 GB of RAM for more than 32 TB storage. ■ An additional 500 MB of RAM for each live mount.	Disk size depends on the size of your backup. Refer to the hardware requirements for NetBackup and Media Server Deduplication Pool (MSDP).

Things to consider before you configure an instant access mount point

The instant access feature enhances the Oracle Intelligent Policy and gives you options to protect an Oracle database using multiple universal shares. This feature gives you better control of backups when an Oracle database backup is placed in a database share by the DBA. This feature also lets you choose a database share as the destination for the primary backup copy. The backup copy is a full set of database data file copies created, incrementally updated, and protected by NetBackup.

For more information about universal share, refer to [NetBackup Administrator's Guide, Volume I](#).

Note the following about the instant access Oracle feature:

- The oracle copilot backup with universal share can only be used for instant access and cannot be used for Oracle Copilot instant recovery.
- For instant access to work following an upgrade of NetBackup, first restart the NetBackup Web Service on the primary server. Run the following commands:
 - `/usr/opensv/netbackup/bin/nbwmc stop`
 - `/usr/opensv/netbackup/bin/nbwmc start`

Backing up an Oracle database using Oracle Copilot policy with a universal share

Before you configure Oracle Copilot instant access, you must back up the Oracle database using the Oracle Copilot policy with a universal share.

To back up Oracle copilot with a universal share

- 1 Create universal shares with the NFS protocol.

Refer to the information on creating a universal share in the *NetBackup Web UI Administrator's Guide*.

- 2 Mount all the universal shares on the Oracle clients.

- 3 Create an Oracle Copilot policy and then in **Backup selections**, select the mount points of the universal shares.

Refer to the *NetBackup Copilot for Oracle Configuration Guide* for information about how to configure an Oracle copilot policy.

See [“Recommendations when you configure an Oracle Copilot with universal shares”](#) on page 146.

See [“Troubleshooting issues related to Oracle Copilot with universal shares”](#) on page 146.

Configure an instant access mount

You can configure an instant access database from a full backup. The full backup must come from an Oracle copilot backup with data file copies in the NetBackup universal share.

You can configure an instant access Oracle database from the web UI or use the REST API.

You can create an instant access mount for a database and for a pluggable database.

Configure an instant access mount

- 1 On the left, select **Workloads > Oracle**.
- 2 Do one of the following:
 - On the **Databases** tab, click the database and then click the **Recovery points** tab if you want to configure an instant access mount for a database.
 - On the **Pluggable databases** tab, click the database and then click the **Recovery points** tab if you want to configure an instant access mount for a pluggable database.
 - On the **Databases** tab, click the database and then click the **Pluggable databases > Recovery points** tab if you want to configure an instant access mount for a pluggable database.
- 3 On the **Recovery points** tab, click the date on which the backup occurred.
The available images appear in rows with the backup timestamp for each image.
- 4 Right-click on the backup image and click **Actions > Configure instant access mount**.

Note: This option is only displayed if the recovery point supports instant access.

- 5 Enter the host names where you want to configure the instant access database.
- 6 (Optional) If you have configured NFS in the Oracle client, select the **Direct NFS** check box for the universal shares support the feature.
- 7 Click **Configure**.
- 8 After the instant access job starts, click on the **Restore activity** tab to view the progress.

See [“View the livemount details of an instant access mount”](#) on page 166.

View the livemount details of an instant access mount

You can view the livemount details of an instant access mount.

To view the livemount details of an instant access mount

- 1
- On the left, select **Workloads > Oracle**.
- 2
- Click the **Instant access databases** tab.
The tab lists the instant access databases.
- 3
- On the **Instant access databases** tab, click on the database name to see its details.

Clone of	Instant access database and pluggable database for which you configured the instant access mount.
Database name	Name of the Oracle database.
Storage server	Name of the storage server.
Mount ID	Unique ID for an instant access livemount.
Export path	Exported instant access livemount path from the storage server.
Live mount path	Exported path with the storage server path. Click on the copy-to-clipboard icon and use the path to mount it in your Oracle environment and see the backed up datafiles. When you create an instant access mount, if you back up an Oracle database to universal shares from a single storage server, a single live mount path is configured. If you back up Oracle database to multiple universal shares from different storage servers, depending on the number of storage servers the live mount paths are configured.
Recovery point time	Date when the recovery point was created.
Created on	Date when the instant access livemount was created.
Retention	Time period for which you want the instant access mount to be retained.

Note: If you have configured an instant access mount to multiple Oracle hosts, mount all livemount paths to all the Oracle hosts.

Configuring Auto Image Replication for Oracle instant access backups

You must add replication to the stream and the snap storage lifecycle policies (SLPs) to ensure that the SLP metadata is replicated for restore from an instant access backup image.

Oracle copilot with universal share's SLP (storage life policy) must include the backup from snapshot operation. The backup can only be used for Instant Access. Instant Recovery (`nborair`) is not supported.

To configure Auto Image Replication for Oracle instant access backups

- 1 Configure Auto Image Replication between two NetBackup master servers.
 For more information about configuring Auto Image Replication, refer to the [NetBackup™ Administrator's Guide, Volume I](#).
- 2 For Oracle instant access create the following storage lifecycle policies (SLPs):
 - stream-slp for backup
 - snap-slp for backup from snapshot
 For more information about configuring storage lifecycle policies, refer to the [NetBackup™ Administrator's Guide, Volume I](#).
- 3 Sign in to the NetBackup web UI.
- 4 On the left, click **Storage > Storage lifecycle policies** and select the **stream-slp** policy.
- 5 In the **Backup** operation, click **Actions** and select **Add child**.
- 6 On the **Properties** pane, select the **Replication** operation and then select the **Target import SLP**.
- 7 Click **Create** to add the child.
- 8 On the NetBackup web UI, click **Storage > Storage lifecycle policies** and select the **snap-slp** policy.
- 9 In the **Snapshot** and **Backup from Snapshot** operations, click **Actions** and select **Add child**.
- 10 On the **Properties** pane, select the **Replication** operation and then select the **Target import SLP**.
- 11 Click **Create** to add the child.

Replication then runs successfully for **Backup** and **Backup from Snapshot** operations in both SLPs.

Delete an instant access mount

You can delete the instant access mount when it is no longer used. Before you delete an instant access mount, if you have mounted one livemount to multiple oracle clients, ensure that the livemount is umounted from all clients.

To delete an instant access mount

- 1 On the left, select **Workloads > Oracle**.
- 2 Click the **Instant access databases** tab.
The tab lists the instant access databases.
- 3 Select **Actions > Remove Instant access mount > Delete**.
Ensure that you run the RMAN crosscheck command after the mount is deleted.
- 4 After an instant access mount share is deleted, perform an RMAN crosscheck of the share before the next backup to prevent failures.
- 5 Specify the "disk" type instead of a "SBT_TAPE" type. Use the default `<NetBackup_polycname>`. Or, if the datafile copy tag is changed in the **Oracle** tab, use that tag name in place of `<NetBackup_polycname>`.

Example command:

```
Run {  
  
Allocate channel ch00 type 'disk';  
  
crosscheck backup tag <Netbackup_polycname>;  
  
crosscheck copy <Netbackup_polycname>;  
  
release channel ch00;  
  
}
```

NetBackup for Oracle terms

The table describes the important terms that might be new to an Oracle database administrator or a NetBackup administrator.

Term	Definition
Full backup	A full backup backs up all the blocks into the backup set, skipping only data file blocks that have never been used. Note that a full backup is not the same as a whole database backup; "full" is an indicator that the backup is not incremental. A full backup has no effect on subsequent incremental backups, which is why it is not considered part of the incremental strategy. In other words, a full backup does not affect which blocks are included in subsequent incremental backups.
Incremental backup	An incremental backup is a backup of only those blocks that have changed since a previous backup. Oracle lets you create and restore incremental backups of data files, tablespaces, and a database. You can include a control file in an incremental backup set, but the control file is always included in its entirety. No blocks are skipped.
Multilevel incremental backup	RMAN lets you create multilevel backups. RMAN can create multilevel incremental backup. A value of 0 or 1 denotes each incremental level. A level 0 incremental backup, which is the base for subsequent incremental backups, copies all blocks containing data. You can create a level 0 database backup as backup sets or image copies. The only difference between a level 0 incremental backup and a full backup is that a full backup is never included in an incremental strategy. Thus, an incremental level 0 backup is a full backup that happens to be the parent of incremental backups whose level is greater than 0. The benefit to performing multilevel incremental backups is that you do not back up all of the blocks all of the time. Incremental backups at a level greater than zero (0) only copy the blocks that were modified. Hence, the backup size can be significantly smaller and the backup might require much less time. The size of the backup file depends solely upon the number of blocks that are modified and the incremental backup level.
Differential incremental backup	In a differential level 1 backup, RMAN backs up all blocks that have changed since the most recent incremental backup at level 1 (cumulative or differential) or level 0. For example, in a differential level 1 backup, RMAN determines which level 1 backup is the most recent backup. RMAN backs up all blocks that have been modified after that backup. If no level 1 is available, then RMAN copies all blocks that have changed since the base level 0 backup.

Term	Definition
Cumulative incremental backup	In a cumulative level 1 incremental backup, RMAN backs up all blocks that have changed since the most recent backup at level 0. Cumulative incremental backups reduce the work that is needed for a restore. The cumulative incremental backup ensures that you only need one incremental backup from any particular level at restore time. Cumulative backups require more space and time than differential incremental backups. Cumulative backups duplicate the work that previous backups did at the same level.

Frequently asked questions

Here are some frequently asked questions for Oracle instant access.

Frequently asked questions	Answer
How can I enable the Oracle instant access feature on BYO after storage is configured or upgraded without the NGINX service installed?	Perform the steps in the following order: 1 Install the required NGINX service version. 2 Run the command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code> 3 Ensure that the new BYO NGINX configuration entry <code>/etc/nginx/conf.d/byo.conf</code> is part of the HTTP section of the original <code>/etc/nginx/nginx.conf</code> file.
How can I resolve the following issue in the <code>vpfs-config.log</code> file that is raised from? Verifying that the MSDP REST API is available via https on port 10087	Perform the steps in the following order: 1 Install the <code>polycoreutils</code> and <code>polycoreutils-python</code> packages through the yum tool. 2 As required by SELinux for NGINX, add the following rules to bind on the 10087 port. ■ <code>semanage port -a -t http_port_t -p tcp 10087</code> ■ <code>setsebool -P httpd_can_network_connect 1</code> 3 Run the following command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Frequently asked questions

Answer

Instant Access for BYO uses a self-signed certificate by default and only supports *.pem external certificate.

How do I replace it with a certificate signed by external CA (*.pem certificate), if required?

To configure the external certificate, perform the following steps:

- 1** If the new certificate is already generated (the certificate must contain long and short host names for the media server), go to step 5.
- 2** Create the RSA public or private key pair.
- 3** Create a certificate signing request (CSR).

The certificate must contain long and short host names for the media server.
- 4** The external Certificate Authority creates the certificate.
- 5** Replace <PDDE Storage Path>/spws/var/keys/spws.cert with the certificate and replace <PDDE Storage Path>/spws/var/keys/spws.key with the private key.
- 6** Run the following command to reload the certificate:


```
/usr/opensv/pdde/vpfs/bin/vpfs_config.sh  
--configure_byo
```

How can I disable media automount for the instant access livemount share in gnome?

If the automount is enabled, the source folder is mounted from the livemount share in gnome and smaller disks appear. In this scenario, the instant access feature does not work properly.

The mounted disk content source is from the .../meta_bdev_dir/... folder under livemount share, while the mount target is in the /run/media/... folder.

Follow the guideline to disable the gnome automount:

<https://access.redhat.com/solutions/20107>

Frequently asked questions**Answer**

How can I resolve the following issue in the `/var/log/vpfs/vpfs-config.log` file?

```
**** Asking the NetBackup  
Webservice to trust the MSDP  
webserver (spws) ****  
/usr/opensv/netbackup/bin/nblibcurlcmd  
failed (1):
```

Perform the steps in the following order:

1 Ensure that your NetBackup primary server is up and there is no firewall blocking the connection between the primary server and storage server.

2 Run the following command on storage server to verify the connection status:

```
/usr/opensv/netbackup/bin/bpclntcmd -pn
```

3 Wait for the NetBackup primary server to start and for an established connection between the NetBackup primary server and storage server. Then run the following command:

```
/usr/opensv/pdde/vpfs/bin/vpfs_config.sh  
--configure_byo
```

NetBackup for Oracle with Snapshot Client

This chapter includes the following topics:

- [About NetBackup for Oracle with Snapshot Client](#)
- [How NetBackup for Oracle with Snapshot Client works](#)
- [About configuring Snapshot Client with NetBackup for Oracle](#)
- [Restoring NetBackup for Oracle from a snapshot backup](#)
- [About configuring NetBackup for Oracle block-level incremental backups on UNIX](#)
- [About Snapshot Client effects](#)
- [About Oracle support for Replication Director](#)

About NetBackup for Oracle with Snapshot Client

To use NetBackup for Oracle with Snapshot Client, NetBackup Snapshot Client and NetBackup for Oracle must both be licensed and installed.

Before you use NetBackup for Oracle with Snapshot Client, confirm that your platform is supported.

See [“Verifying the operating system and platform compatibility”](#) on page 19.

A snapshot is a disk image of the client’s data that is made almost instantaneously. When it is used with NetBackup Snapshot Client, NetBackup for Oracle can back up Oracle objects by taking snapshot images of the component files. Later, it backs up the snapshot version to the storage unit.

Snapshot backup captures the data at a particular instant without having caused significant client downtime. Client operations and user access continue without interruption during the backup. The resulting capture or snapshot can be backed up without affecting the performance or availability of the database.

The following NetBackup Snapshot Client features are available for use with NetBackup for Oracle.

Table 11-1 Snapshot Client features used with NetBackup for Oracle

Feature	Description
Instant recovery	This feature enables instant recovery of backups from disk. It combines snapshot technology with the ability to do rapid disk-based restores. NetBackup creates the image without interrupting user access to data. Optionally, the image is retained on disk as well as backed up to storage. Instant recovery enables block-level restores. The maximum number of instant recovery snapshots to be retained at one time is calculated per client and database name. With the <code>remote_vxfs</code> method, the number of snapshots to be retained at one time is calculated per client, database name, and NetBackup appliance.
Off-host backup	An off-host backup shifts the burden of backup processing onto a separate backup agent, such as an alternate client. This shift reduces the effect on the client's computing resources ordinarily caused by a local backup. The backup agent reads the data from the client disk and writes it to storage.
Block-level incremental backup	On UNIX, a Block-Level Incremental (BLI) Backup uses the change tracking capabilities of the Arctera File System (VxFS) Storage Checkpoint feature. In a BLI backup, only the changed blocks of data are backed up, not the entire file or file system. A BLI backup saves time, decreases the amount of backup media that is required, and significantly reduces CPU and network overhead during backups. You can perform a BLI backup with or without RMAN.

Table 11-1 Snapshot Client features used with NetBackup for Oracle
(continued)

Feature	Description
Proxy copy	A proxy copy is a special type of backup in which the NetBackup for Oracle agent manages the control of the data transfer. During the backup and restore operations, the proxy copy enables the agent to manage the entire data movement between the disks that contain the data files and the storage devices that NetBackup manages. Backups and restores remain tightly integrated with Oracle and its catalog, greatly simplifying administration tasks.
File-based operations	Oracle provides the list of files that require backup or restore to NetBackup for Oracle with Snapshot Client. See “NetBackup for Oracle file-based operations” on page 178.
Stream-based operations	Stream-based operations are the standard NetBackup implementation of conventional NetBackup for Oracle backup and restore. See “NetBackup for Oracle stream-based operations” on page 177.

Proxy copy

A proxy copy is a special type of backup in which the NetBackup for Oracle agent manages the control of the data transfer. During the backup and restore operations, proxy copy enables the agent to manage the entire data movement between the disks that contain the data files and the storage devices that NetBackup manages.

With proxy copy, RMAN provides a list of files that require backup or restore to the NetBackup for Oracle agent. The agent determines how the data is moved and when to move the data. Proxy copy is an extension to Oracle’s Media Management API.

Backups and restores remain tightly integrated with RMAN and its catalog, which greatly simplifies administration tasks.

See [“Optimizing and deduplicating stream-based and proxy copy Oracle backups”](#) on page 240.


See [“Proxy backup examples”](#) on page 184.

NetBackup for Oracle stream-based operations

Stream-based operations are the standard NetBackup implementation of conventional RMAN backup and restore. In a stream-based backup, NetBackup moves the data that the server process provides. NetBackup captures the data stream content that RMAN provides. If the user has specified multiple streams, then RMAN opens multiple streams and NetBackup catalogs them as separate images.

[Figure 11-1](#) represents a stream-based backup or restore.

Figure 11-1 NetBackup for Oracle RMAN stream-based backup or restore



See [“NetBackup for Oracle file-based operations”](#) on page 178.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“About NetBackup for Oracle with Snapshot Client”](#) on page 174.


See [“Proxy copy”](#) on page 176.

NetBackup for Oracle file-based operations

File-based operations are the NetBackup for Oracle with Snapshot Client implementation of Oracle proxy copy backups and restores. In a file-based operation, RMAN provides the list of files that require backup or restore to NetBackup for Oracle with Snapshot Client. NetBackup for Oracle with Snapshot Client performs the data movement.

Figure 11-2 represents a file-based backup or restore.

Figure 11-2 NetBackup for Oracle with Snapshot Client file-based backup or restore



See [“NetBackup for Oracle stream-based operations”](#) on page 177.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Proxy copy”](#) on page 176.

How NetBackup for Oracle with Snapshot Client works

NetBackup users or schedules start database backups or restores. The Oracle Intelligent Policy automatically generates an RMAN script. The script-based policy uses a shell script in the backup selections list of the Oracle policy. The shell script specifies backup or restore commands for the Oracle Recovery Manager (RMAN) to use when you perform the backup or restore on the client.

The RMAN `backup proxy` command initiates a proxy copy backup of the specified objects. The objects that can be backed up using the proxy copy functionality depend on the Oracle version. RMAN translates the objects into the physical file names and provides a list of file names to NetBackup for Oracle.

See [“Database objects supported by advanced backup methods”](#) on page 180.

The agent checks that the policy it uses for the backup is configured with the appropriate Snapshot Client attributes. The agent then initiates file-based backups of the Oracle files and uses the NetBackup Snapshot Client interface to perform the data movement.

When Oracle performs proxy copy backups, it puts the data files being backed up into backup mode. NetBackup then creates a snapshot of the files. After the snapshot has been created, the NetBackup for Oracle agent signals back to Oracle to take the data files out of backup mode. The data files being backed up are in backup mode only for the period of time necessary to capture a snapshot of the data.

See [“About the NetBackup for Oracle backup and restore operations”](#) on page 180.

See [“About configuring Snapshot Client with NetBackup for Oracle”](#) on page 187.

See [“Database objects supported by advanced backup methods”](#) on page 180.

See [“About NetBackup multistreaming”](#) on page 181.

See [“RMAN multiple channels”](#) on page 181.

See [“Restoring data files to a new location”](#) on page 182.

See [“Redirecting a restore to a different client”](#) on page 182.

See [“Symbolic links and raw data files \(UNIX\)”](#) on page 183.

See [“Quick I/O data files \(UNIX\)”](#) on page 183.

See [“RMAN incremental backups”](#) on page 183.

About the NetBackup for Oracle backup and restore operations

For a backup operation, the NetBackup for Oracle agent performs the following steps:

- Receives a list of files to back up from RMAN.
- A unique backup file name identifies each file in the NetBackup catalog. To ensure that this procedure occurs, use the `format` operand to give each data file a unique name.
- Queries the policy to check whether the Snapshot Client policy attributes are specified.
- Initiates a configured number of Snapshot Client backups and waits until the jobs are completed.

See [“About NetBackup multistreaming”](#) on page 181.

For a restore operation, the NetBackup for Oracle agent performs the following steps:

- Receives a list of files to restore from RMAN.
- Sends a restore request to the NetBackup server for all files in the list.
- Waits for NetBackup to restore all files in the file list.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Database objects supported by advanced backup methods”](#) on page 180.

See [“About NetBackup multistreaming”](#) on page 181.

See [“RMAN multiple channels”](#) on page 181.

See [“RMAN incremental backups”](#) on page 183.

See [“Proxy backup examples”](#) on page 184.

Database objects supported by advanced backup methods

Oracle controls the kinds of database objects that can be backed up by proxy copy and, therefore, what NetBackup can back up using Snapshot Client backup methods. Oracle allows proxy copy backups of databases, tablespaces, and data files. With Oracle 10g releases and later, Oracle also allows proxy copy backups of archived redo logs. As a result, NetBackup can use file-based Snapshot Client backup methods to back up these objects.

For control files, Oracle RMAN performs conventional stream-based backups only. NetBackup for Oracle must use stream-based backups for control files even when you use Snapshot Client methods for the other database objects.

The Oracle Intelligent Policy handles both stream-based and file-based components. File-based and stream-based backups require different configurations. When configuring your NetBackup for Oracle with Snapshot Client backups, be sure to configure a policy that allows both stream-based and file-based backups.

See [“About the NetBackup for Oracle backup and restore operations”](#) on page 180.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“About NetBackup for Oracle with Snapshot Client”](#) on page 174.

See [“Proxy backup examples”](#) on page 184.

See [“NetBackup for Oracle stream-based operations”](#) on page 177.

See [“NetBackup for Oracle file-based operations”](#) on page 178.

About NetBackup multistreaming

On the initial call, NetBackup for Oracle with Snapshot Client returns a special entry to RMAN indicating that it supports proxy copy. It also indicates to RMAN that it supports an unlimited number of files to be proxy-copied in a single proxy copy session. The number of channels that are allocated for the RMAN `backup proxy` command does not control the degree of parallelism for proxy backups. RMAN uses only one channel for proxy copy backups except when a specific configuration is used.

The `NB_ORA_PC_STREAMS` variable controls the number of proxy copy backup streams to be started. By default, the agent initiates one backup job for all files. If the RMAN `send` command passes `NB_ORA_PC_STREAMS`, NetBackup for Oracle splits the files into the number of groups that the variable specifies based on the file size. The agent attempts to create streams of equal size and determines the number of processes that run to perform the backup.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“About NetBackup for Oracle with Snapshot Client”](#) on page 174.

See [“Proxy backup examples”](#) on page 184.

RMAN multiple channels

If you allocate multiple channels for an RMAN proxy copy backup session, RMAN uses only one channel to perform a proxy backup of all objects. All other channels can be used for a stream-based (non-proxy) backup of the control file or archived redo logs.

See [“Proxy backup examples”](#) on page 184.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“RMAN incremental backups”](#) on page 183.

Restoring data files to a new location

NetBackup for Oracle with Snapshot Client can restore the data files that are backed up by proxy to a new location. The new location can be specified by using the RMAN `set newname` command or `ALTER DATABASE RENAME DATAFILE` statement before a restore is initiated. For example, to restore a data file for tablespace `TEST` to a new location, you can use the following RMAN commands:

```
RUN
{
  allocate channel t1 'SBT_TAPE';
  sql 'alter tablespace TEST offline immediate'
  # restore the datafile to a new location
  set newname for datafile '/oradata/test.f' to
  '/oradata_new/test.f';
  restore tablespace TEST;
  # make the control file recognize the restored file as current
  switch datafile all;
  recover tablespace TEST;
  release channel t1;
}
```

The RMAN procedure for the data files that are backed up by proxy is the same as for conventionally backed up data files. RMAN knows that the data files were backed up by proxy, and it issues a proxy restore request to NetBackup for Oracle, which restores the data files to the new location. For more information on the required procedure, see your Oracle documentation.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Redirecting a restore to a different client”](#) on page 182.

See [“Performing a snapshot rollback restore using a script or RMAN command”](#) on page 194.

Redirecting a restore to a different client

The procedure for restoring a proxy backup to a different destination client is the same as the procedure for stream-based, non-proxy backups.

Symbolic links and raw data files (UNIX)

NetBackup for Oracle with Snapshot Client backs up and restores the data files that consist of symbolic links and regular files. Both the symbolic link and the file are backed up and restored. However, if you selected **Retain snapshots for instant recovery** then the symbolic link must reside on the same file system as the data file. When you use instant recovery, if the symbolic link resides on a different file system than the data file it points to, the restore fails.

NetBackup for Oracle with Snapshot Client backs up and restores data files created on raw partitions.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Quick I/O data files \(UNIX\)”](#) on page 183.

Quick I/O data files (UNIX)

NetBackup for Oracle with Snapshot Client backs up and restores Quick I/O Oracle data files. A Quick I/O file consists of two components: a hidden file with space allocated for it and a link that points to the Quick I/O interface of the hidden file.

On the backup, NetBackup for Oracle with Snapshot Client follows the symbolic link and backs up both components of the Quick I/O file: the symbolic link and the hidden file.

On the restore, NetBackup for Oracle with Snapshot Client restores both components from the backup image. If one or both of the components are missing, NetBackup for Oracle with Snapshot Client creates the missing component(s).

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Symbolic links and raw data files \(UNIX\)”](#) on page 183.

RMAN incremental backups

You can use proxy copy backups as a part of the incremental strategy with conventional non-proxy RMAN backups. RMAN lets you create a proxy copy incremental level 0 backup. This backup can be the base for subsequent RMAN traditional incremental backups (level 1-*n*). To accomplish this backup, perform a snapshot proxy copy (file-based) level 0 incremental backup and follow with an RMAN traditional (stream-based) level 1-*n* incremental backup.

In Oracle 10g it is possible to track changed blocks using a change tracking file. Enabling change tracking does produce a small amount of database overhead, but it greatly improves the performance of incremental backups. Use the `ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;` `sqlplus` command to enable block change tracking on the database.

In the following example, the first `run` command initiates a proxy copy backup of tablespace `tbs1`. NetBackup for Oracle uses a snapshot file-based backup to perform a full tablespace backup. RMAN designates this backup as eligible for incremental level 1-*n* backups. The second `run` command initiates a traditional non-proxy level 1 incremental backup of the same tablespace `tbs1`. In this case, NetBackup for Oracle performs a stream-based backup.

```
run {
  allocate channel t1 type 'SBT_TAPE';
  backup
    incremental level 0
    proxy
    format 'bk_%U_%t'
    tablespace tbs1;
  release channel t1;
}

run {
  allocate channel t1 type 'SBT_TAPE';
  backup
    incremental level 1
    format 'bk_%U_%t'
    tablespace tbs1;
  release channel t1;
}
```

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“RMAN multiple channels”](#) on page 181.

See [“Proxy backup examples”](#) on page 184.

Proxy backup examples

The Oracle Intelligent Policy automatically creates the RMAN proxy script. In some instances, you need to create a custom script specific to your environment.

The following examples show how to use multiple channels in RMAN scripts with proxy backups.

Table 11-2 Proxy backup examples

Backup example	Sample script
This RMAN sample script initiates a whole database backup, which includes the control file. RMAN starts one proxy copy backup session by sending a list of all data files to the NetBackup for Oracle agent on channel t1.	<pre>run { allocate channel t1 type 'SBT_TAPE'; send 'NB_ORA_PC_STREAMS=3'; backup proxy format 'bk_%U_%t' (database); release channel t1; }</pre> The agent splits the files into three streams and initiates a file-based backup for each stream. After the proxy backup is done, RMAN starts a non-proxy conventional backup of the control file on channel t1.
This RMAN sample script initiates a whole database backup, which includes the control file. RMAN starts one proxy copy backup session by sending a list of all data files to the NetBackup for Oracle agent on channel t1. The agent splits the files into three streams and initiates a file-based backup for each stream. At the same time, RMAN starts a non-proxy conventional backup of the control file on channel t2.	<pre>run { allocate channel t1 type 'SBT_TAPE'; allocate channel t2 type 'SBT_TAPE'; send 'NB_ORA_PC_STREAMS=3'; backup proxy format 'bk_%U_%t' (database); release channel t1; release channel t2; }</pre> If the RMAN recovery catalog is not used, the version of the control file being backed up does not contain information about the current backup. To include the information about the current backup, back up the control file as the last step in the backup operation. This step is not necessary if the recovery catalog is used. <pre>Run { allocate channel t1 type 'SBT_TAPE'; backup format 'cntrl_%s_%p_%t' current controlfile; release channel t1; }</pre>

Table 11-2 Proxy backup examples (*continued*)

Backup example	Sample script
In this sample script, RMAN initiates two proxy copy backups sequentially on channel t1. It starts a proxy backup of tablespace tbs1 data files. After the backup is done, it starts another proxy backup of tablespace tbs2 data files.	<pre>run { allocate channel t1 type 'SBT_TAPE'; backup proxy format 'bk_%U_%t' (tablespace tbs1); backup proxy format 'bk_%U_%t' (tablespace tbs2); release channel t1; }</pre> This configuration can cause problems if the sequential backups create snapshots on the same or a separate volume that share a snapshot resource specification. In such a situation, issue a single <code>backup</code> command such as the following which specifies both tablespaces rather than two separate <code>backup</code> commands: <pre>run { allocate channel t1 type 'SBT_TAPE'; backup proxy format 'bk_%U_%t' (tablespace tbs1, tbs2); release channel t1; }</pre>
In this example, RMAN distributes proxy copy backups over two channels. It creates two proxy copy backup sessions sending tbs1 data files on channel t1 and tbs2 data files on channel t2. Such a method is useful if you want to specify different NetBackup configurations for each channel. In this example, each <code>send</code> command specifies a different policy that is sent to the proxy backups. Each of the proxy backups uses this policy.	<pre>run { allocate channel t1 type 'SBT_TAPE'; send 'NB_ORA_POLICY=policy1'; allocate channel t2 type 'SBT_TAPE'; send 'NB_ORA_POLICY=policy2'; backup proxy format 'bk_%U_%t' (tablespace tbs1 channel t1); (tablespace tbs2 channel t2); release channel t1; release channel t2; }</pre>

See [“RMAN incremental backups”](#) on page 183.

See [“Database objects supported by advanced backup methods”](#) on page 180.

See [“RMAN multiple channels”](#) on page 181.

See [“About NetBackup multistreaming”](#) on page 181.

See [“RMAN incremental backups”](#) on page 183.

About configuring Snapshot Client with NetBackup for Oracle

This topic explains how to configure snapshot and instant recovery backups for the Oracle policy. For information on how a snapshot method is automatically selected and details on the types of backup methods, see the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

Snapshot backups do not back up all database objects. Your backup configuration must include one or more automatic schedules to perform snapshot backups and one or more application schedules to perform stream-based backups. This configuration ensures that the entire database can be restored successfully.

For snapshot or instant recovery backups, configure the following policies and schedules as follows:

- A Oracle policy with the following attributes:
 - Snapshot methods for the file systems in which the database files reside.
 - A backup method on the policy attributes dialog box.
 - An Automatic Full Backup schedule to perform snapshot and off-host backups of the database.
 - (Conditional) For script-based policies: An Application Backup schedule to back up the transaction logs.

To use NAS snapshot with NetBackup for Oracle, the Oracle database must be installed and configured to work in a NAS environment.

If you want to use a SnapVault storage unit, make sure that the storage unit is configured before you start configuring the NAS snapshot policy.

Configuration requirements for snapshot backups with NetBackup for Oracle

Each agent has its own hardware requirements, software requirements, compatibility with certain features, and the snapshot methods that are supported. Special requirements apply for specific types of backups. See the Cohesity Support website for more information. Familiarize yourself with this information before you configure any snapshot backups.

The following list highlights some of the requirements that pertain to database agents:

- Snapshot Client backups do not back up all database objects. Your backup configuration must include schedules to perform snapshot and stream-based backups. This configuration ensures that the entire database can be restored successfully.
- On UNIX, the user identification and group identification numbers (UIDs and GIDs) associated with the files to be backed up must be available. The UID and GID must be available to both the primary client and the alternate backup client. The UID on the primary client and the alternate backup client must be the same. Similarly, the GID on the primary client and the alternate backup client must be the same.

Note: The UID number can be different than the GID number.

- Allocate different areas for data files, archived redo logs, and the control file for database activities. Write the data files to their own repository because it is required for an instant recovery point-in-time rollback. Only data files can exist on the volume or the file system that you want to restore.
- The hardware and software that is required for the appropriate snapshot method must be installed and configured correctly.
- NetBackup Snapshot Client must be installed and configured correctly, and the license for this option must be registered.
- To perform off-host backups, perform any special configuration that is required.

See [“About configuring Snapshot Client with NetBackup for Oracle”](#) on page 187.

See [“Configuring a snapshot policy for NetBackup for Oracle”](#) on page 188.

Configuring a snapshot policy for NetBackup for Oracle

The following procedure shows how to configure a snapshot policy with optional instant recovery, snapshot retention, and off-host backup.

This procedure does not detail how to configure a snapshot policy when a NetBackup appliance is used.

See [“Configuring a snapshot policy using a share on the NetBackup appliance \(Oracle Copilot\)”](#) on page 191.

To configure a snapshot policy

- 1** In the web UI, open the policy you want to configure.
- 2** Click on the **Attributes** tab.
- 3** Select the **Oracle** policy type.
- 4** Select a policy storage unit from the **Policy storage** list.

Select a policy storage unit in this step even if you plan to select **Instant Recovery Snapshots Only** later in this procedure.

NetBackup uses this storage unit for the stream-based backups of the control files and the archived redo logs this policy includes.

- 5** Click **Perform snapshot backups**.
- 6** (Optional) Click **Snapshot options** to choose a snapshot method.

By default NetBackup chooses a snapshot method for you. To choose a snapshot method, click **auto** (the default) or click one of the methods that are presented in the list.

The snapshot method that you can use depends on your hardware environment and software environment. Only certain snapshot methods are supported in certain environments. See the Cohesity Support website for more information.

You can configure only one snapshot method per policy. For example, assume that you want one snapshot method for clients a, b, and c, and a different method for clients d, e, and f. Then you need to create two policies for each group of clients and select one method for each policy.

- 7** (Optional) Select **Retain snapshot for Instant Recovery or SLP management**.

When this option is selected, NetBackup retains the snapshot backup image on disk for later use in recovery.

- 8** (Optional) Select **Perform off-host backup**.

By default, the client that hosts the database performs the backup. If you want to reduce the I/O processing load on the client that hosts the database, specify an alternate client to perform the backup.

9 (Conditional) Select an off-host backup method.

The following off-host backup methods are available:

Alternate client (UNIX and Windows clients) If you select **Alternate client**, also specify the name of the client to perform the backup. This option may require additional configuration. The alternate client must be a client that shares the disk array.

Data mover (UNIX clients only) If you click **Data mover**, select **Network Attached Storage**.

10 Click the **Schedules** tab.

11 Click **Add**.

12 Configure a schedule for the database files.

13 (Conditional) To create only disk images, under **Instant Recovery** select **Snapshots only**.

This setting suppresses NetBackup's default behavior, which is to copy the snapshot to a storage unit. When you select **Snapshots only**, NetBackup creates the on-disk snapshot copy of the database, but it does not copy the snapshot to a storage unit. The on-disk snapshot becomes the only backup copy. Note that the on-disk snapshot is not considered to be a replacement for a traditional backup.

14 (Conditional) On the **Schedules** tab, configure a backup schedule for the control files or archived redo logs.

- Oracle Intelligent Policy backup policy. Configure an **Archived Redo Log Backup** schedule for this policy.
- Script-based backup policy. Configure an **Application Backup** schedule for this policy.

NetBackup uses this storage unit for the stream-based backups of the control files and the logs that are included in this policy. NetBackup copies the database's control files and archived redo logs to the storage unit you select.

15 Script-based backup policy. On the **Clients** tab, specify the clients to be included in this policy.

16 On the **Backup Selections** tab, specify the correct setup depending on policy setup.

- Oracle Intelligent Policy backup policy. Use the radio button to select **Whole Database**, **Partial database – Tablespace**, **Partial database – Datafiles**,

Fast Recovery Area, Database backup shares, or Whole Database - Datafile Copy Share when using this type of policy.

- Script-based backup policy. Specify the backup script when you use this type of policy.

More information is available about how to use scripts for a NetBackup for Oracle policy with Snapshot Client.

See [“About Snapshot Client effects”](#) on page 200.

- 17** Configure other attributes and add any additional schedules and backup selections.

Configuring a snapshot policy using a share on the NetBackup appliance (Oracle Copilot)

Note: This feature requires a NetBackup appliance that runs software version 2.7.1 or later.

Use the following procedure to configure an Oracle snapshot policy that uses **Database Backup Shares** or **Whole Database - Datafile Copy Share** options. This procedure uses the Oracle Intelligent Policy, which makes configuration easier.

To configure a snapshot policy using a NetBackup appliance share

- 1** In the web UI, open the policy you want to configure or create a new policy.
- 2** Select the **Attributes** tab.
- 3** Select **Oracle** as the policy type if this policy is new.
- 4** Select a policy storage unit from the **Policy storage** list.

Oracle combines snapshots (proxy) and stream-based backups as part of the same backup. The storage that is indicated here is used for the stream-based part of the **Database backup shares** or **Whole Database - Datafile Copy Share** backup.

Select a storage lifecycle policy that is configured to contain the stream-based (non-snapshot) part of the database backup. The storage must use a storage lifecycle policy that is configured for non-snapshot backups.

- 5** Select **Perform snapshot backups**.

- 6 Click **Snapshot options** to choose a snapshot method.

When you use the **Database backup shares** or the **Whole Database - Datafile Copy Share** option, `remote_vxfs` is the only valid snapshot method.

Note: If there is more than one backup share that is associated with the database, then **Maximum Snapshots** should be set to number of recovery points you want multiplied with number of backup shares. Example: If you want three recovery snapshot points and the database is associated with two backup shares then **Maximum Snapshots** should be set to 6.

- 7 Select **Retain snapshot for Instant Recovery or SLP management**.

- 8 Select the **Schedules** tab.

- 9 Click **Add**.

- 10 Configure a **Full** schedule for the database backup shares.

- **Type of backup:** Select **Full backup**. The **Full backup** is used for both the snapshot part of the database and the non-snapshot (stream-based) part of the Oracle database.
- **Override policy storage selection:** Enable and select the SLP that is configured for a snapshot. (A snapshot SLP is one in which the first operation is a snapshot operation.) This option must be enabled so that the schedule storage overrides the policy storage with a snapshot SLP.
- **Retention:** The retention for the streamed data is based on the non-snapshot SLP that was indicated as the **Policy storage** in Step 4.
 - The non-snapshot SLP specified on the policy storage in Step 4 determines the retention for the streamed data.
 - The snapshot SLP that is specified as the schedule storage (**Override policy storage selection**) determines the retention for the snapshot data.

When **Database backup shares** or **Whole Database - Datafile Copy Share** is selected, it is recommended that an SLP is set up to backup from the snapshots and copy snapshots to a storage unit. NetBackup retains the snapshot backup image on disk for later use in SLP management.

Click **Add** to save the schedule.

- 11 (Optional) On the **Schedules** tab, configure an **Archived redo log backup** schedule for the archived redo logs.

- 12 Select the **Instances and databases** tab and specify the databases to back up. The policy must include at least one database. To continue to use the Oracle Intelligent Policy (OIP) method, select **Protect non-RAC databases and pluggable databases**.
- 13 On the **Backup selections** tab, use the radio button to select **Database Backup Shares** or **Whole Database - Datafile Copy Share** options.
- 14 (Optional) Configure other attributes and add any additional schedules.

Note: The NetBackup appliance media server can only be configured in a pure IPv4 or dual stack IPv4/IPv6 mode.

Restoring NetBackup for Oracle from a snapshot backup

The following topics describe how to restore files, volumes, and file systems from a snapshot backup:

- See [“About restoring individual files from a NetBackup for Oracle snapshot backup”](#) on page 193.
- See [“About NetBackup for Oracle restores of volumes and file systems using snapshot rollback”](#) on page 194.
- See [“Performing a NetBackup for Oracle point-in-time rollback restore from a SnapVault backup \(UNIX\)”](#) on page 195.

About restoring individual files from a NetBackup for Oracle snapshot backup

Data that is backed up with Snapshot Client methods is restored in the same way as data that is backed up without Snapshot Client methods.

Use this procedure for the files that were backed up with, or without, instant recovery enabled. In all cases, Oracle determines the files that were backed up, and it initiates a corresponding restore request to the database agent.

If instant recovery is enabled, NetBackup attempts to restore the file by using the unique restore methods available with the instant recovery feature. The type of restore method that NetBackup uses depends on your environment and the type of backup performed. If NetBackup is unable to use any of the instant recovery methods, it restores the file in the typical manner. Data is copied from the snapshot to the primary file system. Information on the instant recovery methods that NetBackup

uses is available. See the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

About NetBackup for Oracle restores of volumes and file systems using snapshot rollback

You can request that an entire volume or an entire file system be restored from an instant recovery Snapshot backup. This type of a restore is called a point in time rollback. All the data in the snapshot is restored; single file restore is not available in a rollback.

See the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

The following considerations are relevant for NetBackup for Oracle restores:

- Snapshot rollback overwrites the entire volume.
 - With NetBackup for Oracle, snapshot rollback always performs file verification. The agent checks for the following:
 - The requested files (number and names) are identical to those in the snapshot
 - The primary volume does not contain any files that were created after the snapshot was made
- If verification fails, the rollback aborts with status 249.

Performing a snapshot rollback restore using a script or RMAN command

This topic describes how to perform a snapshot rollback restore using a script or RMAN command.

Note: If the data file you want to restore has not changed since it was backed up, the rollback may fail. Initiate the restore from a script and use the Oracle FORCE option.

To specify a snapshot rollback restore using a script or RMAN command, follow this example:

- If you want to use a shell script or RMAN command, set a new variable, `NB_PC_ORA_RESTORE=rollback`
 - Example:

```
RUN {
    allocate channel t1 'SBT_TAPE';
    send 'NB_ORA_PC_RESTORE=rollback';
```

```

sql 'alter tablespace TEST offline immediate'
restore tablespace TEST;
recover tablespace TEST;
release channel t1;
}

```

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“About restoring individual files from a NetBackup for Oracle snapshot backup”](#) on page 193.

See [“About NetBackup for Oracle restores of volumes and file systems using snapshot rollback”](#) on page 194.

See [“Performing a NetBackup for Oracle point-in-time rollback restore from a SnapVault backup \(UNIX\)”](#) on page 195.

Performing a NetBackup for Oracle point-in-time rollback restore from a SnapVault backup (UNIX)

When you select a point-in-time rollback restore from a SnapVault backup, NetBackup restores the entire subvolume (qtree) to a new subvolume (qtree) on the primary host. The restore does not overwrite the existing subvolume. File verification is not performed.

The format of the new subvolume name is as follows:

mountpointname_restore.timestamp

For example: `subvol1_restore.2005.05.19.10h49m04s`

To perform a NetBackup for Oracle point-in-time rollback restore from a SnapVault backup (UNIX)

- 1 Unmount the original subvolume, which is the subvolume that the restore process did not overwrite.
- 2 Rename the original subvolume.
- 3 Rename the new subvolume with the name of the original.
- 4 Mount the new subvolume on the client. Use the `ALTER DATABASE RENAME DATAFILE` command to point to the restored data file on the newly created subvolume.

About configuring NetBackup for Oracle block-level incremental backups on UNIX

If only a small portion of a database changes on a daily basis, full database backups are costly in terms of time and media. The Block-Level Incremental (BLI) Backup interface extends the capabilities of NetBackup to back up only the file system blocks that contain changed data blocks.

A database BLI backup is done at the file system block level, which means only changed file blocks are backed up. Unchanged blocks within the files are not backed up. The VxFS Storage Checkpoint facility tracks changed blocks in real time. Accordingly, a BLI backup does not need to search the entire volume for the modified blocks at backup time. BLI backup saves time, decreases the amount of backup media that is required, and significantly reduces CPU and network overhead during backups. In addition, BLI backup allows more frequent backups, so backup images are more up to date.

BLI backup is particularly useful for any large databases that are sized in terms of hundreds of gigabytes or terabytes. Most traditional methods for database backup require that any change in the database—no matter how small—requires that the entire database is backed up. With BLI backup, only modified blocks (or file) need to be backed up.

The recommended method for performing BLI backups is the proxy BLI agent with RMAN. This method supports the other features of NetBackup for Oracle, including the policy types and schedules. It also remains tightly integrated with RMAN and its catalog, which greatly simplifies administration tasks.

You can also perform backups with the script-based BLI method without RMAN.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Note: Cohesity recommends that Snapshot Client users who want to perform BLI backups use BLI with RMAN.

NetBackup for Oracle also provides a method for BLI backup without RMAN that uses scripts to put tablespaces into, and take them out of, backup mode. This method is not recommended, and it requires a significantly different configuration. But for Oracle 12c, using script-based BLI backups without the use of RMAN are not supported.

See [“Configuring policies for BLI backups with NetBackup for Oracle”](#) on page 199.

How BLI works with NetBackup for Oracle (UNIX)

NetBackup supports BLI full backups and BLI incremental backups of Oracle databases.

BLI backup supports two types of incremental backups: differential and cumulative. Full, differential incremental, and cumulative incremental backups are specified as part of the policy schedule configuration. When a restore is performed, NetBackup restores an appropriate full backup. Then it applies the changed blocks from the incremental backups.

Restoring any of the incremental backup images requires NetBackup to restore the last full backup image and all the subsequent incremental backups. The restore process continues until the specified incremental backup image is restored. NetBackup performs this restore process automatically, and it is completely transparent. The media that stored the last full backup and the subsequent incremental backups must be available, or the restore cannot proceed.

Note that restoring a file rewrites all blocks in that file. The first subsequent differential incremental backup and or all subsequent cumulative incremental backups back up all the blocks in the restored file. After an entire database is restored, the first subsequent backup results in a full backup.

The restore destination can be a VxFS, UFS (Solaris), JFS (AIX), or HFS (HP-UX) file system. The destination VxFS file system does not need to support the Storage Checkpoint feature to restore files. However, a VxFS file system with the Storage Checkpoint feature is needed to perform BLI backups of the restored data.

This topic uses the following terms to describe BLI backups:

- **Full Backup.**
A backup in which NetBackup backs up each database file completely, not just data blocks that have changed since the last full or incremental backup.
- **Cumulative BLI Backup.**
This type of backup is a backup of all the changed blocks in the database files since the last full backup. A cumulative BLI backup image contains only the data blocks of database files that changed since the last full backup. A cumulative BLI backup can reduce the number of incremental backup images that must be applied during a restore operation. This speeds up the restore process.
- **Differential BLI backup.**
A backup in which NetBackup performs a backup of only those data blocks (within the database files) that changed since the last backup. The previous backup can be of type full, cumulative incremental, or differential incremental.

When NetBackup initiates BLI backups, it creates, manages, and uses the appropriate Storage Checkpoints of the filesystem(s) hosting the Oracle data file systems. These Storage Checkpoints identify and maintain a list of modified blocks.

About the Storage Checkpoint facility and NetBackup for Oracle

The BLI backup methodology uses the Storage Checkpoint facility in the Veritas File System (VxFS). This facility is available through the Storage Foundation for Oracle.

The VxFS Storage Checkpoint facility keeps track of the file blocks modified by the database since the last backup. NetBackup with BLI backup leverages this facility to back up only changed blocks for an incremental backup. The entire volume or file is not backed up.

VxFS Storage Checkpoint is a disk-efficient and I/O-efficient snapshot of file systems. A Storage Checkpoint provides a consistent, stable view of a file system at the instant when the file system was snapped or checkpointed. Instead of making a physically separate copy of the file system, a Storage Checkpoint tracks changed file system blocks. Disk space is saved and I/O overhead is significantly reduced.

Because the changed blocks are tracked, the VxFS Storage Checkpoint enables BLI backups. VxFS Storage Checkpoint facility provides a consistent view of file systems, which allows BLI backup to freeze the database image during database backups.

The Storage Checkpoint operation is similar to the snapshot file system mechanism. However, the Storage Checkpoint persists after a system restart which is unlike a snapshot. Also, the Storage Checkpoint operation is totally transparent to backup administrators. The Checkpoint image is managed and available only through NetBackup or through the VxDBA utility for database backup available with the Veritas Storage Foundation.

For more information on Storage Checkpoints, see the [Cohesity Storage Foundation Administrator's Guide](#).

You can take a Storage Checkpoint while the database is online or offline. To take a Storage Checkpoint while the database is online, you must enable archive log mode. During the creation of the Storage Checkpoint, all tablespaces are placed in backup mode.

Configuration requirements for BLI backups with NetBackup for Oracle

Before you configure BLI backups, make sure that your configuration meets the following requirements:

- NetBackup for Oracle is installed, licensed, and configured.
- NetBackup Snapshot Client is installed and configured, and the primary server must have a valid license for this option.
- Cohesity Storage Foundation for Oracle must be installed and configured.
- Cohesity File System must have Storage Checkpoint licensed.

For more information see the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

Configuring policies for BLI backups with NetBackup for Oracle

This topic explains how to configure BLI backups for Oracle policies. BLI backups do not back up all database objects. Include schedules to perform stream-based backups.

Your backup configuration must ensure that the entire database can be successfully restored.

See [“Configuration requirements for BLI backups with NetBackup for Oracle”](#) on page 198.

To configure a policy for BLI backups, configure the following:

- The BLI backup method on the policy attributes dialog box.
- An **Automatic backup** schedule to perform full and incremental snapshot backups of the data files.
- An **Application backup** schedule to perform a stream-based backup of control files and archived redo logs. These files are backed up with standard RMAN operations.

To configure a policy for BLI backups

- 1 On the left, select **Protection > Policies**.
- 2 Select **Add**.
- 3 Select the **Attributes** tab.
- 4 From the **Policy type** list, select **Oracle**.
- 5 Select the **Policy storage**.
- 6 Select **Perform block-level incremental backups**.
- 7 To configure schedules, select the **Schedules** tab.

Oracle does not support proxy backups of database control files and archived redo logs. To perform a whole database proxy backup, which automatically includes a backup of the control file, configure the following:

- One or more automatic backup schedules to perform proxy BLI backups of the data files.
- An Application backup schedule type to back up the control files and archived redo logs.

8 On the **Clients** tab, specify the clients to back up with this policy.

9 On the **Backup selections** tab, specify the script.

See [“About the Storage Checkpoint facility and NetBackup for Oracle”](#) on page 198.

See [“About configuring NetBackup for Oracle block-level incremental backups on UNIX”](#) on page 196.

About the types of NetBackup for Oracle BLI backups

NetBackup performs BLI backups with Automatic Full Backup, Automatic Differential Incremental Backup, and Automatic Cumulative Incremental Backup schedules.

If a user initiates a backup and the proxy schedule name is not specified on the request with the `NB_ORA_PC_SCHED` environment variable, the NetBackup server starts an Full Backup schedule by default.

NetBackup for Oracle checks that a full backup was performed before it proceeds with an incremental backup. If the NetBackup scheduler or user initiates an incremental backup, and NetBackup for Oracle finds no record of a full backup using the same policy, it performs a full backup.

To ensure that it has a proper set of images to restore, NetBackup performs a full backup when it encounters the following situations:

- If the number of backup streams that is specified changed from the previous backup. This change can be made in the `NB_ORA_PC_STREAMS` environment variable.
- If NetBackup does not have a valid full backup image for the same policy in its database. For example, this situation can occur if images were expired.
- If a new file was added to or deleted from the list of files for an incremental backup.

NetBackup for Oracle always initiates a full backup under these conditions, even if you want to perform an incremental backup.

About Snapshot Client effects

The following topics describe how the Snapshot Client software affects backup types, schedule properties, scripts, and environment variables.

How Snapshot Client software affects backup types

The backup types available on the **Schedules** tab of the policy play a different role for NetBackup for Oracle with Snapshot Client backups.

See [Table 11-3](#) on page 201.

Table 11-3 Backup types for Oracle policies

Backup type	Description
Application Backup	An application backup applies only to script-based policies not the Oracle Intelligent Policies. The Application Backup schedule stores stream-based backups. The Default-Application-Backup schedule is automatically configured as an Application Backup schedule.
Full backup Differential incremental backup, Cumulative incremental backup	The full and incremental backup schedule types automatically start the backups by running the NetBackup for Oracle RMAN scripts. They also store the snapshot backups. Note: For most snapshot types, any automatic backup schedule (full, cumulative, or differential) results in a full volume snapshot. BLI is the only snapshot method that can perform an incremental backup.

See [“How Snapshot Client software affects schedule properties”](#) on page 201.

See [“How Snapshot Client software affects scripts”](#) on page 202.

See [“ Oracle with Snapshot Client environment variables”](#) on page 202.

How Snapshot Client software affects schedule properties

Some schedule properties have a different meaning for Snapshot Client database backups than for a regular database backup. For a description of other schedule properties, see the information that is specific to standard database agent backups.

See [“Schedule properties for Oracle script-based policies ”](#) on page 80.

[Table 11-4](#) explains the properties for Snapshot Client backups.

Table 11-4 Schedule properties

Property	Description
Retention	Automatic Schedules: Determines how long to retain history of the backups that the primary server schedules and also how long to retain snapshot backups. Application Schedules: Determines how long to retain stream-based backups.
Multiple Copies	For snapshot backup, configure Multiple copies on the automatic backup schedule. For stream-based backups, configure Multiple copies on the Application backup schedule.
Frequency	Determines how often an Automatic schedule runs a backup. Does not apply to Application backup schedules.

How Snapshot Client software affects scripts

When you use a script, you must enable the advanced backup method for your clients. Configure this method on the **Attributes** tab of the policy. At run time, the agent checks the policy attributes to determine if a Snapshot Client backup method is configured and performs a proxy file-based backup.

See [“About creating shell scripts”](#) on page 93.

If you use a script, the script must reside on each client that is included in the policy. Include the `RMAN backup proxy` command in the script to perform the advanced backup method. Sample scripts are included with the installation.

See [“How Snapshot Client software affects backup types”](#) on page 201.

See [“Oracle with Snapshot Client environment variables”](#) on page 202.

Oracle with Snapshot Client environment variables

You can use environment variables to change the number of streams the proxy copy session uses or to specify an alternate backup schedule.

The following list shows the variables that you can set that are specific to NetBackup for Oracle with Snapshot Client:

NB_ORA_PC_SCHED

The NetBackup for Oracle schedule NetBackup uses for a proxy copy file-based backup. (This schedule can be Full, Differential Incremental, or Cumulative Incremental backup type).

For scheduled backups, this variable is passed from the scheduler.

NB_ORA_PC_STREAMS

Specifies the number of backup streams that NetBackup starts simultaneously in each proxy copy session. When a backup starts, NetBackup groups all data files into a specified number of backup streams that are based on the file sizes. NetBackup tries to create streams of equal size.

The default value for NB_ORA_PC_STREAMS is 1.

Only a user can set this variable.

For NetBackup for Oracle with Snapshot Client, the order of precedence for environment variables is the same as for standard NetBackup for Oracle. Refer to the instructions for how to configure the NetBackup and the user variables.

See [“About configuring the run-time environment”](#) on page 86.

NetBackup for Oracle installs sample scripts in the following location:

Windows:

`install_path\NetBackup\dbext\oracle\samples\rman`

UNIX:

`/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman`

The following are the scripts for NetBackup for Oracle with Snapshot Client that show how to configure the required variables:

Windows:

`hot_database_backup_proxy.cmd`

UNIX:

`hot_database_backup_proxy.sh`

This script sets the environment and calls RMAN with the appropriate command to perform a whole database proxy backup. When NetBackup runs a schedule, it sets the environment variables that NetBackup for Oracle with Snapshot Client uses. The script shows how to use the RMAN `send` command to pass the NetBackup for Oracle with Snapshot Client variables with a vendor-specific quoted string.

Windows: This script sets the environment and calls RMAN
hot_tablespace_backup_proxy.cmd with the appropriate command to perform a
tablespace proxy backup.

UNIX:
hot_tablespace_backup_proxy.sh

If you use scripts, use the `send` command to pass the environment variables to the agent. The following example uses the `send` command to specify the values for `NB_ORA_PC_SCHED` and `NB_ORA_PC_STREAMS`:

```
run {
  allocate channel t1 type 'SBT_TAPE';
  send 'NB_ORA_PC_SCHED= sched, NB_ORA_PC_STREAMS= number';
  backup proxy
    (database format 'bk_%U_%t');
}
```

For more information, see the sample scripts that are provided with the agent.

See [“How NetBackup for Oracle with Snapshot Client works”](#) on page 179.

See [“Performing a snapshot rollback restore using a script or RMAN command”](#) on page 194.

See [“Proxy backup examples”](#) on page 184.

See [“How Snapshot Client software affects backup types”](#) on page 201.

See [“How Snapshot Client software affects schedule properties”](#) on page 201.

See [“How Snapshot Client software affects scripts”](#) on page 202.

About Oracle support for Replication Director

Replication Director can be used to create snapshots of the Oracle database and replicate the snapshots to other NetApp disk arrays. To use Replication Director, the Oracle database must exist on a NetApp NAS disk array. (It is not supported on SAN storage at this time.)

Oracle snapshot backups that use Replication Director are supported on UNIX platforms only.

The administrator can create an Oracle policy to use Replication Director by using either the following methods:

- The Oracle Intelligent Policy (recommended).
See [“Configuring an Oracle Intelligent Policy using Replication Director”](#) on page 205.

- Create a script-based Oracle policy.
See [“Configure a script-based Oracle policy”](#) on page 209.

[Table 11-5](#) describes the differences between the two methods:

Table 11-5 Differences in Oracle snapshot policy setup

Configuration	Oracle Intelligent Policy	Script-based Oracle policy
Scripts	■ All scripts that are necessary to protect all parts of the database are automatically generated at run-time. ■ The administrator does not need to know how to configure RMAN scripts. ■ The retention levels for the different parts of the database are automatically assigned.	■ NetBackup can continue to use custom scripts to perform database backups. ■ The administrator must know how to configure RMAN scripts. ■ The administrator must set the retention levels for the different parts of the database correctly. ■ The administrator must ensure that a snapshot of the proxy data is created.
Schedules	The administrator configures only one schedule that backs up all parts of the database and sets the correct retention automatically. The Archived Redo Log schedule is not supported with a snapshot backup.	The administrator must configure two schedules with two retentions: ■ One Full Backup schedule to back up the snapshot (proxy) data part of the database. ■ One Application Backup schedule to back up the stream-based part of the Oracle database. The Archived Redo Log schedule is available with a configured script.
Backups	User-directed backups are not supported. To attempt a user-directed backup (results in a status 240 (no schedules of the correct type exist in this policy)).	User-directed backups are supported.
Load balancing	RAC load balancing is not supported.	RAC load balancing is supported.

Configuring an Oracle Intelligent Policy using Replication Director

Use the following procedure to configure an Oracle snapshot policy that uses Replication Director. This procedure uses the Oracle Intelligent Policy, which makes configuration easier.

Only those policy options that are necessary to configure Replication Director are listed in the following procedure. For descriptions of all policy options, see the [NetBackup Administrator's Guide, Volume I](#) or the online Help.

To create an Oracle Intelligent Policy

- 1 Open the NetBackup web UI.
- 2 On the left, click **Protection > Policies**.
- 3 Click **Add**.
- 4 Select the **Attributes** tab.
- 5 Provide a unique name for the policy.
- 6 The following items are specific to creating an Oracle policy for snapshots with Replication Director:
 - **Policy type**
 For NetBackup to perform Oracle backups, select **Oracle**. An **Oracle** tab appears.
 - **Policy storage**
 Oracle combines snapshots (proxy) and stream-based backups as part of the same backup. The storage that is indicated here is used for the stream-based part of the Replication Director backup.
 Select a storage lifecycle policy that is configured to contain the stream-based (non-snapshot) part of the database backup. The storage must use a storage lifecycle policy that is configured for non-snapshot backups.
 - **Use Replication Director**
 Enable **Use Replication Director** to automatically select other options that Replication Director requires:
 - **Perform snapshot backups:** Ensures that the policy creates snapshots of the disk array.
 - **Retain snapshots for Instant Recovery or SLP management:** Ensures that the policy retains the snapshot after the backup completes.
 - **Snapshot options** button

- Snapshot Type**
- **Auto** (default): The OpenStorage partner uses the best snapshot technology available to that partner to create the snapshot.
 - **Differential**: The OpenStorage partner creates a snapshot that is completely dependent on the source. This parameter is based on copy-on-write technology. The device creates a cache object to maintain the original blocks of the snapshot when the blocks are modified.
 - **Plex**: The OpenStorage Partner creates a snapshot that is completely independent of the source snapshot. This option is based on mirror-break-off technology. When a mirror device is attached to the source, the contents of the mirror device is exactly the same as the source device. When the relationship is broken between the two, the mirror device is separated from the source. The mirror device acts as a point-in-time copy.
 - **Clone**: The OpenStorage Partner creates an independent copy of the volume. The copy process can take some time as the entire copy must be complete. The snapshot that is created is independent of the source.

Maximum Snapshots Sets the maximum number of snapshots to be retained at one time.

The default setting is one. Choose the number of snapshots that is appropriate for your environment. Note that the maximum number of snapshots on a NetApp volume is 255.

When the maximum is reached, snapshot rotation occurs: The next snapshot causes the oldest to be deleted.

Managed by SLP retention is automatically selected if the **Fixed** or the **Expire after Copy** retention is currently selected in the SLP.

7 Select the **Schedules** tab. Create one schedule:

- **Type of backup**: Select **Full backup**. The **Full backup** is used for both the snapshot (proxy) part of the database and the non-snapshot (stream-based) part of the Oracle database.
 The Oracle Intelligent Policy does not support the snapshot of an **Archived redo log backup**. To take a snapshot of the archived redo logs, use the script-based Oracle policy method.

Note: Unless creating Block Level Incremental (BLI) backups, always select **Full Backup** to create snapshots of the Oracle database.

- **Override policy storage selection:** Enable and select the SLP that is configured for snapshot replication. (A snapshot SLP is one in which the first operation is a snapshot operation.) This option must be enabled so that the schedule storage overrides the policy storage with a snapshot SLP.
- **Retention:** The retention for the streamed data is based on the non-snapshot SLP that was indicated as the **Policy storage** in step 6.
 - The non-snapshot SLP specified on the policy storage in step 6 determines the retention for the streamed data.
 - The snapshot SLP that is specified as the schedule storage (**Override policy storage selection**) determines the retention for the snapshot data.

Click **Add** to save the schedule.

- 8 Select the **Instances and databases** tab and specify the databases to back up. The policy must include at least one database. To continue to use the Oracle Intelligent Policy method, select either **Protect non-RAC databases and pluggable databases** or **Protect RAC databases and pluggable databases**.
- 9 Select the **Backup selections** tab. Select the parts of the database to back up. Note that the selection applies to all listed instances.

The following can be selected for the policies that use Replication Director:

- **Whole database:** Backs up the entire database (default).
- **Partial database - Tablespaces:** Backs up the tablespaces.
- **Partial database - Datafiles:** Backs up the data files.
- **Fast Recovery Area (FRA):** Do not select for a policy that uses Replication Director.
- **Database backup shares:** Do not select for a policy that uses Replication Director.
- **Whole Database - Datafile Copy Share:** Do not select for a policy that uses Replication Director.

Note: If you back up the partial database, and later want to perform a Point-in-time rollback restore, make sure that you select all of the tablespaces or data files from a partition in the **Backup selections**.

For copy-back restores, this step is not a requirement.

- 10** Select the **Oracle** tab to configure Oracle RMAN properties.
- 11** When the policy configuration is complete, click **Create**.

Configure a script-based Oracle policy

Use the following procedure to configure an Oracle snapshot policy that uses Replication Director. This procedure uses an Oracle policy type, but does not automatically generate the necessary scripts. It allows the administrator to use custom scripts.

Only those policy options that are necessary to configure Replication Director are listed in the following procedure. For descriptions of all policy options, see the [NetBackup Administrator's Guide, Volume I](#), or the online Help.

To create a script-based Oracle policy

- 1** Open the NetBackup web UI.
- 2** On the left, select **Protection > Policies**.
- 3** Select **Add**.
- 4** Select the **Attributes** tab.
- 5** Provide a unique name for the policy.
- 6** The following items are specific to creating an Oracle policy for snapshots with Replication Director:
 - **Policy type**
For NetBackup to perform Oracle backups, select **Oracle**. An **Oracle** tab appears.
 - **Policy storage**
Oracle combines snapshots (proxy) and stream-based backups as part of the same backup. The storage that is indicated here is used for the stream-based part of the Replication Director backup.
Select the storage that is configured to contain the stream-based (non-snapshot) part of the database backup. The storage can be either a storage lifecycle policy that is configured for non-snapshot backups, or a disk or Media Manager unit.
 - **Use Replication Director**
Enable **Use Replication Director** to automatically select other options that Replication Director requires:
 - **Perform snapshot backups**: Ensures that the policy creates snapshots of the disk array.

- **Retain snapshots for Instant Recovery or SLP management:** Ensures that the policy retains the snapshot after the backup completes.
- **Snapshot options** button

Snapshot Type

- **Auto** (default): The OpenStorage partner uses the best snapshot technology available to that partner to create the snapshot.
- **Differential:** The OpenStorage partner creates a snapshot that is completely dependent on the source. This parameter is based on copy-on-write technology. The device creates a cache object to maintain the original blocks of the snapshot when the blocks are modified.
- **Plex:** The OpenStorage Partner creates a snapshot that is completely independent of the source snapshot. This option is based on mirror-break-off technology. When a mirror device is attached to the source, the contents of the mirror device is exactly the same as the source device. When the relationship is broken between the two, the mirror device is separated from the source. The mirror device acts as a point-in-time copy.
- **Clone:** The OpenStorage Partner creates an independent copy of the volume. The copy process can take some time as the entire copy must be complete. The snapshot that is created is independent of the source.

Maximum Snapshots

Sets the maximum number of snapshots to be retained at one time.

The default setting is one. Choose the number of snapshots that is appropriate for your environment. Note that the maximum number of snapshots on a NetApp volume is 255.

When the maximum is reached, snapshot rotation occurs: The next snapshot causes the oldest to be deleted.

Managed by SLP retention is automatically selected if the **Fixed** or the **Expire after Copy** retention is currently selected in the SLP.

- 7 Select the **Instances and databases** tab. Then select **Clients for use with scripts**.
Select **Yes** to continue.
- 8 Select the **Schedules** tab.
- 9 Edit the **Default-Application-Backup** schedule.

- **Override policy storage selection:** Enable and select a non-snapshot storage unit or a non-snapshot SLP. This option is most likely the storage unit that is specified on the **Attributes** tab. Indicating the option here, makes the selection explicit.
- **Retention:** The policy or SLP indicates the retention for the backup:
 - When the storage is an SLP, the SLP determines the retention and no selection is possible here.
 - When the storage is not an SLP, the schedule determines the retention and a selection is possible here.

10 Select **Add** to save the schedule.

11 Create one **Full Backup** schedule:

- Name the schedule.
- **Type of backup:** Select **Full Backup**.

Note: Unless creating Block Level Incremental (BLI) backups, always select **Full Backup** to create snapshots of the Oracle database.

- **Override policy storage selection:** Enable and select the SLP that is configured for snapshot replication.
- **Retention:** The SLP indicates the retention for the backup.

12 Configure the schedule in the **Start window** tab and the **Exclude dates** tab.

13 Select **Add** to save the schedule.

14 Select the **Clients** tab.

15 Add the client names that contain the Oracle database and indicate the operating system of each.

16 Select the **Backup selections** tab. Specify the script that NetBackup should use. Indicate only one script.

17 When the policy configuration is complete, select **Create**.

NetBackup Dedupe Direct for Oracle

This chapter includes the following topics:

- [About NetBackup Dedupe Direct for Oracle plug-in](#)
- [Verifying the operating system and platform compatibility](#)
- [Before you install NetBackup Dedupe Direct for Oracle plug-in](#)
- [Installing NetBackup Dedupe Direct for Oracle plug-in on the Oracle database server](#)
- [Using the NetBackup Dedupe Direct for Oracle plug-in](#)
- [Uninstalling the NetBackup Dedupe Direct for Oracle plug-in](#)

About NetBackup Dedupe Direct for Oracle plug-in

NetBackup Dedupe Direct for Oracle plug-in is a lightweight plug-in that you can use to store the data from RMAN backups to MSDP storage directly. The Oracle DBA can control the whole protection and life-cycle stages without NetBackup client. Oracle database administrator can deploy this plug-in with root or non-root permissions and perform backup or restore jobs. It also enables client-side deduplication to minimize network traffic and improves overall backup speed.

NetBackup Dedupe Direct for Oracle plug-in includes a set of library of functions that implements backup and restore with RMAN.

Verifying the operating system and platform compatibility

Verify that the NetBackup Dedupe Direct for Oracle plug-in is supported on your operating system or platform.

See the following NetBackup compatibility list for the NetBackup server and client requirements:

veritas.com/content/support/en_US/article.100040093

Before you install NetBackup Dedupe Direct for Oracle plug-in

Before you install and configure the NetBackup Dedupe Direct for Oracle plug-in, ensure that the following requirements are met:

Table 12-1 Requirements

Requirement	Description
Database server	■ Oracle Database version: 11g R2 – 19c ■ Red Hat Enterprise Linux 7/8 See “Verifying the operating system and platform compatibility” on page 213.
Limitations or restrictions	Use Oracle user to install and manage the plug-in.
NetBackup deduplication storage server	Create the MSDP storage server. The storage server deduplicates the backups, writes the data to the storage, and manages the storage. A storage server is an entity that writes to and reads from the storage. One host functions as the storage server, and only one storage server exists for each NetBackup deduplication node. The host must be a NetBackup media server. Although the components of the storage server run on a media server, the storage server is a separate logical entity.

Table 12-1 Requirements (*continued*)

Requirement	Description
Database administrator	Prepare a database administrator username and password at NetBackup server. MSDP: Run the following command to create a username and password on the MSDP server. <pre>/usr/opensv/pdde/pdcr/bin/spauser -a -u <username> -p <password> --role app</pre> WORM storage in Flex Appliance Run the following command to create a username and password on the WORM storage server on Flex Appliance: <pre>setting MSDP-user add-MSDP-user username=<username> password=<password> role=app</pre> - For MVG volumes, create the username and password on the MVG server and all MSDP servers which are members of the MVG volume. The credentials must be the same across all servers. For more information, see the <i>NetBackup Deduplication Guide</i>.
Firewall and ports	For information about firewall and ports, see <i>About MSDP port usage</i> topic of the <i>NetBackup Deduplication Guide</i>.

Installing NetBackup Dedupe Direct for Oracle plug-in on the Oracle database server

Install NetBackup Dedupe Direct for Oracle plug-in on all database servers. You can enable the communication between Oracle database server and the MSDP storage server by configuring the environment variables.

This plug-in can co-exist with the NetBackup Oracle client. You can also install the plug-in along with the NetBackup Oracle client on the Oracle server. Following table lists the backup and recovery capability.

Table 12-2 Backup and recovery capability

Backup	Recovery from the plug-in directly	Recovery from NetBackup Oracle client
Backups from the plug-in	Yes	Yes

Table 12-2 Backup and recovery capability (*continued*)

Backup	Recovery from the plug-in directly	Recovery from NetBackup Oracle client
Backups from NetBackup oracle client	No	Yes

To install the plug-in on the database servers:

- 1 Download NetBackup Dedupe Direct for Oracle plug-in.
- 2 Run the following command to install the plug-in.

`command sbt_config with --install option`
- 3 To ensure the communication between NetBackup Dedupe Direct for Oracle plug-in and the MSDP storage server, set the `MSDP_DBA_USER` and `MSDP_DBA_PASSWORD` environment variables before you perform the backups.

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/oracle/sbt/lib:/home/oracle/sbt/ost-plugins
export MSDP_DBA_USER=<Your username>
export MSDP_DBA_PASSWORD=<Your password>
```

Using the NetBackup Dedupe Direct for Oracle plug-in

Configure the following to use NetBackup Dedupe Direct for Oracle plug-in

- Configure the NetBackup Dedupe Direct for Oracle plug-in.
See [“Configuring the NetBackup Dedupe Direct for Oracle plug-in”](#) on page 216.
- Configure External Certificate Authority (ECA).
See [“Configuring External Certificate Authority \(ECA\)”](#) on page 218.
- Configure Oracle RAC.
See [“Configuring Oracle RAC”](#) on page 220.
- Disable Oracle stream handler.
See [“Disabling Oracle stream handler”](#) on page 220.

Perform backup directly to MSDP storage and restore directly from MSDP storage:

- Perform the RMAN backup.
See [“Performing the RMAN backup directly to MSDP server”](#) on page 221.
- Perform the RMAN restore directly from MSDP.

See [“Performing the RMAN restore directly from the MSDP storage server”](#) on page 222.

Configuring the NetBackup Dedupe Direct for Oracle plug-in

Update the `rman.cfg` configuration file to configure the NetBackup Dedupe Direct for Oracle plug-in.

To configure the NetBackup Dedupe Direct for Oracle plug-in:

- 1
- Use a text editor to open the `rman.cfg` file.
- The `rman.cfg` file is available at `<plugin installation path>/etc/`.
- 2
- Configure the parameters in this file that control the behavior of the NetBackup Dedupe Direct for Oracle plug-in.
- See [“NetBackup Dedupe Direct for Oracle `rman.cfg` file parameters”](#) on page 216.
- 3
- Save and close the file.

NetBackup Dedupe Direct for Oracle `rman.cfg` file parameters

The following table lists the `rman.cfg` file parameters. Update this file to configure NetBackup Dedupe Direct for Oracle plug-in.

Table 12-3 `rman.cfg` file parameters

Settings	Description
STORAGE_SERVER	A valid MSDP server. Default value: The host name of the MSDP storage server that the Oracle DBA provided during the plug-in installation.

Table 12-3 rman.cfg file parameters (*continued*)

Settings	Description
RETENTION_LEVEL	Retention level for current backups. Possible values: ■ 0: 1 week ■ 1: 2 weeks ■ 2: 3 weeks ■ 3: 1 month ■ 4: 2 months ■ 5: 3 months ■ 6: 6 months ■ 7: 9 months ■ 8: 1 year Default value: RETENTION_LEVEL=0 The values are equal to the default values in the primary server. If the values are modified on the primary server, they do not work.
CLIENT	The name of the client for the current backups. If it is not configured, host name of the computer is used as a client name. When you redirect a restore to a different client, the client name is the client that backed up the database.
POLICY	The name of the policy for the current backups. If it is not configured, "agentless_policy" is used as the policy name. You are not required to create an actual policy. Do not use the same policy name with existing policy names.
LSU_NAME	A name of the MSDP LSU. Administrator configures Cloud LSU name from MSDP storage server to backup data to the cloud LSU. Default value: LSU_NAME=PureDiskVolume. The default value is the local LSU name in the target MSDP storage server.

Table 12-3 rman.cfg file parameters (*continued*)

Settings	Description
USER	The MSDP app user. The default value is the value that is provided during the plug-in installation. See the topic <i>About MSDP application user support</i> in the <i>NetBackup Deduplication Guide</i>.
PASSWORD	MSDP app user password. The default value is the value that is provided during the plug-in installation. See the topic <i>About MSDP application user support</i> in the <i>NetBackup Deduplication Guide</i>.
LOG_LEVEL	A log level that specifies the amount of information that is written to the log file. Possible values: ■ 0: Fatal ■ 1: Error ■ 2: Warning ■ 3: Info ■ 4: Debug ■ 5: Trace
SLP_NAME	The import SLP name that you have created on the destination storage server with operation IMPORT. NetBackup creates SLP automatically. However, you must create IMPORT SLP manually on Flex WORM.

Configuring External Certificate Authority (ECA)

If remote MSDP servers support ECA, you can configure the plug-in to communicate with remote MSDP server.

For more information on the external CA support, see the *NetBackup Security and Encryption Guide*.

To configure an ECA

- 1
- Ensure that the external certificate files are saved to the local directory on the Oracle server that this plug-in can access.
- 2
- Use a text editor to open the `msdp_app.cfg` file.

The `msdp_app.cfg` file is available under `<plugin installation path>/etc/` directory:

See [“ECA msdp_app.cfg file parameters”](#) on page 219.

For example,

```
# ECA issue name
subjectname="cert subject name"
# ECA cert path
certpath="cert local path"/cert_chain.pem
# ECA private key path
privatekeypath="cert local path"/private/key.pem
# ECA passphrase file
passphrasefile="cert local path"/passphrasefile
# ECA certificate for trusted certs
truststorepath="cert local path"/trusted/cacerts.pem
```

ECA msdp_app.cfg file parameters

The following table lists the `msdp_app.cfg` file parameters.

Table 12-4 msdp_app.cfg file parameters

Settings	Descriptions
subjectname	A name of the ECA that issued a certificate. Default value: The host name of MSDP storage server. The Oracle DBA provided during the plug-in installation.
certpath	ECA certificate path. For example, certpath="cert local path"/cert_chain.pem"
privatekeypath	The ECA private key path. For example, certpath="cert local path"/private/key.pem"

Table 12-4 msdp_app.cfg file parameters (continued)

Settings	Descriptions
passphrasefile	The ECA passphrase file. For example, certpath="cert local path"/passphrasefile"
truststorepath	The ECA certificate for trusted certificates. For example, certpath="cert local path"/trusted/cacerts.pem"

Configuring Oracle RAC

You must install the plug-in on each Oracle RAC node. After the installation, you can update the configuration file to configure Oracle RAC.

To configure Oracle RAC:

- 1 Open the <installation path>/etc/msdp_app.cfg configuration file for editing.
- 2 Set the same client name on all the Oracle RAC nodes.

For example, **CLIENT=RAC_MY_CLIENT_NAME**
- 3 Set the different policy names on each RAC node.

For example,

POLICY=RAC_node1_MY_POLICY_NAME on Oracle RAC node 1.
POLICY=RAC2_POLICY_NAME on Oracle RAC node 2.

Disabling Oracle stream handler

The Oracle stream handler is enabled by default. You can check the status of stream handler using the **cacontrol** command line utility. Run the following command to check the status of stream handler for each client.

```
/usr/opensv/pdde/pdcr/bin/cacontrol --sth get Oracle  
n132-h95.cdc.veritas.com agentless_policy
```

You can disable the Oracle stream handler by updating the `marker.cfg` configuration file. This file is stored at the plug-in installation directory.

To disable the Oracle stream handler:

- 1 Open the `<installation_path>/etc/msdp_app.cfg` configuration file for editing.
- 2 Change the **ENABLE_STH** value to **0** to disable the Oracle stream handler. Default value is **1**.

Performing the RMAN backup directly to MSDP server

Prerequisites to perform the RMAN backup directly to MSDP server:

- Plug-in is installed correctly.
- Plug-in installation path is added to OS environment variable `LD_LIBRARY_PATH`.
- The MSDP application user and password are created on MSDP server and added as OS environment variable.

NetBackup Dedupe Direct for Oracle plug-in includes a set of library of functions that enable RMAN to work with NetBackup MSDP storage server. On UNIX, NetBackup uses the `RMAN SBT_LIBRARY` parameter to link the RMAN server software with the media management API library that NetBackup Dedupe Direct for Oracle plug-in installs.

Before you perform backup, use the `SBT_LIBRARY` parameter in the `PARMS` section of the `allocate channel` in the RMAN script. In the RMAN script, modify the `ALLOCATE` statement so that the `SBT_LIBRARY` parameter points to the NetBackup Dedupe Direct for Oracle library, and specify the `rman.cfg` location in script. The `rman.cfg` location is in `etc` directory under installed path.

For example,

RUN

```
{

  ALLOCATE CHANNEL c1 DEVICE TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=<plugin_install_path>
  /lib/libmsdp_sbt.so, SBT_PARMS=(SBTCONFIG=<plugin_install_path>/etc/rman.cfg)';

  backup FORMAT 'bk_d%d_u%u_s%s_p%p_t%t' DATABASE filesperset 1;

  backup archivelog all format 'archlog_%d_%T_%s_%p';

  RELEASE CHANNEL c1;

}
```

To perform the RMAN backup directly to MSDP server:

- 1 Set the environment variables LD_LIBRARY_PATH, App user, and password.
- 2 Update the default values.
- 3 Perform RMAN backup directly to MSDP storage server.

Performing the RMAN restore directly from the MSDP storage server

Make sure that a backup using the NetBackup Dedupe Direct for Oracle plug-in has completed successfully before you attempt a restore. You can run the `rman` command from a command prompt on the client. Use the appropriate RMAN command file as a parameter.

Before you perform restore, use the `SBT_LIBRARY` parameter in the `PARMS` section of the allocate channel in the RMAN script. In the RMAN script, modify the `ALLOCATE` statement so that the `SBT_LIBRARY` parameter points to the NetBackup Dedupe Direct for Oracle library, and specify the `rman.cfg` location in script. The `rman.cfg` location is in `etc` directory under installed path.

For example,

```
RUN
{
    ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=<plugin_install_path>
    /lib/libmsdp_sbt.so, SBT_PARMS=(SBTCONFIG=<plugin_install_path>/etc/rman.cfg)';
    RESTORE DATABASE;
    RECOVER DATABASE;
    RELEASE CHANNEL ch00;
}
```

To perform the RMAN backup:

- 1 Set the environment variables SBT_LIBRARY, SBTCONFIG.
- 2 Update the default values.
- 3 Perform RMAN restore directly from the MSDP storage server.

Creating an SLP to import the Oracle backups to NetBackup on Flex WORM

You must create an import SLP to import the images backed up by NetBackup Dedupe Direct for Oracle to the NetBackup for the life cycle management and database recovery.

To create an SLP

- 1 Create IMPORT SLP on the WORM storage server in Flex Appliance.
For example, create **my_slp_name** with IMPORT operation.
- 2 Configure SLP name in `rman.cfg` file as follows:

```
### SLP name  
SLP_NAME=my_slp_name
```

Uninstalling the NetBackup Dedupe Direct for Oracle plug-in

NetBackup Dedupe Direct for Oracle plug-in provides the command to uninstall the plug-in.

Run the following command to uninstall the plug-in:

```
Run uninstall ./bin/sbt_config --uninstall
```

Troubleshooting

This chapter includes the following topics:

- [About troubleshooting NetBackup for Oracle](#)
- [About NetBackup for Oracle troubleshooting steps](#)
- [NetBackup debug logs and reports](#)
- [Enabling the debug logs manually \(Windows\) \(NetBackup for Oracle\)](#)
- [Enabling the debug logs manually \(UNIX\) \(NetBackup for Oracle\)](#)
- [About the NetBackup for Oracle log files](#)
- [Setting the debug level on a Windows client](#)
- [Setting the debug level on a UNIX client](#)
- [About RMAN utility logs](#)
- [Troubleshooting RMAN backup or restore errors](#)
- [Troubleshooting NetBackup for Oracle with Snapshot Client](#)
- [Error: Unable to re-create online log](#)
- [Minimizing time-out failures on large database restores](#)
- [Minimizing the loading and unloading of tapes for database backups](#)
- [Enabling the debug logs manually for NetBackup Dedupe Direct for Oracle plug-in](#)

About troubleshooting NetBackup for Oracle

NetBackup, NetBackup for Oracle, and the Oracle Recovery Manager (RMAN) all provide reports on database backup, archive, and restore operations. These reports are useful for finding the errors that are associated with those applications.

For more information about debug logs and reports, refer to the [NetBackup Administrator's Guide, Volume I](#).

About NetBackup for Oracle troubleshooting steps

When you troubleshoot NetBackup for Oracle problems, the following items are referred to as the API:

- On Windows, `orasbt.dll` is called the API.
- On UNIX, the `libobk` module is called the API. Many media manager vendors also call the `libobk` module DMO (Database Module).

To perform this procedure, ensure that NetBackup is properly installed and configured:

To perform general troubleshooting steps

- 1 When verifying your installation, ensure that the NetBackup for Oracle binaries exist.

On UNIX, these are located in `/usr/opensv/netbackup/bin`.

The binaries are as follows:

On Windows: <code>install_path\NetBackup\bin\bphdb.exe</code> On UNIX: <code>bphdb</code>	The binary resides on the client and both the NetBackup scheduler and the graphical user interface uses the binary to start backups. The main purpose of <code>bphdb</code> is to run an Oracle Intelligent Policy or a shell script that calls <code>rman</code>, <code>bporaexp</code>, or <code>bporaimp</code>.
--	---

On Windows: <code>c:\Windows\System32\orasbt.dll</code> On UNIX: <code>libobk</code>	Provides the functions that RMAN can call. A shared library module that contains the functions that RMAN can call. This library is loaded when RMAN is started. The name of this binary depends on the operating system.
---	--

See [“About linking Oracle RMAN with NetBackup for UNIX”](#) on page 21.

- 2 For the Oracle Intelligent Policy, verify that the following binaries exist.

On Windows: `install_path\NetBackup\bin\bpdbsbora.exe`

On Windows: `install_path\NetBackup\bin\dbdsbrman.dll`

On UNIX: `/usr/opensv/netbackup/bin/bpdbsbora`

On UNIX: `/usr/opensv/lib/libdsbrman.so` (`libdsbrman.sl` on HP-UX)

- 3 Check that both the NetBackup server and the client software work properly. That is, check that normal operating system files can be backed up and restored from the client. The NetBackup client must be running the same version of software as the NetBackup server.
- 4 The logs can become very large, especially `bpdbm`. Ensure that enough free disk space exists in the log directory disk partition.
- 5 Check that the following NetBackup log directories exist:
 - On the client: `bpdbsbora`, `bporaexp` (or `bporaexp64`), `bporaimp` (or `bporaimp64`), `dbclient`, `bphdb`, `bpfis`, and `bpcd`.
 - On the primary server: `bprd` and `bpdbm`.

- On the host with the storage unit: `bpbrm` and `bptm`.

The `VERBOSE` level must be 5.

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“Troubleshooting each stage of the backup or restore”](#) on page 234.

See [“About RMAN utility logs”](#) on page 233.

NetBackup debug logs and reports

The NetBackup server and client software let you enable detailed debugging logs. The information in these log files can help you troubleshoot the problems that occur outside of either the database agent or Oracle.

Note the following with regard to these logs:

- These logs do not reveal the errors that occur when Oracle is running unless those errors also affect NetBackup. Oracle may (or may not) write errors in the application to the NetBackup logs. Your best sources for Oracle error information are the logs provided by Oracle.
- Generally, each debug log corresponds to a NetBackup process and executable. However, for an RMAN backup, the debug log is created in the `dbclient` directory, which has no corresponding executable.

More detailed information about the debug log files is available.

See the [NetBackup Troubleshooting Guide](#).

Also refer to the following file:

Windows:

`install_path\NetBackup\logs\README.debug file`

UNIX:

`/usr/opensv/netbackup/logs/README.debug file`

NetBackup provides other reports that are useful in isolating problems. One such report is All Logs Entries on the server. Information on server reports is available.

See the [NetBackup Administrator's Guide, Volume I](#).

Enabling the debug logs manually (Windows) (NetBackup for Oracle)

To create the NetBackup for Oracle for Windows database agent logs manually

1 Create the following directories on the client:

- bphdb

For any backup that is initiated from an automated schedule on the primary server.

```
install_path\NetBackup\logs\bphdb
```

- bpdbsbora

For any OIP backup restore operations.

```
install_path\NetBackup\logs\bpdbsbora
```

- dbclient

For any backup or restore using RMAN.

```
install_path\NetBackup\logs\dbclient
```

- bpbkar

For any snapshot backup.

```
install_path\NetBackup\logs\bpbkar
```

- tar

For any snapshot restore.

```
install_path\NetBackup\logs\tar
```

2 Verify the user or group that the Oracle process (process that loads `orasbt.dll`) has appropriate permissions to access the following directories if they exist. If the following directories do not exist, the directories are created automatically with the correct permissions.

```
install_path\NetBackup\logs\user_ops
```

```
install_path\NetBackup\logs\user_ops\dbext
```

```
install_path\NetBackup\logs\user_ops\dbext\logs
```

Also verify that the user or group that the Oracle process runs as has appropriate permissions to access the log directories in step 1.

- 3 On the NetBackup server or servers, create the debug log directories for the legacy processes that interact with the Oracle agent.

On the primary server:

```
install_path\NetBackup\logs\bprd
```

On the media server or servers:

```
install_path\NetBackup\logs\bpbrm
```

```
install_path\NetBackup\logs\bptm
```

- 4 The debug logs for unified processes on the server and the client hosts are created automatically by NetBackup.

NetBackup writes unified logs to `install_path\NetBackup\logs`.

For information on how to use logs and reports, see the [NetBackup Troubleshooting Guide](#).

Enabling the debug logs manually (UNIX) (NetBackup for Oracle)

To create the NetBackup for Oracle for UNIX database agent logs manually

- 1 Create the following directories on the client:

- **bphdb**

For any backup that is initiated from an automated schedule on the primary server.

```
/usr/opensv/netbackup/logs/bphdb
```

- **bpdbsbora**

For any OIP backup restore operations.

```
/usr/opensv/netbackup/logs/bpdbsbora
```

- **dbclient**

For any backup or restore using RMAN.

```
/usr/opensv/netbackup/logs/dbclient
```

- **bpbkar**

For any snapshot backup.

```
/usr/opensv/netbackup/logs/bpbkar
```

- nbtar

For any snapshot restore.

`/usr/opensv/netbackup/logs/tar`

- 2 Verify the user or group that the Oracle process (process that loads `libobk`) has appropriate permissions to access the following directories if they exist. If the following directories do not exist, the directories are created automatically with the correct permissions.

`/usr/opensv/logs/user_ops`

`/usr/opensv/logs/user_ops/dbext`

`/usr/opensv/logs/user_ops/dbext/logs`

Also verify that the user or group that the Oracle process runs as has appropriate permissions to access the log directories in step 1.

- 3 On the NetBackup server or servers, create the debug log directories for the legacy processes that interact with the Oracle agent.

On the primary server:

`/usr/opensv/logs/bprd`

On the media server or servers:

`/usr/opensv/logs/bpbrm`

`/usr/opensv/logs/bptm`

- 4 The debug logs for unified processes on the server and the client hosts are created automatically by NetBackup.

NetBackup writes unified logs to `/usr/opensv/logs`.

For information on how to use logs and reports, see the [NetBackup Troubleshooting Guide](#).

About the NetBackup for Oracle log files

[Table 13-1](#) describes the logs that are created when you create the log directories. Use a text editor to view the contents of the logs.

The logs are located in the following directories:

Windows: `install_path\NetBackup\logs\<cmd>`

UNIX: `/usr/opensv/netbackup/logs/<cmd>`

For example, the logs for `bphdb` all appear in the `install_path\NetBackup\logs\bphdb` directory (Windows) or the `/usr/opensv/netbackup/logs/bphdb` directory (UNIX).

Table 13-1 Log files

Log directory	Description
bphdb	The <code>bphdb</code> directory contains the following types of logs: ■ Windows: <code>obk_stdout.mmddyy.hhmmss.txt</code> UNIX: <code>obk_stdout.mmddyy</code> Unless it is redirected elsewhere, NetBackup writes the shell script output to this file. ■ Windows: <code>obk_stderr.mmddyy.hhmmss.txt</code> UNIX: <code>obk_stderr.mmddyy</code> Unless it is redirected elsewhere, NetBackup writes any shell script errors to this file. ■ Windows: <code>mmddyy.log</code> UNIX: <code>log.mmddyy</code> This log contains debugging information for the <code>bphdb</code> process. <code>bphdb</code> is the NetBackup database backup binary. It is invoked when an automatic backup schedule is run. NetBackup for Oracle uses this client process for the shell script execution.
dbclient	The <code>dbclient</code> directory contains the following execution log: ■ Windows: <code>mmddyy.log</code> ■ UNIX: <code>log.mmddyy</code> This log contains debugging information and execution status for the Oracle for NetBackup client processes. On Windows, the processes are linked to the library program that is provided with NetBackup for Oracle. On UNIX, this library program is <code>libobk</code>.
bpdbsbora	The <code>bpdbsbora</code> directory contains the following execution log: ■ Windows: <code>mmddyy.log</code> ■ UNIX: <code>log.mmddyy</code>
bporaexp64	The <code>bporaexp</code> (or <code>bporaexp64</code> on UNIX) directory contains the following execution log: ■ Windows: <code>mmddyy.log</code> ■ UNIX: <code>log.mmddyy.log</code>

Table 13-1 Log files (*continued*)

Log directory	Description
bporaimp64	The bporaimp (or bporaimp64 on UNIX) directory contains the following execution log: - Windows: <i>mmdyy.log</i> - UNIX: <i>log.mmdyy</i>

Setting the debug level on a Windows client

The debug logs are located in *install_path*\NetBackup\logs.

To change the amount of debug information in other log directories, set the other debug levels. For instance, Verbose.

To set the debug level on a Windows client

- 1 Open the **Backup, Archive, and Restore** interface.
- 2 Select **File > NetBackup Client Properties**.
- 3 Click the **Troubleshooting** tab.
- 4 Set the **General** debug level.
- 5 Set the **Verbose** debug level.
- 6 Set the **Database** debug level.
- 7 Click **OK** to save your changes.
- 8 Stop and start the Oracle database services. This action is necessary for *orasbt.dll* to pick up the new debug level.

Setting the debug level on a UNIX client

The debug logs are located in */usr/opensv/netbackup/logs*.

To set the debug level

- ◆ Enter the following line in the *bp.conf* file.

```
VERBOSE = X
```

Where *X* is the debug level you want.

About RMAN utility logs

RMAN uses a command language interpreter, and it can be run in interactive or batch mode. You can use the following syntax to specify a log file on the command line to record significant RMAN actions:

```
msglog 'logfile_name'
```

Troubleshooting RMAN backup or restore errors

An RMAN backup error can originate from NetBackup or from Oracle, as follows:

- On the NetBackup side, an error can be from the API, from the NetBackup server or client, or from Media Manager.
- On the Oracle side, an error can be from RMAN or from the target database.

Cohesity suggests that you use the following steps when troubleshooting a failed operation:

- Check the logs to determine the source of the error.
- Troubleshoot each stage of the backup or restore.

Verifying the RMAN script on UNIX

The following procedure describes how to verify that the RMAN script works correctly.

To verify the RMAN script

- 1 Use RMAN to make a backup directly to disk. Do not use NetBackup.
- 2 Use RMAN with NetBackup to create a backup.
- 3 Look for log files under the `/usr/opensv/netbackup/logs/dbclient` directory structure.

On UNIX all Oracle users are non-root so the logs files are in the `/usr/opensv/netbackup/logs/dbclient/<users_dir>` directory.

If no log file exists, `libobk` is not linked into Oracle properly.

See [“Testing configuration settings for NetBackup for Oracle”](#) on page 98.

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Troubleshooting each stage of the backup or restore

The following explains the sequence of events for an action initiated by RMAN and suggests solutions for the problems that can occur at each point in the sequence:

- **rman starts.**

A backup or restore can be started in any of the following ways:

- From an RMAN backup or restore initiated from the operating system prompt such as:

```
rman target user/pwd[@TNS_alias] \
rcvcat user/pwd[@TNS_alias]\
cmdfile RMAN_script_file_name
```

Where the *RMAN_script_file_name* is fully qualified.

- Manually from the administrator interface on the primary server.
- Automatically by an automatic backup schedule.

If an error occurs now, check the RMAN log.

- RMAN verifies its environment and then issues requests to the API.

On Windows, some information, such as the NetBackup version, API versions, and trace file name, is registered with RMAN. An error now is usually due to a problem with client and server communication. Check the messages in the `bprd` and the `bpcd` logs for clues.

On UNIX, some information, such as the NetBackup version, API versions, trace file name, and NetBackup signal handlers, is registered with RMAN. An error now is usually due to a problem with client and server communication. Check the messages in the `bprd` and the `bpcd` logs for clues. Also verify the `bp.conf` entries on the client.

- RMAN issues a backup or restore request.

The API gathers necessary parameters and sends the `backup or restore` request to the NetBackup server. The API waits until both the server and client are ready to transfer data before it returns to the request.

The API then sends this information to the primary server's `bprd` process.

To troubleshoot a problem in this part of the first sequence, examine the following file:

Windows:

```
install_path\NetBackup\logs\dbclient\mmdyy.log
```

UNIX:

```
/usr/opensv/netbackup/logs/dbclient/log.mmdyy
```

If the `bprd` process failed, check the logs for `bprd` and `bpbrm`.

A failure now is frequently due to bad NetBackup server or Oracle policy configuration parameters.

NetBackup can usually select the correct Oracle policy and schedules. But NetBackup can select a policy or schedule in error if there are several Oracle policies in its database.

On Windows, try setting the `SERVER` and `NB_ORA_POLICY` values in the client environment.

On UNIX, try setting the `SERVER` and `POLICY` values in the `bp.conf` file on the client or by setting environment variables.

For example, the following C Shell `setenv` commands specify the Oracle policy, schedule, and server for NetBackup to use:

```
setenv NB_ORA_POLICY policyname
setenv NB_ORA_SCHED application_backup_schedule_name
setenv NB_ORA_SERV Netbackup_server
```

- RMAN issues read or write requests to the API, which then transfers data to or from the NetBackup server.

A failure here is probably due to NetBackup media, network, or time-out errors.

- RMAN tells the API to close the session.

The API waits for the server to complete its necessary actions (for example, it verifies the backup image) and then exits.

An error can originate from either NetBackup or RMAN, as follows:

- RMAN aborts if it encounters an error while it reads a data file during the backup (for example, if Oracle blocks are out of sequence). It also aborts if NetBackup sends a bad backup image during the restore.
- NetBackup might return an error code to the API if for some reason it could not complete the backup successfully.

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“Minimizing time-out failures on large database restores”](#) on page 237.

See [“Troubleshooting each stage of the XML export or XML import”](#) on page 315.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Troubleshooting NetBackup for Oracle with Snapshot Client

Debug logs used for troubleshooting the problems that occur with NetBackup and NetBackup for Oracle have been discussed in previous areas. In addition to those logs, there are debug logs used for troubleshooting NetBackup for Oracle with Snapshot Client.

Snapshot Client backup and debug messages are written to the following subdirectories of:

Windows:

`install_path\NetBackup\logs`

UNIX:

`/usr/opensv/netbackup/logs/`

The logs are as follows:

- The `bpbm` log is on the NetBackup media server.
- The `bptm/bpdm` log is on the NetBackup media server.
- The `bpbkar` log is on the NetBackup client and alternate client.
- The `bpfis` log is on the NetBackup client and alternate client.
- The `bpffi` log is on the NetBackup client or alternate client.

Snapshot Client restore and debug messages are written to the following subdirectories on the NetBackup primary server:

- The `bprestore` is almost always a client log on the NetBackup host that initiated the restore by using the `bprestore` command.
- The `bprd` is on the NetBackup primary server.
- The `bpbm` is on the NetBackup primary server.
- The `bptm/bpdm` is on the NetBackup media server. Both the tape and the disk backup log to `bptm`, disk backups also log to `bpdm`.
- The `tar` is on the NetBackup client or redirected client.

Additional help for troubleshooting most installation and other issues is available in the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Error: Unable to re-create online log

During certain clone operations you may encounter the following error:

```
ORA-00344: unable to re-create online log
'E:\APP\ORACLE\PRODUCT\19.3\DB_HOME1\RAC_RESTORE2\DATA\CDB19R\
ONLINELOG\GROUP_1.264.1126884003'
```

This error may occur in the following situations:

- The backup source is ASM and you are cloning to file system using OMF.
- The backup source is a file system and you are cloning to ASM using OMF.

Cause

NetBackup may not create all of the directories before the clone. This behavior is because the Oracle OMF storage controls the creation of the directories. The exceptions are online redo log location and temp file locations, which you must create manually before the clone operation.

Solution

Before you run the clone operation, create the following directory on the target system. For example:

```
E:\APP\ORACLE\PRODUCT\19.3\DB_HOME1\RAC_RESTORE2\DATA\CDB19R\
ONLINELOG
```

Minimizing time-out failures on large database restores

Large database restores sometimes fail when multiple restore sessions compete for resources. In this situation, a restore session can be delayed while NetBackup waits for media or device access. If the delay is too long, the restore session times out. Use the following procedure to minimize session time-outs and to allow the restore jobs to complete successfully.

To minimize time out failures on large database restores

- 1 Open the NetBackup web UI.
- 2 On the left, select **Hosts > Host properties**.
- 3 Select the check box for the client.
- 4 If necessary, select **Connect**.
- 5 Select **Edit client**.

6 Select **Timeouts**.

7 Set the **Client read timeout** property to a large value.

The default for the **Client read timeout** setting is 300 seconds (5 minutes). For database agent clients, increase the value significantly from the recommended value.

See the [NetBackup Web UI Administrator's Guide](#).

For example, change this setting to 30-60 minutes to minimize time-out errors.

8 Select **Save**.

Note: This change may delay detecting problems during subsequent backups. Consider putting the original value back in place once any restore that requires a change is complete.

Minimizing the loading and unloading of tapes for database backups

You can minimize excessive unloading and reloading of tapes between multistreamed database backups by changing the media settings for the primary or the media server.

See the [NetBackup Administration Guide, Volume 1](#) for details.

To minimize loading and unloading of tapes

1 Open the NetBackup web UI.

2 On the left, click **Hosts > Host properties**.

3 Select the primary or the media server.

4 If necessary, click **Connect**.

5 Click **Edit primary server** or **Edit media server**.

6 Click **Media**.

7 Configure the following settings:

- **Media unmount delay**

- **Media request delay**

Use this variable only with non-robotic drives, such as tape stackers.

Enabling the debug logs manually for NetBackup Dedupe Direct for Oracle plug-in

The configuration files are stored at the installation directory, under the folders `etc` and `ost-plugins`.

You can configure debugging log settings in two configuration files: `/etc/rman.cfg` and `/ost-plugins/pd.conf`

Table 13-2 rman.cfg file settings:

Settings	Description
LOG_LEVEL	Following are the log levels: ■ 0 Fatal ■ 1 Error ■ 2 Warning ■ 3 Info ■ 4 Debug ■ 5 Trace The default log level is 3.
LOG_MAXSIZE	Maximum size per log file. The default value is 100 MB.
LOG_RETENTION	Maximum days log files to keep. The default value is 90 days.

Table 13-3 pd.conf file settings

Settings	Description
LOGLEVEL	Log levels are from 0 to 10. The default log level is 2.
MAX_LOG_MBSIZE	Maximum size per log file. The default value is 100 MB.

Deduplication best practices

This appendix includes the following topics:

- [Optimizing and deduplicating stream-based and proxy copy Oracle backups](#)
- [Oracle stream handler](#)
- [Configuring a stream-based Oracle backup](#)
- [Example RMAN script for a stream-based backup](#)
- [Editing the RMAN script and configuring NetBackup for Oracle for a proxy copy backup](#)
- [Example RMAN script for a proxy copy backup](#)

Optimizing and deduplicating stream-based and proxy copy Oracle backups

NetBackup enables you to perform optimized deduplication of Oracle databases. You can perform either a stream-based backup or a proxy copy backup.

Cohesity recommends that you perform a proxy copy if the database consists of many small tablespaces. A proxy copy is also recommended if the DBA or the backup administrator does not want to set `FILESERSET=1`.

To configure a proxy copy Oracle backup, you need to edit the RMAN script and configure NetBackup for Oracle.

See [“Editing the RMAN script and configuring NetBackup for Oracle for a proxy copy backup”](#) on page 248.

For stream-based backups, Cohesity recommends that you specify `FILESERSET=1` for all Oracle database backups. When `FILESERSET=1` is specified, Oracle generates the backup set identically each time. The backup set is generated with the same data from the same files in the same sequence each time the database is backed up. This uniformity ensures better deduplication. In addition, when `FILESERSET=1` is in effect, Oracle does not perform multiplexing, so Oracle includes only one file in each backup set. If `FILESERSET` is specified with a number other than 1, Oracle groups files together unpredictably and deduplication rates suffer. You may also want to increase the number of channels that are allocated to the backup, if possible.

It is recommended that you test your database backups by running both stream-based backups and proxy copy backups. Measure the deduplication rates and the backup times, and see which method fits best in your environment. The Oracle database files themselves benefit the most from deduplication. Typically, the archive logs and the control files are unique, so they benefit less from deduplication.

Deduplication performs best when used in the following ways:

Stream deduplication

The Oracle Intelligent Policy detects both ASM and non-ASM environments to generate the correct backup scripts ensuring good deduplication rates. In a non-ASM environment, scripts are generated that are a non-snapshot proxy backup. In an ASM environment, scripts are generated that set `FILESERSET=1` if this command has not been modified in a backup policy.

You may need to create a custom script for your environment. However, in most situations, the Oracle Intelligent Policy creates the script that is needed for your environment.

Snapshot deduplication

When you use snapshot deduplication, nothing changes and NetBackup proxy snapshot backup is performed. If ASM is detected, an error is displayed. Snapshot backup is not allowed in an ASM environment.

Note: In OIP when deduplication storage is used and a stream-based backup is selected, the policy overrides and attempts to perform a proxy backup. The override is attempted if NO ASM storage is found in the database. The

ORACLE_OVERRIDE_DATA_MOVEMENT setting in the `bp.conf` file can be used to override this behavior.

Set `ORACLE_OVERRIDE_DATA_MOVEMENT=1` to always do streaming.

Set `ORACLE_OVERRIDE_DATA_MOVEMENT=2` to always do proxy.

Set `ORACLE_OVERRIDE_DATA_MOVEMENT=>2` to maintain standard behavior.

On UNIX you can edit the `/usr/opensv/netbackup/bp.conf` file.

On Windows you can use the `bpsetconfig` command (`install_path\NetBackup\bin\admincmd\bpsetconfig`) on the server to set the client's configuration. See the following example:

```
bpsetconfig -h myoracleclient  
  
ORACLE_OVERRIDE_DATA_MOVEMENT = 1
```

For information about the backup methods, see the following:

- See [“Configuring a stream-based Oracle backup”](#) on page 245.
 - See [“Editing the RMAN script and configuring NetBackup for Oracle for a proxy copy backup”](#) on page 248.
- See [“About Oracle Intelligent Policies \(OIP\)”](#) on page 59.
- See [“Proxy backup examples”](#) on page 184.

Oracle stream handler

The Oracle stream handler is not enabled by default for existing and new Oracle clients in NetBackup 8.3. Also, the Oracle stream handler only supports stream-based backups and you can enable and disable the Oracle stream handler per `<client>` `<policy>` combination using the `cacontrol` command line utility.

In NetBackup 10.0, the Oracle stream handler is enabled (by default) for all new clients that have no existing images. As with previous versions, the Oracle stream handler only supports stream-based backups and you can configure the Oracle stream handler using the `cacontrol` command line utility. You can enable and disable the stream handler per the following:

- Policy and client
- Policy level

- Stream type level

Note: When you use the Oracle stream handler, it is not recommended to use variable-length deduplication.

The `cacontrol` command utility with the `--sth` flag, is used to override the default behavior of NetBackup by creating a `Marker Entry` for a client, policy, or stream type in a configuration file. The `cacontrol` command utility is located in the following locations:

- Windows: `install_path\Veritas\pdde\cacontrol`
- UNIX: `/usr/opensv/pdde/pdcr/bin/cacontrol`

In the following examples for `cacontrol`, `STHTYPE` must be set to `Oracle` to configure the Oracle stream handler.

In NetBackup 8.3, you can configure `cacontrol` using the following options:

- You can query the settings for the stream handler per client and policy.

```
cacontrol --sth get <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can enable the stream handler per client and policy.

```
cacontrol --sth update  
<STHTYPE> <CLIENT> <POLICY> [SPAUSER] <enabled>
```

- You can delete the settings for client and policy (return to default behavior).

```
cacontrol --sth delete <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can disable the stream handler on a client and policy.

```
cacontrol --sth update  
<STHTYPE> <CLIENT> <POLICY> [SPAUSER] <disabled>
```

Note: When you use `cacontrol` to set `<POLICY>` or `<STHTYPE>` to `enabled`, NetBackup enables all the old clients which have existing images. The deduplication rate decreases significantly only at the first backup after enabled. Also, the storage usage increases only in the first backup after enabled. Basically, NetBackup behaves as if you have run a first full backup. Both the deduplication rate and storage usage improve after initial activation of the stream handler.

When using the `cacontrol` command utility to create a Marker Entry in NetBackup 10.0, priority is given to the more granular configuration. For example:

```
Marker Entry 1: <Client1> <Policy1> to enabled
```

```
Marker Entry 2: <Policy1> to disabled
```

The stream handler is enabled because the more granular configuration in Marker Entry 1 has higher priority.

In NetBackup 10.0, you can configure `cacontrol` using the following options:

- You can query the settings for the stream handler per client and policy.

```
cacontrol --sth get <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can enable the stream handler per client and policy.

```
cacontrol --sth update  
<STHTYPE> <CLIENT> <POLICY> [SPAUSER] <enabled>
```

- You can delete the settings for a client and policy (return to default behavior).

```
cacontrol --sth delete <STHTYPE> <CLIENT> <POLICY> [SPAUSER]
```

- You can disable the stream handler on a client and policy.

```
cacontrol --sth update  
<STHTYPE> <CLIENT> <POLICY> [SPAUSER] <disabled>
```

- You can query the settings for the stream handler per policy.

```
cacontrol --sth getbypolicy <STHTYPE> <POLICY> [SPAUSER]
```

- You can enable the stream handler per policy.

```
cacontrol --sth updatebypolicy  
<STHTYPE> <POLICY> [SPAUSER] <enabled>
```

- You can delete the settings for the stream handler per policy (return to default behavior).

```
cacontrol --sth deletebypolicy <STHTYPE> <POLICY> [SPAUSER]
```

- You can disable the stream handler per policy.

```
cacontrol --sth updatebypolicy
<STHTYPE> <POLICY> [SPAUSER] <disabled>
```

- You can query the settings for the stream handler per stream handler type.

```
cacontrol --sth getbytype <STHTYPE> [SPAUSER]
```

- You can enable a stream handler per stream handler type.

```
cacontrol --sth updatebytype <STHTYPE> [SPAUSER] <enabled>
```

- You can delete the settings for a stream handler (return to default behavior).

```
cacontrol --sth deletebytype <STHTYPE> [SPAUSER]
```

- You can disable the stream handler per stream handler type.

```
cacontrol --sth updatebytype <STHTYPE> [SPAUSER] <disabled>
```

Configuring a stream-based Oracle backup

The following procedure explains how to reconfigure an existing Oracle RMAN specification to achieve a stream-based, optimized, deduplicated Oracle backup.

To configure a stream-based Oracle backup

- 1 On the client computer that hosts the Oracle database, open the RMAN backup script in a text editor, and make the following edits:

- Add the `FILESERSET=1` parameter to the part of the RMAN script that specifies how you want to back up the database.

Do not add `FILESERSET=1` to the section of the RMAN script that specifies how to back up the control files or archive logs. Typically, these other data objects are unique to each backup, so there is very little potential gain from optimizing the control file and archive log backups for deduplication.

Example RMAN script after editing:

```
BACKUP
FILESERSET=1
%BACKUP_TYPE%
FORMAT 'bk_u%u_s%s_p%p_t%t'
DATABASE;
```

The addition of `FILESERSET=1` suppresses Oracle multiplexing of more than one data file into a backup set. When you suppress Oracle multiplexing,

Oracle creates the backup set identically each time the backup runs. NetBackup can deduplicate these identical backup sets.

- Specify parallel backup streams for the database backup.
Specify appropriate `ALLOCATE CHANNEL` and `RELEASE CHANNEL` parameters in the backup script.

For an example that shows an edited backup script, see the following:

See [“Example RMAN script for a stream-based backup”](#) on page 247.

- 2 Disable Oracle's optimization and encryption for the duration of the database backup.

By default, Oracle disables optimization and encryption. If optimization and encryption are enabled, run the following RMAN commands from the command line to disable optimization and encryption:

```
RMAN> CONFIGURE BACKUP OPTIMIZATION OFF;  
RMAN> CONFIGURE ENCRYPTION FOR DATABASE OFF;
```

If your site requires encryption, you can specify encryption in the NetBackup for Oracle backup policy.

- 3 Disable Oracle's compression for the duration of the database backup.

By default, Oracle disables compression. If compression is enabled, Oracle compresses unused sections in streams, and the result is unpredictable deduplication rates. If compression is enabled, run the following RMAN command from the command line to disable compression:

```
RMAN> CONFIGURE DEVICE TYPE SBT_TAPE BACKUP TYPE TO BACKUPSET;
```

If your site requires compression, you can specify compression in the NetBackup for Oracle backup policy.

- 4 Configure a NetBackup for Oracle policy.

If you want to compress or encrypt the backup, enable compression and encryption in the NetBackup `pd.conf` file.

Run a full database backup as soon as you can. The policy can perform incremental backups until the full backup can be run.

Note: Make sure that Oracle optimization, encryption, and compression are disabled for the entirety of the database backup. Make sure to check specifications outside of the RMAN backup script, too.

Example RMAN script for a stream-based backup

The following is an example fragment from an RMAN script that performs an optimized, deduplicated, stream-based backup of an Oracle database.

```
RUN {

# Back up the database.
# Use 4 channels as 4 parallel backup streams.

ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL ch02 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL ch03 TYPE 'SBT_TAPE';
SEND ' NB_ORA_SERV=$NB_ORA_SERV';

BACKUP
    $BACKUP_TYPE
    SKIP INACCESSIBLE
    TAG hot_db_bk_level0
    # The following line sets FILESPERSET to 1 and facilitates database deduplication.
    FILESPERSET 1
    FORMAT 'bk_%s_%p_%t'
    DATABASE;
    sql 'alter system archive log current';
RELEASE CHANNEL ch00;
RELEASE CHANNEL ch01;
RELEASE CHANNEL ch02;
RELEASE CHANNEL ch03;

# Back up the archive logs
# The FILESPERSET parameter setting depends on the number of archive logs you have.

ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE';
SEND ' NB_ORA_SERV=$NB_ORA_SERV';
BACKUP
    FILESPERSET 20
    FORMAT 'al_%s_%p_%t'
    ARCHIVELOG ALL DELETE INPUT;
RELEASE CHANNEL ch00;
```

```
RELEASE CHANNEL ch01;
#
# Note: During the process of backing up the database, RMAN also backs up the
# control file. This version of the control file does not contain the
# information about the current backup because "nocatalog" has been specified.
# To include the information about the current backup, the control file should
# be backed up as the last step of the RMAN section. This step would not be
# necessary if we were using a recovery catalog or auto control file backups.
#
ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
SEND ' NB_ORA_SERV=$NB_ORA_SERV';
BACKUP
    FORMAT 'cntrl_%s_%p_%t'
    CURRENT CONTROLFILE;
RELEASE CHANNEL ch00;
}
```

Editing the RMAN script and configuring NetBackup for Oracle for a proxy copy backup

The following procedure explains how to edit the RMAN script on the client.

To edit the RMAN script

- 1 On the client computer that hosts the Oracle database, open the RMAN backup script in a text editor, and make the following edits:

- Add `PROXY` to the list of commands that backs up the data files.

Example RMAN script after editing:

```
BACKUP
FORMAT 'bk_u%u_s%s_p%p_t%t'
PROXY
DATABASE;
```

- Specify the `NB_ORA_PC_STREAMS` parameter in the database backup script. The `NB_ORA_PC_STREAMS` variable controls the number of proxy copy backup streams to be started. By default, the agent initiates one backup job for all files. If the RMAN `send` command passes `NB_ORA_PC_STREAMS`, NetBackup for Oracle splits the files into the number of groups that are specified by the variable based on the file size. The agent attempts to create streams of equal size and determines the number of processes that run to perform the backup.

For an example that shows an edited backup script, see the following:

See [“Example RMAN script for a proxy copy backup”](#) on page 249.

- 2 Disable Oracle's optimization and encryption for the duration of the database backup.

By default, Oracle disables optimization and encryption. If the optimization and encryption are enabled, run the following RMAN commands from the command line to disable optimization and encryption:

```
RMAN> CONFIGURE BACKUP OPTIMIZATION OFF;  
RMAN> CONFIGURE ENCRYPTION FOR DATABASE OFF;
```

If your site requires encryption, you can specify encryption in the NetBackup for Oracle backup policy.

- 3 Disable Oracle's compression for the duration of the database backup.

By default, Oracle disables compression. If compression is enabled, Oracle compresses unused sections in streams, and the result is unpredictable deduplication rates. If compression is enabled, run the following RMAN command from the command line to disable compression:

```
RMAN> CONFIGURE DEVICE TYPE SBT_TAPE BACKUP TYPE TO BACKUPSET;
```

If your site requires compression, you can specify compression in the NetBackup for Oracle backup policy.

- 4 Configure a NetBackup for Oracle policy.

If you want to compress or encrypt the backup, enable compression and encryption in the NetBackup `pd.conf` file.

Run a full database backup as soon as you can. You can perform incremental backups until the full backup can be run.

Note: Make sure that Oracle optimization, encryption, and compression are disabled for the entirety of the database backup. Also, make sure to check specifications outside of the RMAN backup script.

Example RMAN script for a proxy copy backup

The following is an example of an RMAN script that performs an optimized, deduplicated, proxy copy backup of an Oracle database.

```
RUN {

# Back up the database.

ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';

# Specify 2 streams.

SEND 'NB_ORA_PC_STREAMS=2';
BACKUP
    PROXY
    SKIP INACCESSIBLE
    TAG hot_db_bk_proxy
    FORMAT 'bk_%s_%p_%t'
    DATABASE;
    sql 'alter system archive log current';
RELEASE CHANNEL ch00;

# Back up the archive logs.
# The FILESPERSET parameter setting depends on the number of archive logs you have.

ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE';
SEND ' NB_ORA_SERV=$NB_ORA_SERV';
BACKUP
    FILESPERSET 20
    FORMAT 'al_%s_%p_%t'
    ARCHIVELOG ALL DELETE INPUT;
RELEASE CHANNEL ch00;
RELEASE CHANNEL ch01;

#
# Note: During the process of backing up the database, RMAN also backs up the
# control file. This version of the control file does not contain the
# information about the current backup because "nocatalog" has been specified.
# To include the information about the current backup, the control file should
# be backed up as the last step of the RMAN section. This step would not be
# necessary if we were using a recovery catalog or auto control file backups.
#
ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
SEND ' NB_ORA_SERV=$NB_ORA_SERV';
```

```
BACKUP
    FORMAT 'cntrl_%s_%p_%t'
    CURRENT CONTROLFILE;
RELEASE CHANNEL ch00;
}
```

Snapshot Client support of SFRAC

This appendix includes the following topics:

- [About Snapshot Client support of SFRAC](#)
- [NetBackup configuration for an SFRAC environment](#)
- [Configuring the SFRAC environment for a backup operation](#)
- [Performing a rollback restore in an SFRAC environment](#)
- [Troubleshooting NetBackup in an SFRAC environment](#)

About Snapshot Client support of SFRAC

Cohesity Storage Foundation for the Oracle Real Application Clusters (RAC) environment leverages storage management and high availability technologies for deployment of Oracle RAC on UNIX environments.

Storage Foundation is a complete solution for heterogeneous online storage management. Based on VxVM and VxFS, it provides a standard set of integrated tools to centrally manage data growth, maximize storage hardware usage, and provide data protection.

See [“NetBackup configuration for an SFRAC environment”](#) on page 253.

See [“Configuring the SFRAC environment for a backup operation”](#) on page 253.

See [“Performing a rollback restore in an SFRAC environment”](#) on page 255.

See [“Troubleshooting NetBackup in an SFRAC environment”](#) on page 256.

NetBackup configuration for an SFRAC environment

To perform the offhost snapshot backup of Oracle database in the SFRAC environment, the NetBackup client software must be installed on each node of the cluster.

You need to configure the following:

- On the primary server or media server that resides outside of the cluster, you must configure the policy to back up the Oracle RAC database.
- Configure the alternate client so that the snapshot is taken using that offhost. The alternate client should not be part of the cluster.

Note: IPv6 is not supported for SFRAC.

See [“Configuring the SFRAC environment for a backup operation”](#) on page 253.

See [“About Snapshot Client support of SFRAC”](#) on page 252.

Configuring the SFRAC environment for a backup operation

The backup and rollback operations involve the Oracle Agent and the hardware array. The configuration steps required for both of these operations should also be done before taking the snapshot.

The following lists the prerequisites before you perform a backup in the SFRAC environment.

To configure the SFRAC environment for a backup operation

- 1 Configure a virtual IP or virtual name over the cluster. NetBackup refers to the client by using this virtual name.
- 2 The NetBackup client name on each node of the cluster must match the virtual name that is configured on the cluster. Do one of the following:
 - If you have already installed the client, change the CLIENT_NAME entry in the `bp.conf` file of the NetBackup directory to the following:

```
CLIENT_NAME = <virtual_name>
```

- Alternatively, add the following parameter to the RMAN script file that you are using for backup and restore, and keep the default CLIENT_NAME as the hostname:

```
NB_ORA_CLIENT = <host_name>
```

- 3 Specify the required host mode options in the storage array that provides the storage LUNs. For example, in the Host group options of an Hitachi array, enter the type of host (for example, Solaris) and enable the VERITAS Database Edition/Advanced Cluster for Oracle RAC (Solaris) option.
- 4 Add the following line to the `bp.conf` file, on each node in the cluster.

```
PREFERRED_NETWORK = <virtual-host-name>
```

This entry is required when running a rollback operation.

- 5 If the CFS version that you run does not support group quiescence, enable serial quiescence by adding the following lines to the `/usr/opensv/lib/vxfi/configfiles/vxfsfci.conf` file.

```
[QUIESCENCE_INFO]
"QUIESCENCE_SERIAL_QUIESCENCE"=dword:00000001
```

- 6 Ensure that the database is in open (read-write) mode.
- 7 Ensure that the service group for the database in VCS is in the online state.
- 8 Because only the primary node of the cluster supports the rollback restore, change the virtual IP before a rollback operation so that it points to the primary node.
- 9 Configure a snapshot backup policy for the SFRAC environment. In the backup selection tab of the policy, make sure that you provide a path name to the RMAN script. Make sure that the script resides on all the nodes of the cluster. Parameters like ORACLE_SID can differ on each of the nodes. For example, on node 1, the ORACLE_SID can be symc1 and on node 2, the ORACLE_SID can be symc2.

See [“About Snapshot Client support of SFRAC”](#) on page 252.

See [“Performing a rollback restore in an SFRAC environment”](#) on page 255.

Performing a rollback restore in an SFRAC environment

The following procedure describes the manual steps that are needed to restore volumes and file systems by using the snapshot rollback method in an SFRAC environment.

A typical host deployment for running NetBackup for Oracle in an SFRAC environment is as follows: Host A and Host B are in the cluster and Host C is used as an alternate client. The instant recovery snapshot is taken using the Oracle policy and the Hardware Snapshot FIM (frozen image method).

The application I/O stack is built upon the hardware array of VxVM (CVM) and VxFS (CFS).

The Cohesity Cluster Server (VCS) controls the Oracle RAC database and other required essential resources such as shared storage. VCS defines and manages its resources as a single unit called a service group. A service group contains all the necessary components and resources of an application.

The following are entities in the VCS that monitor the application I/O stack:

- CFSMount contains the mount points (cfs) where data files, archive logs, and control files are stored.
- CVMVolDg contains all the Volume Groups (cvm) configured on top of the various array LUNs that participate in the hardware snapshot.
- The Database Resource Group contains the database instance and assists in failover.

To perform a rollback restore in the SFRAC environment

- 1 Ensure that you previously created a virtual IP for the clustered node. Point that virtual IP to the primary node of the cluster.
- 2 On all the clustered nodes, take the VCS database service group (Oracle, CFSMount, and CVMVolDg resources) offline by using the following command:

```
# hagr -offline <DB_Service_Group> -any
```

- 3 Freeze the database service group.

```
# hagr -freeze <DB_Service_Group>
```

- 4 Mount the CFSMount points manually outside VCS on the primary node. This action helps when you start the database in mount state.

```
# mount -F vxfs -o cluster <mntPt>
```

- 5 Start the database with mount option on the clustered primary node using one of these options:

Option 1:

```
# sqlplus /as sysdba
# startup mount;
```

Option 2:

```
# sqlplus /as sysbackup
# startup mount;
```

- 6 Run the rollback restore operation from the RMAN script or the client GUI. A sample RMAN script file (`hot_database_backup_proc`) is located in the following directory path:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman
```

To perform PIT rollback, add the following parameter to the script:

```
NB_ORA_PC_RESTORE=rollback
```

- 7 Unmount the CFS on the primary node.

```
# umount <mntPt>
```

- 8 Unfreeze the VCS database service group.

```
# hagrps -unfreeze <DB_Service_Group>
```

- 9 On all the clustered nodes, take the VCS database service group (Oracle, CFSMount, and CVMVolDg resources) back online:

```
# hagrps -online <DB_Service_Group> -any
```

See [“About Snapshot Client support of SFRAC”](#) on page 252.

See [“Troubleshooting NetBackup in an SFRAC environment”](#) on page 256.

Troubleshooting NetBackup in an SFRAC environment

The following describes some common errors and how to troubleshoot them:

- **Problem:** The backup failed with error code 6 displayed in the GUI. The `rman_script.out` file shows the following error:

```
RMAN-06403: could not obtain a fully authorized session
ORA-01034: ORACLE not available
ORA-27101: shared memory realm does not exist
SVR4 Error: 2: No such file or directory
```

Resolution: Check the `ORACLE_HOME` and `ORACLE_SID` values. `ORACLE_HOME` should not have an extra `'` at the end.

- **Problem:** The backup failed with error code 239. The `dbclient` logs show the following log statement:

```
serverResponse: ERR - server exited with status 239: the specified
client does not exist in the specified policy
01:02:23.844 [4000] <16> CreateNewImage: ERR - serverResponse() fail
```

Resolution: The client name mentioned in the policy and in the `bp.conf` file at the client are different.

- **Problem:** The backup failed with error 156. The `bpfis` logs show the following error.

```
CVxFSPlugin::vxFreezeAll : ioctl VX_FREEZE_ALL failed with errno : 16
CVxFSPlugin::quiesce - Could not quiesce as VX_FREEZE_ALL failed and
VX_FREEZE is not allowed
```

Resolution: Add the following lines to the

`/usr/opensv/lib/vxfi/configfiles/vxfsfi.conf` file:

```
file:[QUIESCENCE_INFO]
"QUIESCENCE_SERIAL_QUIESCENCE"=dword:00000001
```

- **Problem:** The rollback failed with the following error displayed in the GUI:

```
Failed to process backup file <bk_113_1_728619266>
```

The `dbclient` logs show the following error:

```
xbsa_ProcessError: INF - leaving
xbsa_QueryObject: ERR - VxBSAQueryObject: Failed with error: Server
Status: client is not validated to use the server
xbsa_QueryObject: INF - leaving (3)
int_FindBackupImage: INF - leaving
int_AddToFileList: ERR - Failed to process backup file
<bk_113_1_728619266>
```

Resolution : Add the following line in the `bp.conf` file on the primary node of the cluster `PREFERRED_NETWORK = <virtual_name>`

See [“About Snapshot Client support of SFRAC”](#) on page 252.

Script-based Block-Level Incremental (BLI) Backups without RMAN on UNIX and Linux systems

This appendix includes the following topics:

- [About script-based Block-Level Incremental \(BLI\) Backups without RMAN](#)
- [About BLI backup and restore operations](#)
- [Verifying installation requirements for BLI backups without RMAN](#)
- [Creating NetBackup policies for script-based BLI backup](#)
- [Performing backups and restores](#)
- [About troubleshooting backup or restore errors](#)

About script-based Block-Level Incremental (BLI) Backups without RMAN

NetBackup for Oracle with Snapshot Client extends the capabilities of NetBackup to back up only changed data blocks of Oracle database files. NetBackup recommends using RMAN-based BLI backups, which allows tight integration with Oracle administration.

If you choose to use script-based BLI backups without RMAN, you can configure NetBackup for BLI support. A BLI backup backs up only the changed data blocks

of Oracle database files. NetBackup for Oracle script-based BLI performs backups using the Storage Checkpoint facility in the Veritas File System (VxFS) available through the Veritas Storage Foundation for Oracle.

See [“About NetBackup for Oracle with Snapshot Client”](#) on page 174.

About BLI backup and restore operations

A BLI backup performs database backups by obtaining the changed blocks identified by the Storage Checkpoints. BLI backups can also be performed while the database is online or offline. As with Storage Checkpoints, you must enable archive log mode to perform online BLI backups.

A BLI backup places the tablespaces in backup mode, takes a Storage Checkpoint, and then performs the backup. You specify how and when to back up the database when configuring the NetBackup notify scripts.

For example, suppose at 4:00 p.m. you lost a disk drive and its mirrored drive. A number of user tablespaces reside on the disk drive, and you want to recover all committed transactions up to the time you lost the drive. Because the BLI backup facility lets you perform more frequent backups, you did an online differential incremental backup at 1:00 p.m.

You recover by shutting down the database, installing new replacement disk drives, and restoring all the data files with NetBackup. Then you apply the archive logs to recover the tablespaces on the failed drive. If you used Fulldata Storage Checkpoints, the extra redo logs generated during an online backup are small, the media recovery part of the database recovery takes very little time. Moreover, because you have a recent backup, the entire recovery is accomplished quickly.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“Verifying installation requirements for BLI backups without RMAN”](#) on page 260.

Verifying installation requirements for BLI backups without RMAN

Verify the following requirements before you begin the installation.

To verify the installation requirements

- 1 Make sure that the following products are properly installed and configured:
 - NetBackup

- A supported level of Oracle
- NetBackup for Oracle
- Cohesity Storage Foundation for Oracle

2 Verify licensing.

The products must have valid licenses. To check for licenses, enter the following commands based on your version:

For VxFS versions earlier than 3.5:

```
# vxlicense -p
```

For VxFS versions 3.5 or later:

```
# vxlicrep
```

The command displays all the valid licenses that are installed on the system. If you have valid licenses, the Storage Checkpoint feature and the Cohesity Storage Foundation for Oracle appear in the list.

3 Verify that both the NetBackup server (primary and media) and client software work properly.

Particularly, verify that you can back up and restore typical operating system files from the client.

See [“File system and Storage Checkpoint space management”](#) on page 261.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

File system and Storage Checkpoint space management

To support BLI backups, the VxFS file systems need extra disk space to keep track of the block change information. The space that is required depends on the type of checkpoint that is used and the database change rate while the backup is running.

Using Storage Checkpoints has an effect on space in the following ways:

Nodata Storage Checkpoint	If the database is offline during the entire backup window (a cold database backup) or you use this checkpoint type, the additional space is minimal. Each file system requires about 1% of free space. This checkpoint sets a bit to indicate if a file block changed. When you use this checkpoint type, the data files are left in quiesce (write suspend) mode for the duration of the backup.
Fulldata Storage Checkpoint	If the database is online during the backup and using this checkpoint type, then more free space is needed in the file system. NetBackup for Oracle keeps the Oracle containers in quiesce (write suspend) mode only for the time that is needed to create a Storage Checkpoint. During the backup, the checkpoint creates copies of any file blocks immediately before they are changed. The backup up contains only the unchanged blocks and the original copies of the changed blocks. After the backup completes, the Fulldata Storage Checkpoint is converted to a Nodata Storage checkpoint and the copied blocks are returned to the free list. If the workload change rate is light during backup or the backup window is short, 10% free space is usually sufficient for the workload. If the database has a heavy change rate while the backup is running, the file systems may require more than 10% of free space.

Note: The default option that NetBackup uses for backups is Fulldata Storage Checkpoint.

To use Nodata Storage Checkpoint instead of the default option, a user must create the following empty touch file:

```
/usr/opensv/netbackup/ext/db_ext/NODATA_CKPT_PROXY
```

See [“Verifying installation requirements for BLI backups without RMAN”](#) on page 260.

See [“Improving NetBackup backup performance”](#) on page 284.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Creating NetBackup policies for script-based BLI backup

To allow full and incremental backups, you must add at least one Standard type policy to NetBackup and define the appropriate schedules for that policy. NetBackup policies define the criteria for the backup.

These criteria include the following:

- Policy attributes
- Clients and the files or directories to be backed up on the client
- Storage unit to use
- Backup schedules

While most database NetBackup BLI backup policy requirements are the same as for file system backups, the following items have special requirements:

- The number of policies that are required
See [“Number of policies required for BLI backup”](#) on page 263.
- Policy attribute values
See [“About BLI policy attributes”](#) on page 265.
- The BLI client list
See [“About the BLI client list”](#) on page 266.
- The list of directories and files to back up
See [“Backup selections list for BLI backups”](#) on page 266.
- Schedules
See [“About schedules for BLI backup policies”](#) on page 267.

Number of policies required for BLI backup

A database BLI backup requires at least one Standard type policy.

This policy usually includes the following:

- One full backup schedule
- One incremental backup schedule
- One user-directed backup schedule for control files and archive logs

Only one backup stream is initiated for each backup policy during automatic backups. To enable multiple backup streams, define multiple policies for the same database. If you have more than one database SID, configure policies for each SID. If you

intend to do simultaneous backups of more than one SID on the same file system, use Nodata Storage Checkpoints. Set the `METHOD` to `NODATA_CKPT_HOT`.

For example, to back up file systems `F1`, `F2`, `F3`, and `F4` with two streams, you need to define two policies (`P1` and `P2`) with `F1` and `F2` backed up in `P1`, and `F3` and `F4` backed up in `P2`. If you have one large file system that needs to be backed up with multiple streams, divide the files in the file system between different policies. After a file is added to a policy, it should stay in that policy. If you must rearrange the file list, do so only prior to a full backup.

If you have more than one policy defined for an Oracle database instance, NetBackup groups the database instance by the NetBackup keyword phrase. Identify one of the policies as the `POLICY_IN_CONTROL` in the NetBackup notify scripts. This policy performs database shutdowns and restarts. All policies with the same keyword phrase need to be configured to start simultaneously.

Warning: Care must be taken when specifying the keyword phrase. A multistream backup is attempted if the backup process finds more than one policy with the following characteristics: Each policy has the BLI attribute set, each policy is active, each policy contains the same client, and each policy has an identical keyword phrase.

Typical failure status is: "74 - timeout waiting for bpstart_notify to complete."

"See ["NetBackup restore and backup status codes"](#) on page 282.

You can check the file systems on the backup client to see if they are included in one of the NetBackup policies on the server. To see if you need to add any new file systems to the NetBackup policies, run the following commands from the server on a regular basis, perhaps as a `cron(1)` job:

```
# cd /usr/opensv/netbackup/bin/goodies/  
# ./check_coverage -coverage -client mars -mailid \nbadmin
```

The preceding command generates the following output and mails it to the specified mailid:

```
File System Backup Coverage Report (UNIX only)
```

```
-----  
Key:      * - Policy is not active  
          UNCOVERED - Mount Point not covered by an active policy  
          MULTIPLE  - Mount Point covered by multiple active policies
```

```
CLIENT: mars
```

```
Mount Point  Device
```

```
Backed Up By Policy  Notes
```

-----	-----	-----	-----
/	/dev/vg00/lvol3	production_servers	
/home	/dev/vg00/lvol5	production_servers	
/oradata1	/dev/dsk/clt0d0	block_incr1	
/oradata2	/dev/dsk/clt0d0	block_incr1	
/oradata3	/dev/nbuvg/nbuvol	UNCOVERED	
/opt	/dev/vg00/lvol6	production_servers	
/oracle	/dev/vg00/oracle	production_servers	
/stand	/dev/vg00/lvol1	production_servers	
/usr	/dev/vg00/lvol7	production_servers	
/var	/dev/vg00/lvol8	production_servers	

If there is an UNCOVERED file system that is used by Oracle, add it to one of the NetBackup policies so that all the necessary file systems are backed up at the same time.

Note: After a file system is added to a policy, it is a good idea to keep the file system in that policy. If you change the policy, NetBackup performs a full backup the next time backups are run even if an incremental backup is requested.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“About BLI policy attributes”](#) on page 265.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About BLI policy attributes

NetBackup applies policy attribute values when it backs up files.

The following attributes must be set for BLI backup:

Policy Type	Set to Standard.
Perform block level incremental backups	Select to enable BLI backups. If the BLI attribute is not enabled, NetBackup uses the standard method to back up the files in the file list.
Job Priority	Set so that the BLI backup policies run before other policies.

Keyword phrase	Define as the Oracle database instance name (<code>\$ORACLE_SID</code>) in each of the policies for the same instance. Multistream backups start when all the policies with a particular keyword phrase complete their respective startup scripts. If you have multiple Oracle database instances (SIDs) use a separate set of policies for each SID. If the SIDs are backed up simultaneously and any share a common file system for data files, use Nodata Storage Checkpoints. Set the <code>METHOD</code> to <code>NODATA_CKPT_HOT</code> .
-----------------------	---

Note: Do not change a keyword phrase after it is set in a policy. The keyword phrase is used in naming Storage Checkpoints. Changing the keyword phrase necessitates a full backup even if an incremental backup is requested.

The [NetBackup Administrator's Guide, Volume I](#) describes other policy attributes and how to configure them.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“About the BLI client list”](#) on page 266.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About the BLI client list

The client list specifies the clients upon which you configured a BLI backup. For a database backup, specify the name of the machine upon which the database resides. Specify the virtual hostname if clustered.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“Backup selections list for BLI backups”](#) on page 266.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Backup selections list for BLI backups

The backup selections list specifies a list of directories and files to back up. The list must contain all the database files or their directory names. Using directory names, rather than file names, ensures that new database files added to an existing

configuration are backed up without having to update the file list. Use the `check_coverage` script to make sure all file systems are backed up.

If you are using the Quick I/O interface, you need to specify both the Quick I/O file name and the associated hidden file in the file list (for example, `dbfile` and `.dbfile`), or you need to specify the directory that contains both files. NetBackup does not follow the symbolic links to automatically back up the hidden file if you enumerate only the `dbfile` explicitly in the backup selections list. They are both included if you enumerate their common directory.

When the NetBackup scheduler invokes an automatic backup schedule, it backs up the files one at a time, in the same order they appear in the backup selection list.

Oracle does not recommend backing up the online redo log, so it is recommended that you place online redo log files in a different file system than datafiles, archive log files, or database control files. Do not include the online redo log files in the file list.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“About schedules for BLI backup policies”](#) on page 267.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About schedules for BLI backup policies

The NetBackup server starts these schedule types:

- Full Backup
- Differential Incremental Backup
- Cumulative Incremental Backup

Each BLI backup policy must include one full backup schedule and at least one incremental backup schedule. In addition, you must designate one of the BLI backup policies as the `POLICY_IN_CONTROL`. The policies for each stream must have the same types of schedules.

The [NetBackup Administrator's Guide, Volume I](#) describes other schedule attributes and how to configure them.

You can configure the following types of schedules:

- User-directed backup schedule. The user-directed backup schedule encompasses all the days and times when user-directed backups are allowed to occur. Set the backup window as described.

The policies for each stream must have the same types of schedules.

- Automatically initiated backup schedules. Include server-initiated backup schedules to specify the days and times for NetBackup to automatically start backups of the files specified in the policy file list. Set the backup window as described.

For server-initiated full and incremental backup schedules, set the start times and durations to define the appropriate windows for the backups. Follow the same procedure used to define backup schedules for other policies. For more information on these procedures, see the [NetBackup Administrator's Guide, Volume I](#).

The backups are started by the scheduler only within the backup window specified. For the `POLICY_IN_CONTROL`, include in the user-directed backup schedule the time periods when the BLI backup policies complete.

Set the retention level and periods to meet user requirements.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Example Oracle BLI backup policy

The following example shows attributes and schedules for an Oracle BLI backup policy.

```
Policy Name: oracle_backup1
Policy Type: Standard
Active: yes
Block level incremental: yes
Job Priority: 0
Max Jobs/Policy: 1
Residence: oracle_tapes
Volume Pool: NetBackup
Keyword: ORA1
Client List: Sun4 Solaris2.6 mars
              HP9000-800 HP-UX11.00 mars
Backup Selections List: /oradata/oradata1
Schedule: full
Type: Full Backup
Frequency: 1 week
Retention Level: 3 (one month)
Daily Windows:
```

```

        Sunday    18:00:00  -->  Monday    06:00:00
        Monday    18:00:00  -->  Tuesday    06:00:00
        Tuesday   18:00:00  -->  Wednesday  06:00:00
        Wednesday 18:00:00  -->  Thursday   06:00:00
        Thursday   18:00:00 -->  Friday     06:00:00
        Friday     18:00:00 -->  Saturday   06:00:00
        Saturday   18:00:00 -->  Sunday     06:00:00

Schedule:          incr
  Type:            Differential Incremental Backup
  Frequency:       1 day
  Retention Level: 3 (one month)
  Daily Windows:
        Sunday    18:00:00  -->  Monday    06:00:00
        Monday    18:00:00  -->  Tuesday    06:00:00
        Tuesday   18:00:00  -->  Wednesday  06:00:00
        Wednesday 18:00:00  -->  Thursday   06:00:00
        Thursday   18:00:00 -->  Friday     06:00:00
        Friday     18:00:00 -->  Saturday   06:00:00
        Saturday   18:00:00 -->  Sunday     06:00:00

Schedule:          userbkup
  Type:            User Backup
  Retention Level: 3 (one month)
  Daily Windows:
        Sunday    00:00:00  -->  Sunday    24:00:00
        Monday    00:00:00  -->  Monday    24:00:00
        Tuesday   00:00:00  -->  Tuesday   24:00:00
        Wednesday 00:00:00  -->  Wednesday 24:00:00
        Thursday   00:00:00 -->  Thursday   24:00:00
        Friday     00:00:00 -->  Friday     24:00:00
        Saturday   00:00:00 -->  Saturday   24:00:00

```

In this example, the `oracle_backup1` policy backs up all the files in `/oradata/oradata1`. The policy specifies a weekly full backup, a daily differential incremental backup, and a user-directed backup schedule. The archive logs and the control file are backed up using the user-directed schedule at the completion of the full or incremental backup.

Setting the Maximum jobs per client global attribute

Set the **Maximum jobs per client** to the number of policies that have the same keyword phrase. This number can be greater than one when multiple job policies are defined to back up multiple file systems.

About BLI backup methods

You can choose from the following backup methods when configuring BLI notify scripts:

Table C-1 BLI backup terminology

Term	Definition
cold database backup	A cold database backup is taken while the database is offline or closed. BLI backup shuts down the database and performs either full or block-level incremental backups. This backup method is also referred to in Oracle documentation as a "consistent whole database backup" or a "closed backup." The data from a cold backup is consistent, resulting in easier recovery procedures. To select this backup method, set <code>METHOD</code> to <code>SHUTDOWN_BKUP_RESTART</code>. In an offline backup, all database files are consistent to the same point in time (for example, when the database was last shutdown using typical methods). The database must stay shut down while the backup runs.
hot database backup	A hot database backup allows the database to be online and open while the backup is performed. With the Storage Checkpoint facility, this backup method runs database backups in parallel so a database does not need to be in backup mode for a long time. To select this backup method, set <code>METHOD</code> to <code>ALTER_TABLESPACE</code>. Hot backups are required if the database must be up and running 24 hours a day, 7 days a week. To use hot backups, the database must be in <code>ARCHIVELOG</code> mode. BLI backup uses the <code>alter tablespace begin backup</code> command and the <code>alter tablespace end backup</code> command to put the database into and take it out of backup mode. Oracle documentation refers to this method as an inconsistent whole database backup or open backup. Unlike the cold database backup method, the data in hot backups is fuzzy or inconsistent until the appropriate redo log files (online and archived) are applied after the restore operation to make the data consistent.

Table C-1 BLI backup terminology (*continued*)

Term	Definition
Nodata storage checkpoint hot	A Nodata storage checkpoint hot backup puts the tablespaces in backup mode for the duration of the backup. It uses a Nodata Storage Checkpoint to reduce the amount of file system space consumed. To select this backup method, set <code>METHOD</code> to <code>NODATA_CKPT_HOT</code>. Use this method if all of the following conditions are present: ■ You are backing up multiple Oracle database instances. ■ More than one instance shares the file system. ■ The backup of the instances can overlap in time.
quick freeze database backup	The quick freeze database backup is different than an online database backup, because it requires the database to be brought down briefly to take a snapshot or Fulldata Storage Checkpoint of the database image. The Fulldata Storage Checkpoint is created in a few seconds and the database can be restarted immediately. A backup image from a quick freeze database backup is equivalent to a backup image from a cold database backup. You can choose this backup method when you configure BLI notify scripts. To select this backup method, set <code>METHOD</code> to <code>SHUTDOWN_CKPT_RESTART</code>. See “Creating notify scripts for BLI backups” on page 271.

If the database is in `ARCHIVELOG` mode, you can use all four methods to back up the database. If the database is in `NOARCHIVELOG` mode, you can only select the cold backup or quick freeze backup.

When you use the cold and quick freeze database backups, the default shutdown command that you use in the `bpstart_notify.oracle_bli` script is `shutdown` or `shutdown normal`. These commands wait for all users to log off before it initiates the shutdown. In some circumstances, even after all interactive users are logged off, processes such as the Oracle Intelligent Agent (Oracle `dbsnmp` account) can still be connected to the database, preventing the database shutdown. Attempt to use the default shutdown commands to shut down the database cleanly. Alternatively, you can use `shutdown immediate` to initiate the database shutdown immediately.

Creating notify scripts for BLI backups

Create the notify scripts that run on the clients to synchronize the backup operation and the database operation. You need a set of three notify scripts for each policy

that performs BLI backups. The scripts must be in the `/usr/opensv/netbackup/bin` directory on the NetBackup client.

The scripts are named as follows:

- `bpstart_notify.POLICY`
- `post_checkpoint_notify.POLICY`
- `bpend_notify.POLICY`

To create the notify scripts, run the following script as root:

```
/usr/opensv/netbackup/ext/db_ext/oracle/bin/setup_bli_scripts
```

This script copies the sample notify script templates to `/usr/opensv/netbackup/bin` and makes the necessary changes based on the information you provide.

The notify script templates are located on the local computer in the following location:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples
```

When you run `setup_bli_scripts` you need to supply the following information:

- Identify the `POLICY_IN_CONTROL`
See [“Identify the POLICY_IN_CONTROL for BLI backups”](#) on page 272.
- Provide the Oracle environment variables
See [“Oracle environment variables for BLI scripts”](#) on page 272.
- Select a backup method
- Notify scripts for other policies
See [“About BLI notify scripts for other policies”](#) on page 273.

Identify the POLICY_IN_CONTROL for BLI backups

If you have more than one policy defined on the server for one Oracle database instance, identify one of the policies as the `POLICY_IN_CONTROL`. This is the policy that initiates the database `shutdown`, `startup`, or `alter tablespace` commands. The `POLICY_IN_CONTROL` can be any policy (for example, the first policy defined). This variable is stored in the notify scripts.

Oracle environment variables for BLI scripts

If you create notify scripts, or if you run `setup_bli_scripts`, you need to provide values for the Oracle environment variables.

These variables are as follows:

ORACLE_DBA	User name of the Oracle database administrator. Typically, <code>oracle</code> .
ORACLE_BASE	<code>\$ORACLE_BASE</code> of the Oracle database instance.
ORACLE_HOME	<code>\$ORACLE_HOME</code> of the Oracle database instance.
ORACLE_SID	Oracle database instance ID (<code>\$ORACLE_SID</code>) if it is different from the keyword.
ORACLE_LOGS	Directory in which the Oracle archive logs reside.
ORACLE_CNTRL	Location to which a copy of the Oracle control file is written so that it can be backed up.
SQLCMD	<code>sqldba</code> , <code>svrmgrl</code> , or <code>sqlplus</code> command to start up or shut down the database.
ORACLE_INIT	Path name for the Oracle startup parameter file (<code>INIT.ORA</code>). If you are using an Oracle <code>SPFILE</code> as your parameter file, do not set the <code>ORACLE_INIT</code> environment variable.
ORACLE_CONFIG	Path name for the Oracle configuration file (<code>CONFIG.ORA</code>). Some database configurations use the <code>CONFIG.ORA</code> file to specify values for the database parameters that usually do not change. The <code>CONFIG.ORA</code> file can be called by the <code>INIT.ORA</code> file using an include statement.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About BLI notify scripts for other policies

If you have more than one policy defined to support multiple backup streams, create a copy of the notify scripts for each policy defined.

For example, assume that you have two policies defined, `oracle_backup1` and `oracle_backup2`. Also assume that `POLICY_IN_CONTROL` is set to `oracle_backup1`. You also need to create notify scripts for policy `oracle_backup2`. The `setup_bli_scripts` script performs this step automatically.

See [“Sample setup_bli_scripts session”](#) on page 274.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See “About script-based Block-Level Incremental (BLI) Backups without RMAN” on page 259.

See “Example Oracle BLI backup policy” on page 268.

Sample setup_bli_scripts session

The following sample session shows how to use `setup_bli_scripts` to create the notify scripts.

```
#!/usr/opensv/netbackup/ext/db_ext/oracle/bin/setup_bli_scripts
```

```
Please enter the user name of your Oracle administrator? orac901
```

```
ORACLE_BASE is the Oracle environment variable that identifies  
the directory at the top of the Oracle software and administrative  
file structure. The value of this variable is typically  
/MOUNTPOINT/app/oracle
```

```
Please enter your ORACLE_BASE? /dbhome/oracle/orac901
```

```
ORACLE_HOME is the Oracle environment variable that identifies the  
directory containing the Oracle software for a given Oracle server  
release. The value of this variable is typically  
/dbhome/oracle/orac901/product/RELEASE
```

```
Please enter your ORACLE_HOME? /dbhome/oracle/orac901
```

```
sqlplus will be used.
```

```
The default "connect" statement that will be used to connect to the database is:  
"connect / as sysdba"
```

```
Would you like to modify the connect and use a specific login? (y/n) n
```

```
"connect / as sysdba" will be used.
```

```
Please enter the Oracle instance (ORACLE_SID) you want to back up? orac901
```

```
If you are using a CONFIG.ORA file, you need to specify where  
it is, so that it can be backed up. If this does not apply  
apply to your configuration, hit ENTER to go on. If this does  
apply to your configuration, specify the file path.
```

Typically this would be:

```
/dbhome/oracle/orac901/admin/orac901/pfile/configorac901.ora  
but this file could not be found.
```

Enter your Oracle config file path or hit ENTER:

To back up a copy of the Oracle control file, you need to specify a file path where Oracle can write a copy of the control file.

Please enter the file path where Oracle is to write a copy of your control file? /dbhome/oracle/orac901/admin/orac901/pfile/cntrlorac901.ora

To back up the Oracle archive logs, you need to specify their location.

Enter the directory path to your Oracle archive logs?
/dbhome/oracle/orac901/admin/orac901/arch

Do you have more archive log locations? (y/n): n

Do you want the output of successful executions of the NetBackup scripts mailed to you? y
Please enter the mail address to send it to? jdoe@company.com

Do you want the output of unsuccessful executions of the NetBackup scripts mailed to you? y
Please enter the mail address to send it to? jdoe@company.com

There are 4 backup methods to choose from:

- ALTER_TABLESPACE - Use alter tablespace begin backup method
- NODATA_CKPT_HOT - Use alter tablespace begin backup with nodata ckpts
- SHUTDOWN_CKPT_RESTART - Shutdown, create the ckpt clones, and restart
- SHUTDOWN_BKUP_RESTART - Shutdown the DB, backup, and then restart

If one of the methods requiring DB shutdown are selected, you may experience problems with timeouts if the database can't be shut down in a timely manner. You may want to change the shutdown command in the notify scripts to shutdown immediate, or you may have to increase the BPSTART_TIMEOUT value in the bp.conf file on the primary server, or you may want to change the backup method to ALTER_TABLESPACE or NODATA_CKPT_HOT.

Note: the default BPSTART_TIMEOUT value is 300 seconds.

Do you want to use the ALTER_TABLESPACE method? y

You now need to decide on how many NetBackup policies you will have backing up simultaneously. The first one you enter will be known as the `POLICY_IN_CONTROL` in the scripts and will perform any needed DB operations. When you create the policies on the NetBackup server, you will have to divide the filesystems between these policies.

```
Please enter the name of the policy that will be the POLICY_IN_CONTROL? BLI_1
Please enter the name of another policy or DONE to stop? BLI_2
Please enter the name of another policy or DONE to stop? BLI_3
Please enter the name of another policy or DONE to stop? BLI_4
Please enter the name of another policy or DONE to stop? BLI_5
Please enter the name of another policy or DONE to stop? BLI_6
Please enter the name of another policy or DONE to stop? DONE
```

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“Example Oracle BLI backup policy”](#) on page 268.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Performing backups and restores

After the installation and configuration are complete, you can use the NetBackup interfaces to start Oracle backups and restores. You can run backups manually by using schedules that you determine. You can also run a schedule manually.

Note: You must be the root user to perform all operations using the BLI backup software.

See [“About NetBackup for Oracle agent automatic backups”](#) on page 276.

See [“About NetBackup for Oracle manual backups”](#) on page 277.

See [“Backing up Quick I/O files”](#) on page 278.

See [“Restoring BLI backup images”](#) on page 279.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About NetBackup for Oracle agent automatic backups

The best way to back up databases is to set up schedules for automatic backups.

Note: You must be the root user to perform all operations using the BLI backup software.

Note: For HP-UX PA-RISC checkpoints to unmount and be cleaned up, create touch file `/usr/opensv/netbackup/AIO_READS_MAX` that contains the value 1.

HP-UX PA-RISC checkpoints may not be unmounted on Oracle database agents.

See [“Creating NetBackup policies for script-based BLI backup”](#) on page 263.

See [“Backing up Quick I/O files”](#) on page 278.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

See [“Performing backups and restores”](#) on page 276.

About NetBackup for Oracle manual backups

You can also run an Automatic Backup schedule manually. For information about performing manual backups of schedules, see the [NetBackup Administrator's Guide, Volume I](#).

Note: You must be the root user to perform all operations using the BLI backup software.

Note: For HP-UX PA-RISC checkpoints to unmount and be cleaned up, create touch file `/usr/opensv/netbackup/AIO_READS_MAX` that contains the value 1.

To perform a cold (offline) backup, set the environment variable `METHOD` in the `bpstart_notify` script on the client to `SHUTDOWN_BKUP_RESTART`. The `bpstart_notify` script shuts down the database before the backup begins and the `bpend_notify` script restarts the database after the backup completes.

To perform a hot (online) backup using Fulldata Storage Checkpoints, make sure the database is running in `ARCHIVELOG` mode and set the variable `METHOD` to `ALTER_TABLESPACE`. The `bpstart_notify` script changes the tablespaces to online backup mode before the backup begins, and the `post_checkpoint_notify` script changes the tablespaces back to normal mode after the Fulldata Storage Checkpoints are created.

To perform a Nodata Storage Checkpoint Hot (online) backup, make sure the database is running in `ARCHIVELOG` mode and set the environment variable `METHOD`

in the `bpstart_notify` script to `NODATA_CKPT_HOT`. The `bpstart_notify` script changes the tablespaces to online backup mode before the backup begins. The `bpend_notify` script changes the tablespaces back to normal mode after the backup completes.

To perform a quick freeze backup, set the environment variable `METHOD` in the `bpstart_notify` script to `SHUTDOWN_CKPT_RESTART`. The `bpstart_notify` script shuts down the database and the `post_checkpoint_notify` script restarts it immediately after the Fulldata Storage Checkpoints are created. Taking VxFS Fulldata Storage Checkpoints is very fast (within a minute), and with the NetBackup queuing delay for scheduling the backup jobs, the database down time is typically only a few minutes.

Backing up Quick I/O files

A Quick I/O file consists of two components: a hidden file with the space allocated for it, and a link that points to the Quick I/O interface of the hidden file. Because NetBackup does not follow symbolic links, you must specify both the Quick I/O link and its hidden file in the list of files to be backed up.

Note: You must be the root user to perform all operations using the BLI backup software.

For example:

```
ls -la /db02
total 2192
drwxr-xr-x 2 root  root    96 Jan 20 17:39 .
drwxr-xr-x 9 root  root  8192 Jan 20 17:39 ..
-rw-r--r-- 1 oracle dba 1048576 Jan 20 17:39 .cust.dbf
lrwxrwxrwx 1 oracle dba    22 Jan 20 17:39 cust.dbf ->\
    .cust.dbf::cdev:vxfv:

```

The preceding example shows that you must include both the symbolic link `cust.dbf` and the hidden file `.cust.dbf` in the backup file list.

If you want to back up all Quick I/O files in a directory, you can simplify the process by only specifying the directory to be backed up. In this case, both components of each Quick I/O file is properly backed up. In general, you should specify directories to be backed up unless you only want to back up some files in those directories.

Note: For HP-UX PA-RISC checkpoints to unmount and be cleaned up, create touch file `/usr/openv/netbackup/AIO_READS_MAX` that contains the value 1.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

See [“Performing backups and restores”](#) on page 276.

See [“Restoring BLI backup images”](#) on page 279.

Restoring BLI backup images

Restoring the backup images that a BLI backup creates is no different than restoring the backup images that are created using the default NetBackup configuration. Restoring to any of the incremental backup images requires NetBackup to restore the last full backup image and all the subsequent incremental backups until the specified incremental backup image is restored. NetBackup does this automatically. The media that stored the last full and the subsequent incrementals must be available, or the restore cannot proceed.

You can start the restore operations from the NetBackup client by using the Backup, Archive, and Restore interface. To restore the latest copy of each file, select either the files or parent directories with the latest backup date, and click **Restore**. For more information on restoring, see the [NetBackup Backup, Archive, and Restore Getting Started Guide](#).

If the operation is to restore files from an incremental backup image, NetBackup issues multiple restore operations beginning from the last full backup image and the subsequent incremental backup images until the selected date. The activity of multiple restores is logged in the Progress Log.

If you plan to restore files backed up by another client or to direct a restore to another client, start the restore from the NetBackup server using the Backup, Archive, and Restore interface. Before you initiate a restore, a backup must have successfully completed or an error occurs during the execution.

For Solaris, the restore destination can be a VxFS or UFS file system. The destination file system does not need to support the Storage Checkpoint feature, but to be able to perform BLI backups of the restored data, a VxFS file system with the Storage Checkpoint feature is required.

For HP-UX, the restore destination can be a VxFS or HFS file system. The destination file system does not need to support the Storage Checkpoint feature to restore files. However, a VxFS file system with the Storage Checkpoint feature is required to perform BLI backups of the restored data.

For AIX, the restore destination can be a VxFS or JFS file system. The destination file system does not need to support the Storage Checkpoint feature to restore files. However, a VxFS file system with the Storage Checkpoint feature is required to perform BLI backups of the restored data.

Note that restoring a file causes all blocks in that file to be rewritten. Thus, all the blocks in the file are considered to have been modified. Thus, the first subsequent differential incremental backup and all subsequent cumulative incremental backups back up all of the blocks in the restored file. If you are restoring an entire database or a file system, the first subsequent backup backs up all blocks that are restored.

To restore a Quick I/O file, if both the symbolic link and the hidden file already exist, NetBackup restores both components from the backup image. If either one of the two components is missing, or both components are missing, NetBackup creates or overwrites as needed.

Oracle database recovery might be necessary after restoring the files. See the Oracle documentation for more information on doing database recovery.

See [“Performing backups and restores”](#) on page 276.

See [“Backing up Quick I/O files”](#) on page 278.

See [“About BLI backup and database recovery”](#) on page 285.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About NetBackup backup and restore logs

NetBackup provides logs on the database backup and restore operations and are useful for finding problems with those operations. The following table describes the most useful logs and reports for troubleshooting backup and restore operations.

Table C-2 NetBackup backup and restore logs

Log file type	Description
NetBackup progress logs	This description only applies to a File System Restore (script-based BLI). For user-directed backups and restores performed with Backup, Archive, and Restore interface, the most convenient log to use for NetBackup is the progress log. The progress log file is written to the user's home directory, by default in <code>/usr/opensv/netbackup/logs/user_ops/username/logs</code>. This log indicates whether NetBackup was able to complete its part of the operation. You can view the progress log from the Backup, Archive, and Restore interface, or you can use a file editor such as <code>vi(1)</code>.
NetBackup debug logs	The NetBackup server and client software provide debug logs for troubleshooting any problems that occur outside of BLI backups. To enable these debug logs on the server or client, create the appropriate directories under the following directory: <code>/usr/opensv/netbackup/logs</code> For more information on debug logs, see the NetBackup Troubleshooting Guide or see the <code>/usr/opensv/netbackup/logs/README.debug</code> file.

Table C-2 NetBackup backup and restore logs (*continued*)

Log file type	Description
NetBackup reports	In addition to logs, NetBackup provides a set of reports that help isolate problems. One report is All log entries on the server. The All logs entries report is also available in the NetBackup web UI. For a description of all reports, see the NetBackup Administrator's Guide, Volume I.

See [“Restoring BLI backup images”](#) on page 279.

See [“About troubleshooting backup or restore errors”](#) on page 281.

See [“About BLI backup and database recovery”](#) on page 285.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About troubleshooting backup or restore errors

A backup or restore error can originate from NetBackup for Oracle, from the NetBackup server or client, from the Media Manager, or from VxFS. In addition to examining log files and reports, you should determine at which stage of the backup or restore operation the problem occurred. You can also use NetBackup status codes to determine the cause of the problem.

See [“About NetBackup backup and restore logs”](#) on page 280.

See [“About BLI backup and database recovery”](#) on page 285.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Troubleshooting stages of backup and restore operations

Refer to the following list to determine the source of a backup or restore error:

- A backup or restore can be started in either of the following ways:
 - Manually from the administrator interface on the primary server
 - Automatically by a NetBackup server using a full schedule or incremental schedule

If an error occurs during the start operation, examine the Java reports window for the possible cause of the error.
- If the backup or restore starts successfully but eventually fails, one of the following can be the cause:

- Server/Client communication problem
- Schedule error
- Media-related error
- VxFS errors

For more information, see the [NetBackup Troubleshooting Guide](#).

- There can be insufficient disk space for the VxFS Fulldata Storage Checkpoints to keep track of changed block information. Check the `All Log Entries` report for errors.

If there is a file system out-of-space condition, increase the size of the file system so it is large enough for Fulldata Storage Checkpoints or use the Nodata Storage Checkpoint Hot backup method. This error does not affect the integrity of the backup images because a full backup of the affected file system occurs after the condition is fixed.

- If an incremental backup is intended, but the whole file system is backed up instead, one of the following conditions might be present:
 - Storage Checkpoints that keep track of changes have been removed
 - The **Block level incremental** attribute is not selected
 - Other errors with a nonzero status code

The most common cause of this problem is the file system removed the Storage Checkpoint that keeps track of the block changes. This action might occur if the file system runs out of space, and there are no volumes available to allocate to the file system. The integrity of the backup images is not affected, because a full backup of the file system occurs at the next backup opportunity after NetBackup detects that a Storage Checkpoint is missing.

See “[NetBackup restore and backup status codes](#)” on page 282.

See “[About NetBackup backup and restore logs](#)” on page 280.

See “[About BLI backup and database recovery](#)” on page 285.

See “[About script-based Block-Level Incremental \(BLI\) Backups without RMAN](#)” on page 259.

NetBackup restore and backup status codes

The status codes and their meanings are as follows:

- Status Code 9. An extension package is needed but was not installed. The client does not have the NetBackup binaries required to do BLI backups. Use `update_clients` on the server to push out new binaries. Also, use

`vxlicense -p` to verify that the Storage Checkpoint feature [83] and the Cohesity Storage Foundation for Oracle [100] are installed.

- Status Code 69. Invalid file list specification.

Look for a message such as the following in the error log on the server:

```
FTL - /oradata is not in a VxFS file system. A block incremental backup of it is not possible.
```

This indicates that there was an attempt to back up a file system that is not a VxFS file system with the **Block level incremental** attribute. This error can also occur if the file system is not mounted.

- Status Code 73. `bpstart_notify` failed.

When running the notify scripts, the `bpstart_notify` script exited with a nonzero status code, or the permission bits are set wrong on the `bpstart_notify` script. The script must have execute permission. If the permission bits are set, check the `bpstart_notify_output.Oracle_SID` file in the `/usr/opensv/netbackup/bin/BLOCK_INCR` directory.

- Status Code 74. Client timed out waiting for `bpstart_notify` to complete.

Check the `BPSTART_TIMEOUT` setting on the NetBackup server. The `BPSTART_TIMEOUT` specified did not allow enough time for the script to complete. The shutdown database operation might be taking too long, or the script might be waiting for other streams to start. Check the `bpstart_notify_output.Oracle_SID` file and the `post_checkpoint_notify_output.Oracle_SID` file in the `/usr/opensv/netbackup/bin/BLOCK_INCR` directory. Make sure that the policies and schedules are configured with appropriate multiplexing factors and that the required storage units that allow all streams to start at the same time are configured. Check to see if all needed tape drives are working and available. Make sure that the database is not processing transactions so that the instance cannot be shut down immediately (if you are using one of the backup methods where the database is shut down).

Finally, make sure that the priority on the BLI policies is higher than other policies, so they get access to the tape drives before the other policies.

- Status Code 75. Client timed out waiting for `bpend_notify` to complete.

Check the `BPEND_TIMEOUT` setting on the NetBackup server. The `BPEND_TIMEOUT` specified did not allow enough time for the script to complete. The restart database operation might be taking too long, or the script might be waiting for other streams to call the `bpend_notify` script. Check the `bpend_notify_output.Oracle_SID` file and the `post_checkpoint_notify_output.Oracle_SID` file in the `/usr/opensv/netbackup/bin/BLOCK_INCR` directory. Make sure that the policies

and schedules are configured with appropriate multiplexing factors and that the required storage units that can allow all streams to be started at the same time are configured. Verify that all needed tape drives are working and available during backup.

- Status Code 77. Execution of the specified system command returned a nonzero status code.

Check the `post_checkpoint_notify_output.KEYWORD` file in the `/usr/opensv/netbackup/bin/BLOCK_INCR` directory for the possible cause. The `post_checkpoint_notify` script exited with a nonzero status code.

- Status Code 143. Invalid command protocol.

Check to see if the **Block level incremental** policy attribute is selected without a keyword specified. Set the **Keyword phrase** in the policies to the Oracle database instance name (`$ORACLE_SID`).

See [“Troubleshooting stages of backup and restore operations”](#) on page 281.

See [“About NetBackup backup and restore logs”](#) on page 280.

See [“About BLI backup and database recovery”](#) on page 285.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

Improving NetBackup backup performance

If backups are running slowly, check to see if the database has an excessive workload. BLI backups allow hot database backups and quick freeze database backups. Because the database is running during both of these backup methods while NetBackup is backing up the database files, Oracle I/O can affect the backup performance.

If the database is not running with a high transaction volume, troubleshoot NetBackup. If the incremental backup takes a long time to finish, it could mean that there are more changed blocks since the last incremental backup. Verify whether the size of the incremental backup image has increased, and consider increasing the frequency of incremental backups.

Finally, you can improve the speed at which backup is performed by using multiplexed backups. Assigning multiple policies to the same backup device is helpful when devices are not writing at their maximum capacity.

See [“Troubleshooting stages of backup and restore operations”](#) on page 281.

See [“Verifying installation requirements for BLI backups without RMAN”](#) on page 260.

See [“File system and Storage Checkpoint space management”](#) on page 261.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

About BLI backup and database recovery

A BLI backup does not perform automatic database recovery. This process includes restoring the database files from NetBackup images and applying the Oracle redo log files to the database files. Follow the Oracle documentation to perform database recovery after a restore.

See [“NetBackup restore and backup status codes”](#) on page 282.

See [“Troubleshooting stages of backup and restore operations”](#) on page 281.

See [“About NetBackup backup and restore logs”](#) on page 280.

See [“Restoring BLI backup images”](#) on page 279.

See [“About script-based Block-Level Incremental \(BLI\) Backups without RMAN”](#) on page 259.

XML Archiver

This appendix includes the following topics:

- [NetBackup for Oracle XML export and XML import](#)
- [About the environment variables set by a user in the XML export parameter file](#)
- [About XML export shell scripts](#)
- [Performing an XML export archive](#)
- [Browsing XML export archives using bporaimp parameters](#)
- [Browsing XML export archives using bplist](#)
- [Restoring an XML export archive](#)
- [Troubleshooting XML export or XML import errors](#)
- [Additional XML export and import logs](#)

NetBackup for Oracle XML export and XML import

While Oracle RMAN performs backup, restore, and recovery of physical Oracle database objects (data files, tablespaces, control files, and archived redo logs), the NetBackup for Oracle XML export and XML import utilities provide backup and restore of logical database objects (tables, users, and rows).

The XML format is used to provide a self-identifying and system-independent format ideal for database archiving.

See [“NetBackup for Oracle XML export and import archiving features”](#) on page 287.

See [“XML export archive process”](#) on page 288.

See [“XML import restore process”](#) on page 291.

See [“Sequence of operation: XML import restore”](#) on page 292.

See [“Sequence of operation: XML export archive”](#) on page 289.

See [“NetBackup for Oracle features”](#) on page 14.

NetBackup for Oracle XML export and import archiving features

[Table D-1](#) describes NetBackup for Oracle XML export and XML import archiving features.

Table D-1 NetBackup for Oracle XML export and XML import archiving features

Feature	Description
System- and database-independent archive format	NetBackup for Oracle uses the eXtensible Markup Language (XML) standard to represent relational database table data that is extracted from an Oracle database. The eXtensible Markup Language (XML) is a universal format for structured documents and data. The XML 1.0 standards are produced by the World Wide Web Consortium and include the XML Schema standard. Unicode UTF-8 is the character set encoding generated by NetBackup for Oracle. Standard XML processors support UTF-8. US7ASCII is a strict subset of UTF-8.
Self-identifying archive format	The XML Schema standard is used to describe the table data that is included in an archive. In this way, the archive contains the key to understanding the format of the data as well as the data itself.
Command line interfaces that allow export and import at row-level granularity	Parameter files specify the table data to include in an archive and the table data to extract from an archive for import into an Oracle database.
Restore destination option	NetBackup for Oracle can either restore XML data to an operating system directory or import the data back into the Oracle database.
Flexible archive image searches	The NetBackup catalog contains information on the contents of the archive that can be searched by using flexible search criteria, such as tablename or user.

See [“XML export archive process”](#) on page 288.

See [“XML import restore process”](#) on page 291.

See [“NetBackup for Oracle XML export and XML import”](#) on page 286.

See [“Sequence of operation: XML import restore”](#) on page 292.

See “About NetBackup for Oracle” on page 13.

See “NetBackup for Oracle features” on page 14.

XML export archive process

Figure D-1 shows the XML export archive process.

Figure D-1 XML export archives




Table D-2 describes the archive activity.

Table D-2 Archive activity

Activity	Process
Oracle XML archive	NetBackup for Oracle extracts database table data, converts it into XML format, and stores XML data to either of the following types of repositories: ■ A directory ■ A storage unit
XML export	NetBackup converts Oracle table data to XML format (XML schema or metadata and XML instance or data).
Archive	NetBackup stores the XML data on a NetBackup storage unit.
<code>bpوراexp/bporaexp64</code> command	NetBackup for Oracle's XML export utility converts Oracle database table data into a self-identifying XML schema document and instance document. NetBackup archives them or redirects them to an OS directory.

See [“NetBackup for Oracle XML export and XML import”](#) on page 286.

See [“Sequence of operation: XML export archive”](#) on page 289.


See [“Sequence of operation: XML import restore”](#) on page 292.

See [“XML import restore process”](#) on page 291.

Sequence of operation: XML export archive

[Figure D-2](#) shows data flow.

Figure D-2 XML export archive data flow



NetBackup for Oracle users or automatic schedules start the database XML export archives. This operation is done by performing a manual backup of an Oracle policy, by invoking a script at the command line on the client.

For an XML export archive:

- The NetBackup for Oracle script calls the `bpوراexp` utility with a specified parameter file.
- The query processor uses the parameters in the specified file to build an SQL query for each table.
- Oracle's OCI API executes the queries on the Oracle database instance to be archived.
- The query processor passes the output (including metadata and data for a single table or multiple tables) to the XML Generator.

- For each table passed, the XML generator builds one or more sets of XML schema and XML instance documents.
- XML data streams are backed up by NetBackup.
- Alternately, `bporaexp` allows the files to be saved to an operating system directory.

See “XML export archive process” on page 288.

See “XML import restore process” on page 291.

See “NetBackup for Oracle XML export and XML import” on page 286.

See “Sequence of operation: XML export archive” on page 289.

XML import restore process

Figure D-3 shows the XML import restore process.




Table D-3 describes the restore activity.

Table D-3 Restore activity

Activity	Process
Oracle XML Restore	NetBackup for Oracle manages the retrieval of archived database table data, the parsing of the XML format, and the insertion of the data back into the Oracle database.
Restore	NetBackup retrieves the XML-formatted data from the storage unit.
XML import	NetBackup for Oracle parses XML-formatted Oracle table data and inserts data into the Oracle database.

Table D-3 Restore activity (*continued*)


Activity	Process
bporaimp/bporamip64 commands	NetBackup for Oracle's XML import utility can parse the XML-formatted data for re-insertion into the database or can redirect the data to an OS directory.

- See [“XML export archive process”](#) on page 288.
- See [“NetBackup for Oracle XML export and XML import”](#) on page 286.
- See [“Sequence of operation: XML import restore”](#) on page 292.
- See [“Sequence of operation: XML export archive”](#) on page 289.
- See [“NetBackup for Oracle features”](#) on page 14.

Sequence of operation: XML import restore

[Figure D-4](#) shows data flow.

Figure D-4 XML import restore data flow



NetBackup for Oracle users start database XML import restores by invoking a NetBackup for Oracle script at the client command line.

For an XML import restore:

- The NetBackup for Oracle script calls the `bporaimp` utility with a specified parameter file.
- The input parameters that identify the XML archive to restore are passed to NetBackup.
- NetBackup locates and reads the set of XML schema and instance documents from the NetBackup storage unit.
- The XML data stream is passed to an XML parser, which passes the data to the XML data loader.
- The XML data loader uses Oracle's OCI API to insert the data into the database.

Optionally, `bporaimp` allows the XML data stream to bypass the XML parser and be sent to an operating system directory. In addition, users can restore the table metadata only into an operating system directory. `bporaimp` also allows import from an operating system directory into Oracle.

See [“XML export archive process”](#) on page 288.

See [“XML import restore process”](#) on page 291.

See [“NetBackup for Oracle XML export and XML import”](#) on page 286.

See [“Sequence of operation: XML export archive”](#) on page 289.

About the environment variables set by a user in the XML export parameter file

You can set the XML export parameter file in the Oracle user’s environment.

See [“About the environment variables set by a user in the XML export parameter file”](#) on page 294.

[Table D-4](#) shows the NetBackup for Oracle environment variables.

Table D-4 NetBackup for Oracle environment variables

Environment variable	Purpose
NB_ORA_SERV	Specifies the name of NetBackup primary server.
NB_ORA_CLIENT	Specifies the name of the Oracle client. On Windows, this variable is useful for specifying a virtual client name in a cluster.
NB_ORA_POLICY	Specifies the name of the policy to use for the Oracle backup. To define NB_ORA_POLICY, use the RMAN PARMS statement or send statement in Oracle shell scripts. For example: <code>ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE' ;</code> <code>send 'NB_ORA_POLICY=Oracle_Backup' ;</code> <code>BACKUP</code>
NB_ORA_SCHED	Specifies the name of the Application Backup schedule to use for the Oracle backup.

See [“About the bp.conf file on UNIX systems”](#) on page 92.

About XML export shell scripts

The following sections describe XML export shell scripts.

The user writes the shell scripts and they must conform to the operating system's shell syntax. Sample XML export and import shell scripts are installed on the client with the NetBackup for Oracle agent. Modify these scripts to meet your individual requirements.

See [“About creating RMAN scripts manually”](#) on page 94.

See [“About storing shell scripts”](#) on page 97.

Creating XML export scripts manually

When the database agent was initially installed, the installation software wrote example scripts to the following locations:

- For export:

Windows:

```
install_path\NetBackup\dbext\Oracle\samples\bporaexp
```

UNIX:

```
/usr/openv/netbackup/ext/db_ext/oracle/samples/bporaexp
```

- For import:

Windows:

```
install_path\NetBackup\dbext\Oracle\samples\bporaimp
```

UNIX:

```
/usr/openv/netbackup/ext/db_ext/oracle/samples/bporaimp
```

The example export scripts that are installed in `bporaexp` are as follows:

Windows:

```
data_archiver_export.cmd
```

UNIX:

```
data_archiver_export.sh
```

```
data_archiver_export64.sh
```

```
bporaexp_help.param  
bporaexp_partitions.param  
bporaexp_table_to_files.param  
bporaexp_tables.param  
bporaexp_tables_rows.param
```

The example import scripts that are installed in `bporaimp` are as follows:

Windows:

```
data_archiver_import.cmd
```

UNIX:

```
data_archiver_import.sh  
data_archiver_import64.sh  
  
bporaimp_archive.param  
bporaimp_archive_schema_to_files.param  
bporaimp_archive_to_users.param  
bporaimp_bfile_table.param  
bporaimp_help.param  
bporaimp_ignore_rows_table.param  
bporaimp_large_table.param  
bporaimp_list.param  
bporaimp_old_archive.param  
bporaimp_partitions.pram  
bporaimp_table_from_files.param  
bporaimp_table_to_files.param  
bporaimp_table_to_user.param  
bporaimp_tables.param
```

To use the example scripts

- 1 Copy the example scripts to a different directory on your client. Oracle scripts can be located anywhere on the client.
- 2 Modify each script for your environment.
- 3 On UNIX, make sure that the `su` command logs into the correct user.

If you do not include an `su - user` (*user* is Oracle administrator account) in your Oracle scripts, they do not run with the proper permissions and environment variables. The result is problems with your database backups and restores.

See [“About XML export shell scripts”](#) on page 295.

See [“About storing shell scripts”](#) on page 97.

See [“About the NetBackup for Oracle sample scripts”](#) on page 95.

Performing an XML export archive

The following sections describe how to perform an XML export archive.

Table D-5 Tasks and commands

Task	The commands that are used to accomplish the task
Automatic backup of an Oracle policy	As with Oracle backups using RMAN, the most convenient way to create archives that consist of XML exports of data from your database is to create schedules for automatic backups. The Oracle policy runs NetBackup for Oracle shell scripts. When the NetBackup scheduler invokes a schedule for an automatic backup, the NetBackup for Oracle XML export shell scripts run as follows: ■ In the same order as they appear in the file list ■ On all clients in the client list The NetBackup for Oracle XML export shell scripts start the XML export by running the NetBackup <code>bporaexp</code> or <code>bporaexp64</code> utility.
Manual backup of an Oracle policy	The administrator can use the NetBackup server software to manually run an automatic backup schedule for the Oracle policy. For more information, see the NetBackup Administrator's Guide, Volume I. See “Testing configuration settings for NetBackup for Oracle” on page 98.
User-directed XML exports from the client	The following sections describe procedures for performing user-directed XML exports. ■ See “Running the NetBackup for Oracle XML export script on the client” on page 297. ■ See “Running bporaexp on the client as an Oracle user” on page 298. ■ See “Writing to a directory versus writing to a storage unit” on page 299. ■ See “About bporaexp parameters” on page 300.

Running the NetBackup for Oracle XML export script on the client

You can initiate a database XML export from the operating system command prompt: Type the full path to the shell script that performs the export. For example:

Windows:

```
install_path\oracle\scripts\data_archiver_export.cmd
```

UNIX:

```
/oracle/scripts/data_archiver_export.sh
```

The operating system shell starts the database XML export archive by running the XML export script. The XML export script contains commands to run `bporaexp`.

The NetBackup for Oracle installation script installs sample scripts in the following location:

Windows:

```
install_path\NetBackup\dbext\oracle\samples\bporaexp
```

UNIX:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/bporaexp
```

See [“Performing an XML export archive”](#) on page 297.

Running bporaexp on the client as an Oracle user

As an Oracle user you can also run the `bporaexp` command (`bporaexp64` on some platforms) from the operating system command prompt and specify a parameter file.

To run bporaexp on the client as an Oracle user

- 1 Create a parameter file that specifies the settings that determine how the backup is to be performed. Information is available about the `bporaexp` parameters.

See [“About bporaexp parameters”](#) on page 300.

- 2 Run the following command to specify the parameter file:

```
# bporaexp [username/password] parfile = filename | help=y
```

- 3 Configure the runtime environment, because this method does not call the full script that includes the runtime configuration.

On UNIX and Linux, check the sample scripts for runtime environment details.

See [“About configuring the run-time environment”](#) on page 86.

`bporaexp` creates a set of XML schema and instance documents that can be used to archive Oracle table data. For each archive, one primary XML schema (`.xsd`) document is generated. In addition, `bporaexp` generates a table-specific schema (`.xsd`) document and a table specific instance (`.xml`) document for each table. Additional files are created if the table contains `LONG` or `LOB` columns.

See [“Performing an XML export archive”](#) on page 297.

Writing to a directory versus writing to a storage unit

One important aspect of the parameter file is the `DIRECTORY` parameter. If you specify the `DIRECTORY` parameter, the `bp oraexp` (`bp oraexp64` on some platforms) command writes the backup files to the operating system directory you specify. NetBackup does not write the files to a storage unit.

For example, assume that the archive `test1` contains one table, `USER1`. If the `directory` parameter is specified, NetBackup creates certain files when you run the `bp oraexp` command.

Windows:

```
DIRECTORY=\db\netbackup\xml
```

UNIX:

```
DIRECTORY=/db/netbackup/xml
```

Table D-6 shows the files NetBackup creates when you run the command.

Table D-6 NetBackup files for example table `USER1`

File	Content
Windows: \db\netbackup\xml\test1\test1.xsd UNIX: /db/netbackup/xml/test1/test1.xsd	Master XML schema for table <code>USER1</code>
Windows: \db\netbackup\xml\test1\USER1\TEST1.xsd UNIX: /db/netbackup/xml/test1/USER1/TEST1.xsd	Table schema for table <code>USER1</code>
Windows: \db\netbackup\xml\test1\USER1\TEST1.xml UNIX: /db/netbackup/xml/test1/USER1/TEST1.xml	XML document for table <code>USER1</code>

If the `DIRECTORY` parameter is not specified, NetBackup writes the backup images to a storage unit. A NetBackup backup set is created and cataloged under the name:

Windows:

```
\Oracle\XMLArchive
```

UNIX:

```
/Oracle/XMLArchive
```

All NetBackup for Oracle `bp oraexp` backups are cataloged using this convention.

Alternatively, if the parameter file does not contain the `DIRECTORY` parameter, NetBackup creates and catalogs the following files:

Windows:

```
\Oracle\XMLArchive\test1\test1.xsd
\Oracle\XMLArchive\test1\USER1\TEST1.xsd
\Oracle\XMLArchive\test1\USER1\TEST1.xml
```

UNIX:

```
/Oracle/XMLArchive/test1/test1.xsd
/Oracle/XMLArchive/test1/USER1/TEST1.xsd
/Oracle/XMLArchive/test1/USER1/TEST1.xml
```

In production, do not use the `DIRECTORY` parameter in the `bporaexp` parameter file. When you write to a storage unit, NetBackup offers the features that include searching and cataloging with the NetBackup catalog and automatic handling of output that exceeds file system limits. With the `DIRECTORY` parameter, file system limits, such as a 2 GB maximum, can cause an error.

To run `bporaexp` on the client, run the following command:

```
bporaexp [username/password] parfile = filename | help=y
```

On some UNIX platforms, the `bporexp64` command is used.

See [“Performing an XML export archive”](#) on page 297.

About bporaexp parameters

This topic describes the available `bporaexp` (`bporexp64` on some platforms) parameters.

Note the following:

- Use the NetBackup parameters `NB_ORA_SERV`, `NB_ORA_CLIENT`, `NB_ORA_POLICY`, and `NB_ORA_SCHED` to specify the NetBackup runtime configuration. Otherwise, the order of precedence for the runtime configuration variable settings is used.
- Some parameters are valid only when you write to a storage unit. Other parameters are valid only when you write to a directory. In the following table, the Target Location column contains either Storage Unit or Directory to indicate whether the parameter in that row applies to writing to a storage unit or to a directory. Parameters that are recognized when you write to a directory are ignored when you write to a storage unit.

[Table D-7](#) shows the available `bporaexp` parameters with their default values.

Table D-7 bpوراexp parameters and default values

Parameter	Required?	Default	Description	Target location
CONSISTENT	N	N	Specifies if bpوراexp uses the SET TRANSACTION READ ONLY statement to ensure that the data from all tables is consistent to a single point in time and does not change during the execution of the bpوراexp command. If the default of CONSISTENT=N is used, each table is exported as an independent transaction.	Directory
DIRECTORY	N	no default	Optionally specifies a directory for the output of the bpوراexp utility.	Directory
HELP	N	N	Displays a help message with descriptions of bpوراexp parameters. Does not export data if HELP=Y.	Directory
KEYWORD	N	no default	Optionally specifies a keyword phrase that NetBackup associates with the image being created by the archive operation. Values for KEYWORD must be in double quotes.	Storage Unit
LOG	N	no default	Optionally specifies a file name to receive informational and error messages. If so, messages are logged on the log file and not displayed to the terminal display.	Directory
NAME	Y	no default	The name of the primary XML schema file.	Directory
NB_ORا_SERV	N	default primary server	Optionally specifies the name of the NetBackup primary server.	Storage Unit
NB_ORا_CLIENT	N	default client	Optionally specifies the name of the NetBackup for Oracle client.	Storage Unit
NB_ORا_POLICY	N	default Oracle policy	Optionally specifies the name of the NetBackup for Oracle policy.	Storage Unit
NB_ORا_SCHED	N	default backup policy schedule	Optionally specifies the name of the backup policy schedule to use.	Storage Unit

Table D-7 bpوراexp parameters and default values (*continued*)

Parameter	Required?	Default	Description	Target location
OWNER	N	no default	Lists the Oracle schema owners to export. For each owner, the tables, partitions, and views that are owned by that Oracle account are exported by default. The <code>PARTITIONS</code> and <code>VIEWS</code> parameters can be used to exclude partitions and views.	Directory
PARTITIONS	N	Y	Optionally specifies whether or not table partitions are included. Only valid when used with the <code>OWNER</code> parameter.	Directory
QUERY	N	no default	Selects a subset of rows from a set of tables. The value of the query parameter is a string that contains a <code>WHERE</code> clause for a SQL select statement that is applied to all tables and table partitions listed in the <code>TABLES</code> parameter. For example, if <code>TABLES = emp,bonus</code> and <code>QUERY = "where job = 'SALESMAN' and sal < 1600"</code>, two SQL statements are run: ■ <code>SELECT*FROM emp where job='SALESMAN' and sal<1600;</code> ■ <code>SELECT*FROM bonus where job='SALESMAN' and sal<1600;</code> Each query that runs refers to a single table at a time in the <code>FROM</code> clause, so it is illegal to have a join in the <code>WHERE</code> clause.	Directory
ROW_BUFFER	N	1000	Specifies the size, in rows, of the buffer used to fetch rows. Tables with <code>LONG</code> columns are fetched one row at a time. The maximum value allowed is 32767.	Directory
TABLES	Y	no default	Lists the table names, view names, and partition names to export. The <code>USERID</code> must have <code>SELECT</code> privilege on the tables and views. The syntax used is: <code>schema.table: partition name or schema.view name</code>	Directory

Table D-7 bporaexp parameters and default values (*continued*)

Parameter	Required?	Default	Description	Target location
USERID	Y	no default	Specifies the username/password (and optional connect string) of the user initiating the export. If a connect string is not provided, the ORACLE_SID environment variable is used.	Directory
VIEWS	N	Y	Optionally specifies whether or not views are included. Only valid when used with the OWNER parameter.	Directory

See [“Performing an XML export archive”](#) on page 297.

Browsing XML export archives using bporaimp parameters

To use the `bporaimp` (`bporaimp64` on some platforms) command to browse XML export archives created by using `bporaexp` (`bporaexp64` on some platforms), create a parameter file with the desired search criteria. First, set the variables `LIST=Y` and `USERID=username/`. Only the archives created using the Oracle `USERID` are listed.

The Oracle password is not required. The operating system account that is running `bporaimp` has access only to archives that were created using the same account.

Note: Only XML export archives created using NetBackup mode are searched. Exports stored in an operating system directory using the `DIRECTORY` parameter are not searched.

Use the `NB_ORA_SERV` and `NB_ORA_CLIENT` parameters to specify the NetBackup server and client. Otherwise, the order of precedence for the runtime configuration variable settings is used. You can also include the `LOG` parameter.

Information is available on the `LIST`, `LOG`, `NB_ORA_CLIENT`, `NB_ORA_SERV`, and `USERID` parameters.

See [“About bporaimp parameters”](#) on page 307.

[Table D-8](#) shows other parameters you can include in the parameter file.

Table D-8 Parameters you can include in a parameter file

Parameter	Default	Description
ARCHIVE_DATE_FROM	no default	Optionally specifies a start date for the archive search. Used with ARCHIVE_DATE_TO to specify a range. The date format is <i>mm/dd/yyyy [hh:mm:ss]</i> .
ARCHIVE_DATE_TO	no default	Optionally specifies an end date for the archive search. Used with ARCHIVE_DATE_FROM to specify a range. The date format is <i>mm/dd/yyyy [hh:mm:ss]</i> .
KEYWORD	no default	Optionally specifies a keyword phrase for NetBackup to use when searching for archives.
NAME	no default	The name of the primary XML schema file.
FROMUSER	no default	Optionally specifies a comma-separated list of table owners.
TABLES	no default	Optionally specifies a list of table and partition names that were included in an archive.

For example, assume you named the list parameter file `bporaimp_list.param`. At the command prompt, type the following:

```
bporaimp parfile = bporaimp_list.param
```

Note: On some UNIX platforms, the `bporaimp64` command is used.

See [“Browsing XML export archives using bplist”](#) on page 304.

See [“Restoring an XML export archive”](#) on page 306.

See [“Browsing backups using the bplist command”](#) on page 108.

Browsing XML export archives using bplist

For a higher level view of the Oracle XML export archive list, you can use the `bplist` command. The result is the list of XML schema and instance document file names.

Note: Only XML export archives created using NetBackup mode are searched. Exports stored in an operating system directory using the `DIRECTORY` parameter are not searched.

The following UNIX or Linux example uses `bplist` to search all Oracle archives for a client named `jupiter`. The sample output is produced for two archives, `test1` and `little_sales`, where each archive has one Oracle table (`test1` has `USER1.TEST1` and `little_sales` has `USER1.LITTLE_SALES`).

```
/usr/opensv/netbackup/bin/bplist -C jupiter -t 4 -R /Oracle/XMLArch/  
/Oracle/XMLArchive/test1/test1.xsd  
/Oracle/XMLArchive/test1/USER1/TEST1.xsd  
/Oracle/XMLArchive/test1/USER1/TEST1.xml  
/Oracle/XMLArchive/little_sales/little_sales.xsd  
/Oracle/XMLArchive/little_sales/USER1/LITTLE_SALES.xsd  
/Oracle/XMLArchive/little_sales/USER1/LITTLE_SALES.xml  
/exb_n2bm5bco_1_1392342936  
/exb_mabm02ko_1_1392170136  
/exb_lqbltds6_1_1392083334
```

The following Windows example uses `bplist` to search all Oracle archives for a client named `jupiter`. The sample output is produced for one archive, `test`.

```
install_path\NetBackup\bin\bplist -C jupiter -t 4 -R Oracle:\XMLArch\  
Oracle:\XMLArchive\test\test.xsd  
Oracle:\XMLArchive\test\SCOTT\BONUS.xsd  
Oracle:\XMLArchive\test\SCOTT\BONUS.xml  
Oracle:\XMLArchive\test\SCOTT\DEPT.xsd  
Oracle:\XMLArchive\test\SCOTT\DEPT.xml  
Oracle:\XMLArchive\test\SCOTT\EMP.xsd  
Oracle:\XMLArchive\test\SCOTT\EMP.xml  
Oracle:\XMLArchive\test\SCOTT\SALGRADE.xsd  
Oracle:\XMLArchive\test\SCOTT\SALGRADE.xml
```

The `-t 4` on this command specifies the Oracle backups or archives. The `-R` specifies the default number of directory levels to search, 999.

For more information on this command, see the `bplist` man page in the [NetBackup Commands Reference Guide](#).

See “[Browsing XML export archives using bporaimp parameters](#)” on page 303.

See “[Restoring an XML export archive](#)” on page 306.

See “[Browsing backups using the bplist command](#)” on page 108.

Restoring an XML export archive

Before you attempt to restore an archive, make sure that the XML archive has successfully completed. You can identify the correct archive to restore by browsing the XML export archives. NetBackup generates an error if an archive backup history does not exist.

The following sections describe procedures for performing user-directed restores.

See [“Browsing XML export archives using bporaimp parameters”](#) on page 303.

See [“Browsing XML export archives using bplist”](#) on page 304.

See [“Running an XML import script on the client”](#) on page 306.

See [“Running bporaimp on the client in NetBackup for Oracle”](#) on page 307.

See [“About bporaimp parameters”](#) on page 307.

See [“About redirecting a restore of an XML export archive to a different client”](#) on page 311.

Running an XML import script on the client

You can initiate a restore from the operating system command prompt by typing the full path to the XML import script that initiates the restore. For example:

Windows:

```
install_path\oracle\scripts\data_archiver_import.cmd
```

UNIX:

```
/oracle/scripts/data_archiver_import.sh
```

The operating system shell starts the database restore by running the XML import script file. The XML import script file contains commands to run `bporaimp` (`bporaimp64` on some platforms).

The NetBackup for Oracle installation script writes sample scripts to the following location:

Windows:

```
install_path\NetBackup\dbext\oracle\samples\bporaimp
```

UNIX:

```
/usr/opensv/netbackup/ext/db_ext/oracle/samples/bporaimp
```

See [“Restoring an XML export archive”](#) on page 306.

See [“Running bporaimp on the client in NetBackup for Oracle”](#) on page 307.

See [“About bporaimp parameters”](#) on page 307.

Running bporaimp on the client in NetBackup for Oracle

Run the `bporaimp` command from the operating system command line on the client using the appropriate parameter file.

The Windows account that runs `bporaimp` has access only to XML export archives that were created using the same Windows account.

The UNIX account that runs `bporaimp` has access only to XML export archives that were created using the same UNIX account. Be sure to configure the runtime environment, because this method does not call the full script that includes the runtime configuration. Check the sample scripts for runtime environment details.

To run bporaimp on the client

- ◆ At the command prompt, type the `bporaimp` command in the following format:

```
bporaimp [username/password] parfile = filename | help=y
```

On some UNIX platforms, the `bporaimp64` command is used.

See [“About bporaimp parameters”](#) on page 307.

About bporaimp parameters

Use the NetBackup parameters `NB_ORA_SERV` and `NB_ORA_CLIENT` to specify the NetBackup runtime configuration. Otherwise, the order of precedence for the runtime configuration variable settings is used.

Some parameters are valid only when writing to a storage unit. Other parameters are valid only when writing to a directory. In the following table, the right-most column contains either “Storage Unit” or “Directory” to indicate whether the parameter in that row is applicable for either writing to a storage unit or to a directory. Parameters that are recognized when writing to a directory are ignored when writing to a storage unit.

[Table D-9](#) describes the `bporaimp` (`bporaimp64` on some platforms) parameters and default values.

Table D-9 bporaimp parameters and default values

Parameter	Required?	Default	Description	Target location
ARCHIVE_DATE_FROM	N	no default	Optionally specifies a start date for the archive to be imported. Used with ARCHIVE_DATE_TO to specify a range. If not used, the most recent archive is imported. If the range used results in more than one archive, the most recent from the range is used. The date format is <i>mm/dd/yyyy [hh:mm:ss]</i> .	Storage Unit
ARCHIVE_DATE_TO	N	no default	Optionally specifies an end date for the archive to be imported. Used with ARCHIVE_DATE_FROM to specify a range. If not used, the most recent archive is imported. If the range used results in more than one archive, the most recent from the range is used. The date format is <i>mm/dd/yyyy [hh:mm:ss]</i> .	Storage Unit
BFILE_DIRECTORY	Y (if any table being imported has BFILE columns)	no default	Specifies a directory for the output of any BFILE columns being imported. Oracle's CREATE DIRECTORY command can be used to create the DIRECTORY in Oracle, and the name should match the name used in the export file.	Directory
COMMIT	N	N	Specifies whether bporaimp should commit after each array insert. The size of the array is determined by ROW_BUFFER. By default, bporaimp commits only after loading each table, and performs a rollback when an error occurs, before continuing with the next object.	Directory
DIRECTORY	N	no default	Optionally specifies a directory for the input of the bporaimp utility.	Directory
FROMUSER	N	no default	Optionally specifies a comma-separated list of users to import from an archive containing multiple users' tables. If not specified, all of the tables are imported.	Directory
HELP	N	N	Displays a help message with descriptions of bporaimp parameters.	Directory

Table D-9 bporaimp parameters and default values (*continued*)

Parameter	Required?	Default	Description	Target location
IGNORE_ROWS	N	N	Specifies whether or not rows should be inserted into a table that is not empty. The default is that the table already exists and that it is empty. If it is not empty, IGNORE_ROWS = N causes an error to be reported, and the table is skipped with no rows inserted. IGNORE_ROWS = Y causes rows to be inserted with errors reported in the log file. If IGNORE_ROWS = Y and an error such as a primary key constraint violation occurs, no data is inserted if COMMIT = N. However, if COMMIT = Y, the array of rows (size determined by ROW_BUFFER) is not inserted, but bporaimp continues to process additional row arrays in the order in which they were exported. To cause all rows that do not violate a primary key constraint to be inserted, set COMMIT = Y, ROW_BUFFER = 1, and IGNORE_ROWS = Y.	Directory
KEYWORD	N	no default	Optionally specifies a keyword phrase for NetBackup to use when searching for archives from which to restore files.	Storage Unit
LIST	N	N	LIST = Y queries the NetBackup catalog and lists the archives available. Does not import the data if LIST = Y.	Storage Unit
LOG	N	no default	Optionally specifies a file name to receive informational and error messages. If this parameter is specified, messages are logged in the log file and not displayed to the terminal display.	Directory
NAME	Y	no default	The name of the primary XML schema file. This parameter is required if LIST = N.	Directory
NB_ORA_SERV	N	default primary server	Optionally specifies the name of the NetBackup primary server.	Storage Unit
NB_ORA_CLIENT	N	default client	Optionally specifies the name of the NetBackup for Oracle client.	Storage Unit

Table D-9 *bporaimp* parameters and default values (continued)

Parameter	Required?	Default	Description	Target location
RESTORE_SCHEMA_ONLY	N	N	Used with <code>RESTORE_TO_DIRECTORY</code> to restore the XML schema files only to a directory.	Storage Unit
RESTORE_TO_DIRECTORY	N	no default	Optionally specifies a directory for the output of the <i>bporaimp</i> utility. If used, the XML data is not parsed and inserted into Oracle.	Storage Unit
ROW_BUFFER	N	1000	Specifies the size, in rows, of the buffer used to insert rows. Tables with <code>LONG</code> or <code>LOB</code> columns are inserted one row at a time. The maximum value allowed is 32767.	Directory
TABLES	N	no default	Optionally specifies a list of table, view, and partition names to import. If not used, all objects in the archive are imported. The objects must already exist, and the <code>USERID</code> must have <code>INSERT</code> privilege on the objects. The object names cannot be qualified with owner names, and the <code>FROMUSER</code> parameter is used to specify a particular owner. If a partition name is specified, it indicates the exported partition only and the rows are inserted according to the partitioning scheme of the target table. If the export contains partitions, and the import does not specify them, all are inserted.	Directory
TOUSER	N	no default	Optionally specifies a comma-separated list of users to import to that can be used with the <code>FROMUSER</code> parameter to change the table owners. The <code>TOUSER</code> Oracle accounts must already exist, and the <code>USERID</code> must have <code>INSERT</code> privilege on the tables that must also exist.	Directory
USERID	Y	no default	Specifies the username/password (and optional connect string) of the user initiating the import. If a connect string is not provided, the <code>ORACLE_SID</code> environment variable is used.	Directory

See [“Restoring an XML export archive”](#) on page 306.

See [“Running bporaimp on the client in NetBackup for Oracle”](#) on page 307.

See [“Browsing XML export archives using bporaimp parameters”](#) on page 303.

See “[Example - Using bporaimp for a redirected restore](#)” on page 312.


About redirecting a restore of an XML export archive to a different client

With NetBackup for Oracle you have the option to restore an XML export archive to a client other than the one that originally performed the XML export. The process of restoring data to another client is called a redirected restore.

Before you redirect the restore, see the following topic:

[Figure D-5](#) illustrates a typical redirected restore.

Figure D-5 Redirected restore of an XML export archive to a different client



The user on client A cannot initiate a redirected restore to client B. Only the user on client B, which is the client receiving the backup image, can initiate the redirected restore. Any user who belongs to the database group that performed the backup can restore it, unless the `BKUP_IMAGE_PERM` variable is set to `USER`.

Redirecting a restore of an XML export archive to a new client using bporaimp

On UNIX and Linux, any user who belongs to the database group that performed the archive can restore XML export archive. The `BKUP_IMAGE_PERM` variable must be set to `GROUP` or `ANY`, not `USER`.

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its

group is `srvgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

Perform the following procedure on the new client host if you want to restore XML export archives that are owned by another client.

To redirect a restore of an XML export archive to a new client using `bporaimp`

- 1 Set the environment variables for `bporaimp` (`bporaimp64` on some platforms) on the new client, including `ORACLE_HOME` and `ORACLE_SID`.
- 2 In the `bporaimp` parameter file, include the following lines:

```
nb_ora_serv = Netbackup_server
nb_ora_client =
original_client_where_XML_export_occurred.
```

- 3 Specify any other `bporaimp` parameters.
See [“Running bporaimp on the client in NetBackup for Oracle”](#) on page 307.
- 4 Run `bporaimp`.

Example - Using `bporaimp` for a redirected restore

For example, assume the following:

- Original client is `jupiter`
- New client is `saturn`
- Server is `jupiter`
- `ORACLE_SID` is `test` on both `saturn` and `jupiter`
- Windows user is `ora` on both `jupiter` and `saturn`
- UNIX user is `ora` on both `jupiter` and `saturn`
- Archive name is `sales`

To use bporaimp for a redirected restore (example)

- 1 Create the following file on server `jupiter`:

Windows:

```
install_path\NetBackup\db\altnames\saturn
```

UNIX:

```
/usr/opensv/netbackup/db/altnames/saturn
```

- 2 Edit the preceding file to contain the name `jupiter`.

- 3 Log on to `saturn` as `ora`.

- 4 Create file `bporaimp.param`.

Include the following parameters:

```
NAME = sales
NB_ORA_SERV = jupiter
NB_ORA_CLIENT = jupiter
USERID = orauser/orapasswd
```

- 5 Run `bporaimp parfile=bporaimp.param` to restore sales archive to saturn and to import the data into the `test` database on saturn.

See [“Redirecting a restore of an XML export archive to a new client using bporaimp”](#) on page 311.

See [“About bporaimp parameters”](#) on page 307.

See [“Restoring an XML export archive”](#) on page 306.

Troubleshooting XML export or XML import errors

An XML export or import error can originate from NetBackup or from Oracle, as follows:

- On the NetBackup side, an error can be from the `bporaexp` or `bporaimp` programs, the NetBackup server or client, or Media Manager.
- On the Oracle side, an error can be from the target database.

Use the following steps when troubleshooting a failed operation:

- Check the logs to determine the source of the error.
- Troubleshoot each stage of the XML export or XML import.
The following sections describe these steps in detail.

On UNIX and Linux, these sections describe the log files from the `bporaexp` and `bporaimp` commands. The logs are created in

`/usr/opensv/netbackup/logs/bporaexp` or

`/usr/opensv/netbackup/logs/bporaimp`.

See [“Checking the logs to determine the source of an error”](#) on page 314.

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“Troubleshooting each stage of the XML export or XML import”](#) on page 315.

See [“Additional XML export and import logs”](#) on page 318.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Checking the logs to determine the source of an error

This topic describes how to check the logs to determine the source of an error.

To check the logs

1 Check the `bporaexp` or `bporaimp` log.

If the `LOG` parameter is specified in the `bporaexp` or `bporaimp` command's `parfile`, the commands write logs to the file that is specified as the argument to the `LOG` parameter. The commands write log information to the screen if `LOG` is not specified.

For example, incorrect installation or configuration causes the following common problems:

- The `ORACLE_HOME` environment variable was not set.
- The `bporaexp` or `bporaimp` program was unable to connect to the target database.

When `bporaexp` and `bporaimp` are used and the backup images are written to an operating system directory, these logs are the only source of error logging and tracking.

2 Check the NetBackup logs.

On Windows, the first NetBackup log to check is

`install_path\NetBackup\logs\bporaexp\log.mmddyy` or

`install_path\NetBackup\logs\bporaimp\log.mmddyy`.

On UNIX, the first NetBackup log to check is

`/usr/opensv/netbackup/logs/bporaexp/log.mmddyy` or

`/usr/opensv/netbackup/logs/bporaimp/log.mmddyy`.

Examine these logs for any messages that show how to determine the source of an error.

The NetBackup client writes these logs and they contain the following:

- Requests from `bporaexp` and `bporaimp`
- Activities between `bporaexp` and `bporaimp` and NetBackup processes

If the logs do not contain any messages, the following conditions could be present:

- `bporaexp` or `bporaimp` terminated before service is requested from NetBackup.
- `bphdb` (if started by the scheduler or graphical user interface) did not start the shell script successfully. Check the `bphdb` logs for `stderr` and `stdout` files.

Try to run the XML export or XML import script file from the command line to determine the problem.

On UNIX, the error is usually due to a file permission problem for `bphdb` itself or for the export or the import script file.

Ensure that the full XML export or import script file name is entered correctly in the **Backup Selections** list of the Oracle policy configuration. Also verify that this script name is correct.

For more information about debug logs and reports, refer to the [NetBackup Administrator's Guide, Volume I](#).

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Troubleshooting each stage of the XML export or XML import

The information in this section does not apply to you if `DIRECTORY` is specified in `bporaexp` or `bporaimp` command's `parfile`.

The following explains the sequence of events for an action that `bporaexp` or `bporaimp` initiates in NetBackup mode. This situation occurs when `DIRECTORY` is not specified in the `bporaexp` or `bporaimp` command's `parfile`. It suggests solutions for the problems that can occur at each point in the sequence.

To troubleshoot by stage

- 1 `bporaexp` or `bporaimp` starts.

An export or import can be started in any of the following ways:

- Command line from the system prompt.
For example:

```
bporaexp parfile = parameter_filename  
bporaimp parfile = parameter_filename
```

- Automatically by an automatic export schedule.

If an error occurs now, check the `bporaexp` or `bporaimp` log.

- 2 `bporaexp` or `bporaimp` verifies its environment and then connects to Oracle and NetBackup.

An Oracle environment problem, a database problem, an incorrect user ID, or an incorrect password can cause Oracle connect errors.

A NetBackup error now is usually due to a problem with client and server communication. Check the messages in the `bprd` and `bpbd` logs for clues.

Also verify the `bp.conf` entries on the UNIX or Linux client.

- 3 `bporaexp` or `bporaimp` issues a backup or restore request.

Before the backup or restore request proceeds, `bporaexp` or `bporaimp` commands perform three functions:

- Gather necessary parameters
- The backup or restore request is sent to the NetBackup server
- Wait until the server and client are ready to transfer data

The NetBackup client interfaces gather information from the following places:

- The environment, including `bporaexp` and `bporaimp` parameter files.
If you use scripts, you have to generate the parameter file manually.
- Server configuration parameters on Windows.
- The user's `bp.conf` and `/usr/opensv/netbackup/bp.conf` files on the UNIX or Linux client.

This information is sent to the primary server's `bprd` process.

To troubleshoot a backup problem in this part of the sequence, examine the following file:

Windows:

```
install_path\NetBackup\logs\bporaexp\log.mmddyy
```

UNIX:

```
/usr/opensv/netbackup/logs/bporaexp/log.mmddyy
```

If the `bprd` process failed, check the `bprd` and `bpbrm` logs.

During this sequence, most failures occur because of incorrect NetBackup server or Oracle policy configuration parameters.

NetBackup can usually select the correct Oracle policy and schedules. However, NetBackup can select a policy or schedule in error if there are several Oracle policies in its database.

In Windows, try setting the `SERVER` and `POLICY` values in the client environment or by setting the following `bporaexp` parameters:

```
NB_ORA_POLICY=policyname
NB_ORA_SCHED=schedule
NB_ORA_SERV=Netbackup_server
NB_ORA_CLIENT=Netbackup_client
```

In UNIX, try setting the `SERVER` and `POLICY` values in the `bp.conf` file on the client or by setting the following `bporaexp` parameters:

```
NB_ORA_POLICY=policyname
NB_ORA_SCHED=schedule
NB_ORA_SERV=Netbackup_server
NB_ORA_CLIENT=Netbackup_client
```

To troubleshoot a restore, examine the following log file:

Windows:

```
install_path\NetBackup\logs\bporaimp\mmdyy.log
```

UNIX:

```
/usr/opensv/netbackup/logs/bporaimp/log.mmdyy
```

Make sure that the correct NetBackup server and NetBackup client values are used by setting the following `bporaimp` parameters:

```
NB_ORA_SERV=Netbackup_server
NB_ORA_CLIENT=Netbackup_client
```

Set these parameters to the same values that were used for the XML export operation.

- 4 `bpوراexp` or `bpوراimp` issues read or write requests to the NetBackup client, which then transfers data to or from the NetBackup server.

`bpوراexp` builds an SQL query for each table being archived, and it uses the Oracle Call Interface (OCI) to run the query. The query results are translated into XML. The XML output is passed to the NetBackup client interfaces.

`bpوراimp` uses the reverse process. That is, XML data is restored, parsed, and inserted into the database.

A failure here is probably due to an Oracle error, or to a NetBackup media, network, or time-out error.

- 5 `bpوراexp` or `bpوراimp` tells the NetBackup client to close the session and disconnects from the Oracle database.

The NetBackup client waits for the server to complete its necessary actions (backup image verification and so on) and then exits.

See [“About troubleshooting NetBackup for Oracle”](#) on page 225.

See [“Minimizing time-out failures on large database restores”](#) on page 237.

See [“Troubleshooting each stage of the backup or restore”](#) on page 234.

See [“About NetBackup for Oracle troubleshooting steps”](#) on page 225.

Additional XML export and import logs

The `bpوراexp` and `bpوراimp` utilities perform error logging and tracing in the file that is specified by the `LOG` parameter. The log files contain Oracle errors and other errors that are not related to NetBackup.

When `bpوراexp` and `bpوراimp` are used and the backup images are written to a storage unit, these errors are also logged in the NetBackup debug logs. These logs appear in the following directories:

Windows:

```
install_path\NetBackup\logs\bpوراexp  
install_path\NetBackup\logs\bpوراimp
```

UNIX and Linux:

```
/user/opensv/netbackup/logs/bporaexp  
/user/opensv/netbackup/logs/bporaimp
```

When you use `bpوراexp` and `bpوراimp` and the backup images are written to an operating system directory, the file that is specified by the `LOG=` parameter is the only source of error logging and tracing.

Register authorized locations

This appendix includes the following topics:

- [Registering authorized locations used by a NetBackup database script-based policy](#)

Registering authorized locations used by a NetBackup database script-based policy

During a backup, NetBackup checks for scripts in the default script location and any authorized locations. The default, authorized script location for UNIX is `usr/opensv/netbackup/ext/db_ext` and for Windows is `install_path\netbackup\dbext`. If the script is not in the default script location or an authorized location, the policy job fails. You can move any script into the default script location or any additional authorized location and NetBackup recognizes the scripts. You need to update the policy with the script location if it has changed. An authorized location can be a directory and NetBackup recognizes any script within that directory. An authorized location can also be a full path to a script if an entire directory does need to be authorized.

If the default script location does not work for your environment, use the following procedure to enter one or more authorized locations for your scripts. Use `nbsetconfig` to enter an authorized location where the scripts reside. You can also use `bpsetconfig`, however this command is only available on the primary or the media server.

Note: One recommendation is that scripts should not be world-writable. NetBackup does not allow scripts to run from network or remote locations. All scripts must be stored and run locally. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

For more information about registering authorized locations and scripts, review the knowledge base article:

https://www.veritas.com/content/support/en_US/article.100039639

To add an authorized location

- 1 Open a command prompt on the client.
- 2 Use `nbsetconfig` to enter values for an authorized location. The client privileged user must run these commands.

The following examples are for paths you may configure for the Oracle agent. Use the path that is appropriate for your agent.

- On UNIX:

```
[root@client26 bin]# ./nbsetconfig
nbsetconfig>DB_SCRIPT_PATH = /Oracle/scripts
nbsetconfig>DB_SCRIPT_PATH = /db/Oracle/scripts/full_backup.sh
nbsetconfig>
<ctrl-D>
```

- On Windows:

```
C:\Program Files\Veritas\NetBackup\bin>nbsetconfig
nbsetconfig> DB_SCRIPT_PATH=c:\db_scripts
nbsetconfig> DB_SCRIPT_PATH=e:\oracle\fullbackup\full_rman.sh
nbsetconfig>
<ctrl-Z>
```

Note: Review the [NetBackup Command Reference Guide](#) for options, such as reading from a text file and remotely setting clients from a NetBackup server using `bpsetconfig`. If you have a text file with the script location or authorized locations listed, `nbsetconfig` or `bpsetconfig` can read from that text file. An entry of `DB_SCRIPT_PATH=none` does not allow any script to run on a client. The `none` entry is useful if an administrator wants to completely lock down a server from running scripts.

Registering authorized locations used by a NetBackup database script-based policy

- 3** (Conditional) Perform these steps on any clustered database or agent node that can perform the backup.
- 4** (Conditional) Update any policy if the script location was changed to the default or authorized location.