



# ORACLE BACKUP & RECOVERY GUIDE

How to backup and restore an Oracle database

Alan Clark  
[alan.clark@nextw.com](mailto:alan.clark@nextw.com)



## ORACLE BACKUP & RECOVERY GUIDE

*Lessons from the Front Lines of Implementation and Integration*

By Alan Clark





## Forward

*My past and present team members who are in the  
trenches daily making it all work.*



## Revision Log

Revision	Date	Change Description
1	11/1/2025	Created original revision
2	11/13/2025	Added full Linux RMAN script and shell command example in Appendix G.



## Contents

Forward .....	2
Revision Log .....	3
Oracle Database Recovery Guide .....	14
PART I – FOUNDATIONS OF BACKUP AND RECOVERY .....	15
Chapter 1 – Introduction and Recovery Principles .....	15
1.1 Purpose of This Guide .....	15
1.2 Scope and Supported Platforms .....	15
1.3 Intended Audience and Prerequisites .....	16
1.4 The Role of the DBA in Recovery .....	16
1.5 Recovery Objectives: RTO and RPO .....	17
1.6 Types of Failures in Oracle Databases .....	17
1.7 Key Components of Oracle Recovery .....	18
1.8 Data Integrity and Backup Validation .....	18
1.9 Overview of the Guide’s Structure .....	19
1.10 Summary .....	19
Chapter 2 – Oracle Backup and Recovery Concepts .....	21
2.1 Purpose and Objectives .....	21
2.2 Instance Recovery vs. Media Recovery .....	21
2.3 Complete and Incomplete Recovery .....	22
2.4 ARCHIVELOG and NOARCHIVELOG Modes .....	22
2.5 Physical vs. Logical Backups .....	23
2.6 Flashback Technologies Overview .....	24
2.7 Recovery Catalogs and Metadata Management .....	24
2.8 Retention Policies and Backup Obsolescence .....	25
2.9 Applying Little’s Law to Backup Throughput .....	25
2.10 Backup Validation and Integrity Checks .....	26
2.11 Best Practices for Backup and Recovery Planning .....	26



2.12 Summary .....	27
<b>PART II – RMAN BACKUP OPERATIONS .....</b>	<b>28</b>
Chapter 3 – RMAN Configuration and Preparation .....	28
3.1 Overview of RMAN .....	28
3.2 RMAN Architecture and Components .....	28
3.3 Connecting to Target, Catalog, and Auxiliary Databases .....	29
3.4 Configuring Persistent RMAN Settings .....	30
3.5 Setting Up the Fast Recovery Area (FRA) .....	30
3.6 Configuring Channels and Parallelism .....	31
3.7 Control-File Autobackup and Naming Conventions .....	31
3.8 Using and Maintaining the Recovery Catalog .....	32
3.9 Backup Validation and Crosscheck .....	33
3.10 RMAN Environment Files and Logging .....	33
3.11 Security – Backup Encryption and Access Control .....	34
3.12 Verification Checklist .....	34
3.13 Summary .....	35
Chapter 4 – Performing Backups .....	36
4.1 Introduction .....	36
4.2 Types of RMAN Backups .....	36
4.3 Full Database Backup .....	37
4.4 Incremental and Cumulative Backups .....	38
4.5 Archived Redo Log Backups .....	38
4.6 Control File and SPFILE Backups .....	39
4.7 Compressed and Encrypted Backups .....	40
4.8 Incrementally Updated Image Copies .....	40
4.9 Backup of Specific Tablespaces or Datafiles .....	41
4.10 Parallel Backups and Performance Tuning .....	41
4.11 Validating and Reporting Backups .....	42
4.12 Scheduling and Automation .....	42



4.13 Post-Backup Validation and Testing .....	43
4.14 Monitoring and Housekeeping .....	44
4.15 Summary .....	44
<b>PART III – RMAN RESTORE AND RECOVERY PROCEDURES .....</b>	<b>45</b>
Chapter 5 – Instance and Database Recovery.....	45
5.1 Introduction .....	45
5.2 Understanding Recovery Workflow .....	45
5.3 Instance Recovery .....	46
5.4 Determining the Type of Recovery Required.....	46
5.5 Complete Database Recovery .....	47
5.6 Incomplete Database Recovery.....	48
5.7 Control-File Recovery During Database Recovery .....	49
5.8 Datafile Media Recovery .....	49
5.9 Block Media Recovery.....	50
5.10 Using Datafile Copies for Fast Recovery .....	51
5.11 Recovery Using Archived Logs Only.....	51
5.12 Recovering a Dropped Tablespace .....	51
5.13 Testing Recovery Scenarios .....	52
5.14 Post-Recovery Verification .....	52
5.15 Preventing Future Recovery Events .....	52
5.16 Summary .....	53
Chapter 6 – Control File and SPFILE Recovery .....	54
6.1 Introduction .....	54
6.2 Understanding the Control File .....	54
6.3 Detecting Control File Failure .....	55
6.4 Restoring Control Files from Autobackup .....	55
6.5 Restoring Control Files from a Known Backup Piece.....	56
6.6 Recreating the Control File Manually.....	56
6.7 Multiplexing Control Files.....	57



6.8 Understanding the Server Parameter File (SPFILE) .....	58
6.9 Creating a Backup of the SPFILE .....	58
6.10 Restoring the SPFILE from RMAN .....	59
6.11 Recreating the SPFILE Manually.....	59
6.12 Protecting the Control File and SPFILE .....	60
6.13 Post-Recovery Validation .....	60
6.14 Troubleshooting Common Errors .....	61
6.15 Recommended RMAN CONFIGURE Settings .....	61
6.16 Nightly Backup Job (Windows to Disk).....	62
6.17 Summary .....	63
Chapter 7 – Tablespace and Table-Level Recovery .....	64
7.1 Introduction .....	64
7.2 Understanding Tablespace Recovery .....	64
7.3 Offline Tablespace Recovery .....	64
7.4 Recovering a Dropped Tablespace .....	65
7.5 Partial Datafile Recovery within a Tablespace .....	66
7.6 Tablespace Point-in-Time Recovery (TSPITR).....	66
7.7 Recovering a Single Table Using RMAN.....	68
7.8 Recovering Tables Across Schemas .....	69
7.9 Recovering Tablespace or Table to a Different Database .....	69
7.10 Validating Tablespace and Table Recovery.....	70
7.11 Common Restrictions and Considerations .....	70
7.12 Example: End-to-End Table Recovery .....	70
7.13 Performance Optimization During Recovery .....	71
7.14 Summary .....	72
PART IV – FLASHBACK AND ADVANCED RECOVERY.....	73
Chapter 8 – Flashback Technologies .....	73
8.1 Introduction .....	73
8.2 Flashback Architecture Overview .....	73



8.3 Enabling Flashback Features.....	74
8.4 Types of Flashback Operations.....	75
8.5 Flashback Query .....	75
8.6 Flashback Table .....	76
8.7 Flashback Drop (Recycle Bin) .....	77
8.8 Flashback Transaction Query and Backout.....	77
8.9 Flashback Database.....	78
8.10 Managing Restore Points.....	79
8.11 Verifying Flashback Success .....	80
8.12 Troubleshooting Flashback Operations .....	80
8.13 Best Practices for Flashback Usage .....	81
8.14 Summary .....	81
Chapter 9 – Advanced Recovery and Duplication .....	82
9.1 Introduction .....	82
9.2 Duplication Methods Overview.....	82
9.3 Preparing for Duplication .....	83
9.4 Active Database Duplication .....	84
9.5 Backup-Based Duplication .....	85
9.6 Duplicating for Testing or QA .....	86
9.7 Creating Test Databases from Backups .....	87
9.8 Encrypted and Compressed Backups in Duplication .....	87
9.9 Multitenant Database Duplication (CDB/PDB) .....	88
9.10 Cross-Platform Duplication and Conversion.....	88
9.11 Converting Backups into Data Pump Exports .....	89
9.12 Verifying and Securing Duplicated Databases.....	89
9.13 Common Duplication Errors .....	90
9.14 Best Practices for Advanced Recovery .....	90
9.15 Summary .....	91
PART V – EXPERT TECHNIQUES AND TROUBLESHOOTING .....	92



Chapter 10 – Recovery Troubleshooting and Diagnostics.....	92
10.1 Introduction .....	92
10.2 Common Causes of Recovery Failure .....	92
10.3 Understanding RMAN Log Files.....	93
10.4 Using the Database Alert Log.....	94
10.5 Diagnostic Data Viewer (ADRCI) .....	94
10.6 Troubleshooting Backup Failures .....	95
10.7 Troubleshooting Restore and Recovery Failures .....	95
10.8 Troubleshooting Flashback Operations .....	97
10.9 Using Dynamic Performance Views.....	97
10.10 Data Recovery Advisor (LIST, ADVISE, REPAIR) .....	98
10.11 Interpreting Common Error Codes .....	98
10.12 Verification and Post-Recovery Checks .....	99
10.13 Automation and Logging Standards.....	100
10.14 Escalation and Support Readiness.....	100
10.15 Best Practices for Continuous Improvement.....	101
10.16 Summary .....	101
Chapter 11 – Performance, Optimization, and Validation.....	103
11.1 Introduction .....	103
11.2 Key Performance Concepts .....	103
11.3 Factors Influencing Backup Performance .....	104
11.4 Measuring RMAN Performance .....	104
11.5 Parallelism and Channel Optimization .....	105
11.6 Compression and Deduplication .....	106
11.7 Incremental Merge Optimization.....	107
11.8 Block Change Tracking.....	107
11.9 Restore and Recovery Performance .....	108
11.10 Monitoring and Analyzing Backup History .....	108
11.11 Validating Backups .....	109



11.12 Periodic Test Restores .....	110
11.13 Proactive Recovery Readiness Tests.....	110
11.14 Retention Policy Optimization.....	111
11.15 Automating Performance Reports .....	111
11.16 Best Practices for Optimal Backup and Recovery Performance .....	112
11.17 Summary .....	112
<b>Chapter 12 – Periodic Testing and Lab Drills .....</b>	<b>114</b>
12.1 Introduction .....	114
12.2 Objectives of Recovery Testing .....	114
12.3 Recommended Testing Frequency .....	114
12.4 Preparing a Test Environment .....	115
12.5 Categories of Recovery Tests .....	116
12.6 Simulating Common Failure Scenarios .....	116
12.7 Measuring Recovery Performance.....	117
12.8 Establishing a Test Procedure Template.....	118
12.9 Testing Flashback Operations.....	119
12.10 Recovery Validation with Auxiliary Instances.....	119
12.11 Recovery Checklist and Audit Log .....	120
12.12 Post-Test Review.....	121
12.13 Reporting and Continuous Improvement .....	121
12.14 Integrating Recovery Testing into Change Management.....	121
12.15 Example Lab Exercise: Full Database Recovery Simulation .....	122
12.16 Common Mistakes During Recovery Testing.....	122
12.17 Continuous Learning and Team Preparedness .....	123
12.18 Summary .....	123
<b>PART VI – APPENDICES .....</b>	<b>125</b>
Appendix A – RMAN Command Reference.....	125
A.1 Introduction.....	125
A.2 Connection Commands .....	125



A.3 Configuration Commands .....	125
A.4 Backup Commands .....	126
A.5 Restore and Recovery Commands .....	126
A.6 Maintenance and Validation.....	127
A.7 Duplication and Cloning .....	127
A.8 Flashback Commands .....	127
A.9 Diagnostic Commands.....	128
A.10 Notes .....	128
Appendix B – Recovery Catalog Views and Metadata.....	129
B.1 Overview .....	129
B.2 Common Catalog Views .....	129
B.3 Example Metadata Queries .....	129
B.4 Catalog Maintenance .....	130
Appendix C – Automation and Script Samples .....	131
C.1 Purpose .....	131
C.2 Daily Backup Script (Linux) .....	131
C.3 Weekly Validation Script.....	132
C.4 Recovery Drill Script.....	132
C.5 Automated Log Rotation.....	132
C.6 Windows RMAN Batch Example .....	133
C.7 Scheduled Job Verification .....	133
Appendix D – Troubleshooting Reference .....	134
D.1 Purpose.....	134
D.2 Common RMAN Errors .....	134
D.3 Common ORA Errors .....	134
D.4 Troubleshooting Sequence .....	135
Appendix E – Lab Scenarios .....	136
E.1 Introduction.....	136
E.2 Scenario 1 – Control File Loss.....	136



E.3 Scenario 2 – Datafile Corruption.....	136
E.4 Scenario 3 – Redo Log Loss .....	136
E.5 Scenario 4 – Tablespace Point-in-Time Recovery.....	137
E.6 Scenario 5 – Flashback and Restore Point Test .....	137
E.7 Scenario 6 – RMAN Validation Drill.....	137
E.8 Documentation.....	138
<b>Appendix F – Version Compatibility and Feature Map .....</b>	<b>139</b>
F.1 Purpose .....	139
F.2 Feature Evolution Summary.....	139
F.3 Deprecated or Replaced Features .....	140
F.4 Operating System Considerations .....	140
F.5 Version-Specific Enhancements .....	140
F.6 Summary .....	140
<b>Appendix G – Linux RMAN Backup Script and Shell Example.....</b>	<b>142</b>
1. Create the RMAN Backup Script: .....	142
2. Make the Script Executable:.....	143
3. Schedule with Cron: .....	143
Important Considerations: .....	143
<b>Appendix H – RMAN Backup Script Example for Veeam on a Linux Operating System .....</b>	<b>145</b>
Default RMAN Settings .....	145
RMAN Veeam Script Example .....	145
<b>PART VII – QUICK REFERENCE.....</b>	<b>147</b>
Quick Reference Overview .....	147
RMAN Emergency Command Blocks .....	147
A. Full Database Restore and Recovery .....	147
B. Incomplete Recovery to a Specific Time or SCN .....	148
C. Restore Control File from Autobackup .....	149
D. Restore a Single Datafile .....	149
E. Restore and Recover Specific Tablespace .....	150



F. Restore SPFILE from Backup .....	150
G. Block Media Recovery .....	151
H. Recreate Control File Manually (No Backup Available).....	151
RESTORE DATABASE and RECOVER DATABASE Templates .....	152
Template 1 – Standard Full Restore and Recovery.....	152
Template 2 – Full Restore with Validation and Log.....	152
Template 3 – Point-in-Time (Incomplete) Recovery .....	153
Template 4 – Tablespace Recovery .....	153
Template 5 – Datafile-Level Restore.....	153
Template 6 – Control File Restore and Database Recovery .....	153
Template 7 – Validation Only (No Restore).....	154
Flashback and Table Recovery Quick Commands.....	155
A. Flashback Table .....	155
B. Flashback Table to Before Drop .....	155
C. Flashback Query.....	155
D. Flashback Transaction Backout .....	155
E. Flashback Database.....	156
F. Drop Restore Point After Success .....	156
Emergency Recovery Decision Matrix .....	157
Emergency Recovery Notes .....	158



# Oracle Database Recovery Guide

**Scope:** On-premise Oracle Databases (12c–21c)

**Editions:** Standard and Enterprise

**Operating Systems:** Windows and Linux

**Storage:** Disk / shared network storage (no ASM, no cloud)

**Approach:** Manual and procedural (no automation by default)

**Structure:** Beginner → Expert progression, with flashback, shortcut code blocks, and testing labs



---

## PART I – FOUNDATIONS OF BACKUP AND RECOVERY

---

### Chapter 1 – Introduction and Recovery Principles

---

#### 1.1 Purpose of This Guide

The purpose of this guide is to provide database administrators with a practical, procedural reference for performing **backup, restore, and recovery operations** on on-premise Oracle Database environments.

This guide focuses on **Oracle 12c through 21c**, supporting both **Standard** and **Enterprise Editions**, and assumes the use of **disk-based** or **network-shared** storage rather than Oracle ASM, cloud, or third-party backup appliances.

The design of this guide emphasizes **hands-on recovery** rather than theory. Each procedure is written to reduce decision-making time during a critical event, minimize data loss, and restore service availability as quickly as possible.

Where appropriate, examples and quick-reference command blocks are provided for RMAN, SQL\*Plus, and operating-system utilities.

---

#### 1.2 Scope and Supported Platforms

This document applies to:

- Oracle Database 12c Release 1 (12.1) through 21c
- Operating systems: Microsoft Windows Server and Linux (RHEL, OL, SLES, Ubuntu LTS)
- Storage: local disks or network file shares mounted to the host
- Database architectures: non-ASM single-instance, PDB/CDB multitenant
- Backup medium: RMAN disk backups, control file autobackups, archivelogs, and flat-file parameter backups

Excluded from this guide:



- ASM, Exadata, or ZFS Storage Appliance integration
- Oracle Cloud Infrastructure (OCI) backup modules
- Third-party tape or snapshot solutions
- Data Guard and RAC disaster recovery automation

---

## 1.3 Intended Audience and Prerequisites

This guide is intended for **database administrators, DBA trainees, and system engineers** responsible for maintaining the recoverability of Oracle databases.

Readers should already understand:

- Relational-database fundamentals
- Oracle instance architecture (SGA, PGA, control files, redo logs, datafiles)
- Basic SQL and command-line operation under Windows or Linux
- Oracle Environment Variables (ORACLE\_HOME, ORACLE\_SID)
- Concepts of **archive logging** and **redo generation**

Those new to RMAN will benefit from reviewing *Oracle Database Backup and Recovery User's Guide* for their specific release before using the expert procedures later in this guide.

---

## 1.4 The Role of the DBA in Recovery

A database administrator must ensure that the **backup strategy** aligns with the organization's **business continuity objectives**—specifically the *Recovery Time Objective (RTO)* and *Recovery Point Objective (RPO)*.

The DBA's role during a recovery event includes:

1. **Assessing the failure.** Determine which components are affected (control file, redo log, datafile, or instance).
2. **Selecting the appropriate recovery type.** Choose complete, incomplete, flashback, or point-in-time recovery based on data availability and RPO.
3. **Executing the recovery procedure with verified backups.**
4. **Validating the restored database and returning it to service.**



## 5. Documenting the incident and revising backup policies to prevent recurrence.

A prepared DBA treats recovery as a **routine procedure**, not a crisis response.

Testing and documentation ensure that recovery steps can be followed by any trained team member under pressure.

## 1.5 Recovery Objectives: RTO and RPO

Recovery planning revolves around two measurable goals:

Objective	Definition	Example
<b>RTO – Recovery Time Objective</b>	The maximum acceptable downtime before the database must be restored to service.	“Critical applications must resume within 2 hours of a major failure.”
<b>RPO – Recovery Point Objective</b>	The maximum amount of data loss, measured in time between the last valid backup and the failure.	“Data loss must not exceed 15 minutes of transactions.”

These objectives determine backup frequency, archivelog configuration, and redo transport strategy.

Meeting both targets requires continuous monitoring of backup health and recovery performance.

## 1.6 Types of Failures in Oracle Databases

Oracle databases may fail for numerous reasons. Each type requires a specific recovery response:

1. **Instance Failure** – A crash or power loss causing an abnormal shutdown. Oracle performs automatic instance recovery upon startup by applying redo logs.
2. **Media Failure** – Damage to a datafile, control file, or redo log on disk. Requires RMAN restore and media recovery.
3. **User Error** – Accidental table deletion or data modification. Handled via flashback or point-in-time recovery.
4. **Logical Corruption** – Block corruption detected during read operations. Resolved by block media recovery or table export/import.



5. **Environmental Failure** – Operating-system or hardware failure affecting database availability.
6. **Catastrophic Failure** – Loss of entire database or host. Requires full restore from validated backups.

Each subsequent chapter in this guide provides detailed, tested procedures for responding to these failure types.

---

## 1.7 Key Components of Oracle Recovery

Effective recovery depends on understanding how Oracle stores and protects data:

Component	Function	Recovery Relevance
<b>Control File</b>	Tracks physical database structure and backup metadata.	Essential for RMAN operations; must be multiplexed and backed up frequently.
<b>Redo Logs</b>	Record all changes made to the database.	Required for instance and media recovery; archived for point-in-time restores.
<b>Undo Tablespace</b>	Maintains before-image data for transactions.	Supports read consistency and flashback operations.
<b>Datafiles</b>	Store actual table and index segments.	Primary target for backup and restore operations.
<b>Parameter Files (PFILE/SPFILE)</b>	Define instance configuration and memory allocation.	Required to recreate or start instances during disaster recovery.

Regular validation of each component ensures that a recovery can be executed without delay.

---

## 1.8 Data Integrity and Backup Validation

A backup has no value until it is **tested and verified**. Oracle provides several tools to confirm recoverability:

1. **RMAN VALIDATE** – Verifies the physical integrity of backup sets and datafiles.



2. **CROSSCHECK** – Confirms that backups listed in the control file or catalog exist on disk.
3. **REPORT OBSOLETE** – Identifies backups eligible for deletion based on the retention policy.
4. **LIST BACKUP SUMMARY** – Generates a quick inventory of available backups.

Periodic restore tests using **auxiliary instances** validate the entire recovery path and should be performed quarterly or after any infrastructure change.

---

## 1.9 Overview of the Guide's Structure

This guide is organized into progressive sections:

Part	Focus	Description
<b>I – Foundations</b>	Concepts and Planning	Recovery principles, terminology, and architecture.
<b>II – Backup Operations</b>	RMAN Configuration and Execution	How to configure and perform reliable disk backups.
<b>III – Restore &amp; Recovery</b>	Practical Recovery Scenarios	Procedures for restoring instances, control files, and tablespaces.
<b>IV – Flashback and Advanced Recovery</b>	Modern Recovery Methods	Flashback database, point-in-time recovery, and duplication.
<b>V – Expert Techniques and Troubleshooting</b>	Optimization and Diagnostics	Performance tuning and error resolution.
<b>VI – Appendices</b>	Reference and Lab Exercises	Commands, catalog views, and testing procedures.
<b>VII – Quick Reference and Bibliography</b>	Rapid Recovery Checklist	Shortcut commands and source citations.

---

## 1.10 Summary

Chapter 1 introduced the purpose, scope, and guiding principles for Oracle database recovery. The key points are:

- Recovery must be **planned, tested, and repeatable**.



- DBAs must understand each **component's role in recovery**.
- RMAN and Flashback provide complementary approaches to data protection.
- **Backup validation and testing** are essential to true recoverability.

Subsequent chapters build upon this foundation, moving from configuration of RMAN backups to advanced flashback and diagnostic techniques.



---

## Chapter 2 – Oracle Backup and Recovery Concepts

---

### 2.1 Purpose and Objectives

The goal of this chapter is to provide a working understanding of how Oracle performs backup and recovery internally.

Before running any RMAN command or recovery script, a DBA must understand *what* is being protected, *when* data becomes recoverable, and *how* the database engine coordinates redo and checkpoint information during restoration.

By the end of this chapter, you should be able to:

- Distinguish between **instance**, **crash**, and **media** recovery.
- Identify when to perform **complete** or **incomplete** recovery.
- Understand the difference between **physical** and **logical** backups.
- Describe how **ARCHIVELOG mode** enables point-in-time recovery.
- Recognize when to use **flashback** features instead of full recovery.
- Explain the function of the **recovery catalog** and **retention policies**.
- Apply **Little's Law** to estimate throughput and backup windows.

---

### 2.2 Instance Recovery vs. Media Recovery

Oracle distinguishes two fundamental recovery operations:

#### 2.2.1 Instance Recovery

Occurs automatically after an abnormal shutdown or system crash.

When the instance starts, Oracle uses the **redo log files** to roll forward all committed changes that were not written to datafiles and then rolls back uncommitted transactions using undo segments.

#### Key points

- Performed automatically by SMON during startup.



- Requires only the **online redo logs**—no backup restore needed.
- Duration depends on the amount of uncheckpointed data.

### 2.2.2 Media Recovery

Restores lost or corrupted physical files such as datafiles, control files, or archived logs. RMAN restores copies from backup and then applies archived redo logs to make the data consistent.

#### Typical cases

- Disk failure or accidental deletion of a datafile.
- Block corruption detected by DBVERIFY or RMAN VALIDATE.
- Control-file loss requiring RESTORE CONTROLFILE.

Media recovery may be **complete** (recover to the latest SCN) or **incomplete** (recover to an earlier SCN, time, or log sequence).

## 2.3 Complete and Incomplete Recovery

Type	Description	When Used
<b>Complete Recovery</b>	Restores all required files and applies all redo to bring the database to the most recent committed transaction.	After hardware failure or datafile loss when all redo logs and archives are available.
<b>Incomplete Recovery</b>	Recovers the database to a specific <b>SCN</b> , <b>timestamp</b> , or <b>log sequence</b> , then opens with RESETLOGS.	When archived redo after the target point is missing, or to reverse user error or data corruption.

Incomplete recovery always produces a **new incarnation** of the database. The previous incarnation's backups remain usable but must be referenced through the recovery catalog if further recovery is required.

## 2.4 ARCHIVELOG and NOARCHIVELOG Modes

Oracle can operate in one of two redo management modes:



## ARCHIVELOG Mode

- Each filled online redo log is archived to a designated directory or Fast Recovery Area (FRA).
- Enables point-in-time recovery using archived redo.
- Recommended for all production and test environments.

SQL> ALTER DATABASE ARCHIVELOG;

SQL> ARCHIVE LOG LIST;

## NOARCHIVELOG Mode

- Redo logs are overwritten in a circular fashion.
- Permits only **cold backups** (database must be closed).
- Suitable only for development systems where data loss is acceptable.

Switching between modes requires a **mount** state and a full backup afterward to preserve recoverability.

## 2.5 Physical vs. Logical Backups

Backup Type	Description	Tool	Recovery Use
Physical Backup	Copy of database files—datafiles, control files, redo, and parameter files.	RMAN, OS copy, or third-party snapshot.	Used for complete or incomplete recovery.
Logical Backup	Extract of database objects as SQL or Data Pump dump files.	Data Pump (expdp/impdp)	Used for table-level or cross-platform migration.

Physical backups guarantee transactional consistency; logical backups provide portability but require import utilities. Both should be part of a layered recovery strategy.



## 2.6 Flashback Technologies Overview

Flashback features supplement RMAN by allowing *rewind* operations inside the same database instance.

Feature	Description	Typical Use
<b>Flashback Query</b>	View past data using AS OF TIMESTAMP or SCN.	Investigate historical row values.
<b>Flashback Table</b>	Restore one or more tables to a prior state.	Undo accidental deletes or updates.
<b>Flashback Drop</b>	Retrieve dropped objects from the Recycle Bin.	Recover tables dropped by mistake.
<b>Flashback Database</b>	Roll back the entire database to a restore point or SCN.	Rapid recovery from logical corruption.
<b>Guaranteed Restore Point</b>	Preserves redo for flashback until explicitly dropped.	Pre-upgrade or major-change safety net.

Flashback operations require **UNDO tablespace**, **flashback logs**, and **ARCHIVELOG mode**. They are most effective when changes need reversal without restoring backups.

---

## 2.7 Recovery Catalogs and Metadata Management

RMAN maintains its metadata in two locations:

1. **Control File Repository** – Default; contains backup history and configuration.
2. **Recovery Catalog Database** – Optional schema in a separate database that stores historical metadata beyond the control-file retention period.

### Advantages of a Recovery Catalog

- Centralized management of multiple target databases.
- Retains records even after control-file reuse.
- Enables stored scripts and reporting queries (RC\_\* views).



## Key Views

- RC\_BACKUP\_SET, RC\_BACKUP\_PIECE, RC\_DATABASE, RC\_RESTORE\_POINT.
- Accessed via SQL plus or data dictionary tools for auditing backup history.

For critical environments, maintain the recovery catalog on independent storage with its own backup policy.

---

## 2.8 Retention Policies and Backup Obsolescence

A retention policy defines which backups are preserved and which may be deleted.

RMAN uses either:

- **Recovery Window Policy** (default) – keeps all backups needed to recover to a point  $n$  days in the past.
- **Redundancy Policy** – retains a fixed number of backup copies.

RMAN> CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;

RMAN> REPORT OBSOLETE;

RMAN> DELETE NOPROMPT OBSOLETE;

Choosing an appropriate policy balances storage cost with recovery flexibility.

For systems with daily backups, a seven-day window ensures one week of restore capability.

---

## 2.9 Applying Little's Law to Backup Throughput

**Little's Law** (from queuing theory) expresses the relationship between average throughput, concurrency, and service time:

$$L = \lambda \times W$$

Where:

- $L$  = Average number of tasks in the system (e.g., active channels).
- $\lambda$  = Average throughput (MB/sec).
- $W$  = Average response time (seconds per operation).

For RMAN, this translates to:



$$\text{Effective Throughput} = \frac{\text{Data Volume (MB)}}{\text{Backup Duration (s)}}$$

### Example

If a 500 GB database backup completes in 25 minutes (1,500 s), throughput = 333 MB/s.

With four RMAN channels, each channel processes about 83 MB/s.

Using these figures, you can estimate the number of channels or disks needed to meet RTO/RPO targets.

## 2.10 Backup Validation and Integrity Checks

Recovery depends on **valid, readable backups**. Oracle provides multiple validation methods:

1. **RMAN VALIDATE DATABASE** – Scans all data blocks in backups or live datafiles.
2. **RESTORE ... VALIDATE** – Simulates a restore without writing files.
3. **DBVERIFY utility** – Checks physical block structure of datafiles.
4. **Checksum comparison** – Ensures redo and data consistency.

Validation should be scheduled weekly and logged to a central repository.

Invalid or corrupt backups must be replaced immediately.

## 2.11 Best Practices for Backup and Recovery Planning

- Enable **ARCHIVELOG mode** on all production databases.
- Use **multiplexed control files** on separate disks.
- Place archived logs and backups on a different volume from datafiles.
- Automate **RMAN CROSSCHECK** and **REPORT OBSOLETE** tasks.
- Test **restore and recovery** quarterly using an auxiliary instance.
- Maintain **backup logs** and RMAN output for audit compliance.
- Document every recovery test, including elapsed time and bottlenecks.



## 2.12 Summary

Backup and recovery are not single actions but an integrated **data-protection lifecycle**:

1. Configure a consistent **backup policy** aligned with business RTO/RPO.
2. Validate backups to ensure physical integrity.
3. Understand the relationship between redo, checkpoints, and archived logs.
4. Apply RMAN and Flashback techniques as complementary tools.

This conceptual foundation prepares you for **Chapter 3 – RMAN Configuration and Preparation**, where the theory becomes actionable procedures using real RMAN commands.



---

## PART II – RMAN BACKUP OPERATIONS

---

### Chapter 3 – RMAN Configuration and Preparation

---

#### 3.1 Overview of RMAN

Recovery Manager (RMAN) is Oracle's integrated utility for creating, managing, and validating backups.

Unlike manual file-copy methods, RMAN tracks every backup through metadata in the control file or an optional recovery catalog, ensuring that restoration can occur accurately and consistently.

RMAN communicates directly with the Oracle instance to read datafiles, produce backup sets, and apply redo during recovery.

It automates checksum validation, retention enforcement, and incremental updates, reducing manual intervention and error risk.

Key capabilities include:

- Block-level corruption detection during backup and restore.
- Automatic control-file and SPFILE backups.
- Incremental and cumulative backup optimization.
- Configurable retention policies and channel parallelism.
- Script automation through stored procedures or command files.

---

#### 3.2 RMAN Architecture and Components

RMAN consists of coordinated elements that perform specific roles during backup and recovery:



Component	Description
<b>RMAN Client</b>	Command-line interface that interprets commands and issues calls to the target database.
<b>Target Database</b>	The database being backed up or recovered. It runs server sessions that execute RMAN operations.
<b>Recovery Catalog</b>	Optional schema in a separate database that stores historical metadata, configuration settings, and stored scripts.
<b>Channels</b>	Server sessions that move data between the database and the storage medium. Each channel performs one stream of I/O.
<b>Fast Recovery Area (FRA)</b>	Disk location managed by Oracle to store archived redo logs, flashback logs, and RMAN backups.

For on-premise environments that use only disk storage, all channels are configured with DEVICE TYPE DISK.

---

### 3.3 Connecting to Target, Catalog, and Auxiliary Databases

RMAN connections can be local (OS authentication) or remote (password file authentication).

# Local connection

```
rman TARGET /
```

# Remote connection using TNS

```
rman TARGET sys/password@ORCL
```

# With a recovery catalog

```
rman TARGET sys/password@ORCL CATALOG rmanuser/rmanpw@CATDB
```

For operations that require a temporary database—such as duplication or tablespace point-in-time recovery—connect an auxiliary instance:



rman TARGET / AUXILIARY sys/password@AUXDB

Always verify ORACLE\_HOME and ORACLE\_SID before starting RMAN to ensure the session connects to the intended database.

---

## 3.4 Configuring Persistent RMAN Settings

RMAN stores configuration parameters in the control file (and optionally the catalog). These persistent settings remain in effect for all sessions until changed.

### Common configuration commands

CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;

CONFIGURE CONTROLFILE AUTOBACKUP ON;

CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/%U';

CONFIGURE DEVICE TYPE DISK PARALLELISM 4;

CONFIGURE ENCRYPTION FOR DATABASE ON;

SHOW ALL;

Persistent configuration ensures that every backup conforms to the organization's retention, naming, and security standards.

---

## 3.5 Setting Up the Fast Recovery Area (FRA)

The Fast Recovery Area is a unified storage location for archived redo, flashback logs, and backup pieces.

ALTER SYSTEM SET DB\_RECOVERY\_FILE\_DEST\_SIZE = 200G SCOPE=BOTH;

ALTER SYSTEM SET DB\_RECOVERY\_FILE\_DEST = '/u01/app/oracle/FRA' SCOPE=BOTH;

Guidelines:

- Size the FRA to hold at least one full backup and the archived redo required to meet the recovery window.
- Place the FRA on a different physical disk than datafiles.
- Monitor usage with V\$RECOVERY\_FILE\_DEST and V\$FLASH\_RECOVERY\_AREA\_USAGE.



- Schedule CROSSCHECK and DELETE OBSOLETE to reclaim space.

If the FRA becomes full, Oracle halts archiving until space is freed, so continuous monitoring is essential.

---

## 3.6 Configuring Channels and Parallelism

Channels determine the degree of parallelism in RMAN operations.

Each channel performs one stream of I/O between the database and the backup destination.

### Manual channel allocation

```
RUN {
  ALLOCATE CHANNEL c1 DEVICE TYPE DISK FORMAT '/backup/DB_%U.bkp';
  BACKUP DATABASE PLUS ARCHIVELOG;
  RELEASE CHANNEL c1;
}
```

### Automatic configuration

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 4 BACKUP TYPE TO BACKUPSET;
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/DF_%U.bkp';
```

Recommendations:

- Allocate one channel per physical disk or storage path.
- Increase gradually to determine optimal performance.
- Avoid oversubscription, which can saturate CPU or I/O bandwidth.

---

## 3.7 Control-File Autobackup and Naming Conventions

Control-file autobackup ensures that every backup operation produces a current copy of the control file and SPFILE.

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;
```

```
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'/backup/CTL_%F.bkp';
```



Best practices:

- Store at least two copies on separate volumes.
- Retain pre-upgrade control-file copies.
- Verify autobackups periodically with LIST BACKUP OF CONTROLFILE;.

The %F substitution variable automatically generates a unique name that includes the DBID and timestamp.

---

## 3.8 Using and Maintaining the Recovery Catalog

### Creating the Catalog

1. Create a dedicated user and tablespace in the catalog database:
2. CREATE USER rman IDENTIFIED BY rmanpw DEFAULT TABLESPACE rman\_ts QUOTA UNLIMITED ON rman\_ts;
3. GRANT RECOVERY\_CATALOG\_OWNER TO rman;
4. Connect through RMAN and build the catalog:
5. rman CATALOG rman/rmanpw@CATDB
6. RMAN> CREATE CATALOG;
7. Register the target database:
8. RMAN> CONNECT TARGET sys/password@ORCL;
9. RMAN> REGISTER DATABASE;

### Maintaining the Catalog

- Synchronize metadata with RESYNC CATALOG;.
- Upgrade schema after database upgrades with UPGRADE CATALOG;.
- Back up the catalog database as part of the enterprise backup plan.

A recovery catalog retains history far beyond the control-file record-keeping period, supporting audits and long-term retention.

---



## 3.9 Backup Validation and Crosscheck

Regular validation ensures that every backup recorded in RMAN metadata exists and is readable.

CROSSCHECK BACKUP;

REPORT OBSOLETE;

VALIDATE DATABASE;

LIST BACKUP SUMMARY;

Interpretation of status:

- **AVAILABLE** – file exists and is usable.
- **EXPIRED** – missing or deleted.
- **OBSOLETE** – superseded by newer backups under the retention policy.

Integrate validation into weekly maintenance or after any storage migration.

## 3.10 RMAN Environment Files and Logging

All RMAN operations should generate a timestamped log.

### Example

SPOOL LOG TO '/var/log/rman/backup\_&DATE..log';

BACKUP DATABASE PLUS ARCHIVELOG;

SPOOL LOG OFF;

Guidelines:

- Store logs in a secured location separate from backups.
- Rotate or compress logs older than the retention period.
- On Windows, send summary results to the Event Viewer; on Linux, forward them to syslog using the logger command.

Comprehensive logs allow performance analysis and provide audit evidence of backup completion.



## 3.11 Security – Backup Encryption and Access Control

Even when backups remain on-premise, encryption is recommended to protect against unauthorized access.

### Transparent Encryption

Uses an Oracle Wallet:

```
CONFIGURE ENCRYPTION ALGORITHM 'AES256';
```

```
CONFIGURE ENCRYPTION FOR DATABASE ON;
```

```
BACKUP DATABASE;
```

### Password-Protected Encryption

Useful when no wallet is configured:

```
BACKUP AS COMPRESSED BACKUPSET DATABASE
```

```
FORMAT '/backup/secure_%U.bkp'
```

```
ENCRYPTION PASSWORD = 'StrongPassword';
```

Apply strict OS permissions to backup directories and never store wallet files or password scripts alongside the encrypted backups.

## 3.12 Verification Checklist

Before scheduling production backups, confirm:

1. Database runs in ARCHIVELOG mode.
2. FRA exists and is writable.
3. RMAN configuration verified with SHOW ALL;.
4. Control-file autobackup is enabled.
5. A full backup completed and validated successfully.
6. RMAN logs stored and reviewed.
7. Scheduler or CRON jobs configured and tested.

Document these confirmations as part of the system build record and audit trail.



## 3.13 Summary

This chapter established the foundation for all subsequent recovery procedures.

You have:

- Defined a persistent RMAN configuration.
- Created and tested the Fast Recovery Area.
- Set up channels, parallelism, and control-file autobackups.
- Built and synchronized a recovery catalog.
- Verified backup integrity and implemented encryption where appropriate.

With these preparations complete, RMAN is now a fully operational backup infrastructure.

The next chapter demonstrates how to **perform backups** using these configurations to achieve both performance and recoverability objectives.



# Chapter 4 – Performing Backups

## 4.1 Introduction

After RMAN has been configured and tested, the next critical responsibility of the DBA is to design and execute a consistent backup strategy.

A well-defined backup routine guarantees that every datafile, control file, and archive log can be restored in the shortest possible time with minimal data loss.

This chapter describes the major backup types, scheduling methods, and verification procedures that ensure recoverability in on-premise Oracle databases using only disk or network-share storage.

All examples assume the database operates in **ARCHIVELOG** mode and that the Fast Recovery Area (FRA) is configured and accessible.

## 4.2 Types of RMAN Backups

RMAN supports multiple backup structures.

Understanding when to use each type allows the DBA to balance storage efficiency and recovery performance.

Backup Type	Description	Typical Use
<b>Full Backup</b>	Copies every data block from the specified datafiles into a backup set.	Initial baseline or periodic full refresh.
<b>Incremental Level 0</b>	Equivalent to a full backup but marks a new incremental baseline.	First step of an incremental strategy.
<b>Incremental Level 1</b>	Captures only blocks changed since the most recent Level 0 or 1 backup.	Daily differential backups.
<b>Cumulative Incremental</b>	Captures blocks changed since the last Level 0 backup only.	Simplifies restore sequence at the expense of size.



Backup Type	Description	Typical Use
<b>Image Copy</b>	Byte-for-byte copy of datafiles or control files.	Used for fast recovery or merge operations.
<b>Archived Redo Log Backup</b>	Archives all redo generated since last backup.	Required for point-in-time recovery.
<b>Control File and SPFILE Backup</b>	Protects database structure and parameter state.	Automatically included if autobackup is enabled.

## 4.3 Full Database Backup

A full database backup creates a consistent restore point for every datafile.

### Procedure

1. Confirm the database is open and operating in ARCHIVELOG mode.
2. Connect to RMAN and execute:

```
rman TARGET /
```

```
RUN {
```

```
  BACKUP DATABASE TAG 'FULL_DB_BACKUP';
```

```
}
```

3. Verify the output log for successful completion.
4. Review the backup summary:

```
LIST BACKUP SUMMARY;
```

5. Validate the backup set:

```
RESTORE DATABASE VALIDATE;
```

### Recommendations

- Run a full backup after any major structural change (tablespace creation, upgrade, parameter adjustment).



- Store at least two generations of full backups on separate volumes.
- Use meaningful **TAG** names to simplify recovery selection.

---

## 4.4 Incremental and Cumulative Backups

Incremental backups reduce time and storage by recording only changed blocks.

A typical weekly pattern is one Level 0 backup on Sunday and Level 1 incrementals Monday through Saturday.

### Differential Incremental Example

RUN {

```
BACKUP INCREMENTAL LEVEL 1 DATABASE TAG 'INCR_L1';
```

}

### Cumulative Incremental Example

RUN {

```
BACKUP INCREMENTAL LEVEL 1 CUMULATIVE DATABASE TAG 'INCR_CUM';
```

}

### Best Practices

- Maintain a consistent schedule so that each Level 1 references the proper Level 0 baseline.
- Validate baseline existence before running incremental jobs.
- Periodically refresh the baseline by taking a new Level 0 to prevent long incremental chains.
- When possible, enable **block change tracking** to improve incremental performance:

```
ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE '/u01/bct/bct.trk';
```

---

## 4.5 Archived Redo Log Backups

Archived logs are mandatory for media recovery and point-in-time restoration.

They should be backed up frequently and deleted only after the backup is verified.

### Procedure



```

RUN {
  BACKUP ARCHIVELOG ALL
  FORMAT '/backup/arch/AL_%U.bkp'
  DELETE INPUT
  TAG 'ARCHIVE_BACKUP';
}

```

### Guidelines

- Back up archived logs at least hourly for OLTP databases or every time redo usage approaches 80 % of available space.
- The DELETE INPUT clause removes logs only after they are successfully copied.
- To back up logs generated in the last day only:

```
BACKUP ARCHIVELOG FROM TIME 'SYSDATE-1';
```

- Monitor FRA usage carefully; failure to archive due to space exhaustion halts redo generation.

## 4.6 Control File and SPFILE Backups

Control files record the physical structure and backup history of the database.

The SPFILE stores parameter settings necessary to start the instance.

Although autobackup typically protects both, manual copies are recommended before upgrades or configuration changes.

### Manual Backup

```
BACKUP CURRENT CONTROLFILE FORMAT '/backup/CTL_%U.bkp';
```

```
BACKUP SPFILE FORMAT '/backup/SPFILE_%U.bkp';
```

### Best Practices

- Place control-file backups on multiple disks or network shares.
- After restoring an older control file, always run RECOVER DATABASE to synchronize metadata.



- Maintain text copies of parameter settings for reference:

```
CREATE PFILE='/backup/init_backup.ora' FROM SPFILE;
```

---

## 4.7 Compressed and Encrypted Backups

Compression and encryption optimize storage while protecting data.

### Compressed Backup

```
BACKUP AS COMPRESSED BACKUPSET DATABASE PLUS ARCHIVELOG TAG
'COMPRESSED_FULL';
```

Compression reduces backup size significantly on low-change data but increases CPU consumption; adjust parallelism accordingly.

### Encrypted Backup

If an Oracle Wallet is configured:

```
CONFIGURE ENCRYPTION ALGORITHM 'AES256';
```

```
BACKUP DATABASE PLUS ARCHIVELOG TAG 'ENCRYPTED_FULL';
```

For password-based encryption:

```
BACKUP DATABASE ENCRYPTION PASSWORD='StrongPassword' TAG 'PW_ENCRYPTED';
```

Always secure the wallet or password in accordance with the organization's data-protection policy.

---

## 4.8 Incrementally Updated Image Copies

An incrementally updated backup merges Level 1 changes into a static image copy, creating a rolling full backup.

### Configuration

1. Create the initial image copy:

```
BACKUP AS COPY DATABASE FORMAT '/backup/df_%U.bkp';
```

2. Apply incremental changes nightly:

```
BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG 'ROLLING_FULL' DATABASE;
```



This approach eliminates the need for periodic full backups and allows near-instant recovery by switching to the copy.

---

## 4.9 Backup of Specific Tablespaces or Datafiles

Selective backups reduce maintenance time for large databases.

BACKUP TABLESPACE users, hr;

BACKUP DATAFILE 5, 6 FORMAT '/backup/DF\_users\_%U.bkp';

Use this method for testing environments or to protect high-change tablespaces independently.

Always include the **SYSTEM** and **UNDO** tablespaces in full backups; partial sets alone are not sufficient for complete recovery.

---

## 4.10 Parallel Backups and Performance Tuning

RMAN can perform multiple I/O streams simultaneously by allocating several channels.

The number of channels should reflect available CPU, I/O bandwidth, and disk targets.

RUN {

ALLOCATE CHANNEL c1 DEVICE TYPE DISK FORMAT '/backup/disk1\_%U.bkp';

ALLOCATE CHANNEL c2 DEVICE TYPE DISK FORMAT '/backup/disk2\_%U.bkp';

BACKUP DATABASE PLUS ARCHIVELOG;

}

### Optimization Tips

- Separate data and log backups onto different disks.
- Use **multiple channels** rather than OS-level striping for flexibility.
- Enable compression only if CPU utilization is below 60 % during backup windows.
- Schedule backups during low-I/O periods to minimize contention.

---



## 4.11 Validating and Reporting Backups

Validation ensures that backups are complete and restorable.

In addition to routine monitoring, validation should occur after every structural change or patch application.

### Commands

VALIDATE DATABASE;

LIST BACKUP SUMMARY;

REPORT NEED BACKUP;

REPORT OBSOLETE;

Interpret results as follows:

- *VALIDATE OK* – all blocks verified successfully.
- *NEED BACKUP* – datafile requires a backup to meet retention policy.
- *OBSOLETE* – eligible for deletion under the current recovery window.

Regular validation ensures that the recovery plan remains executable without surprises.

## 4.12 Scheduling and Automation

Backups should be fully automated to ensure consistency and reliability.

### Example Shell Script (Linux)

```
#!/bin/bash

export ORACLE_HOME=/u01/app/oracle/product/19c/dbhome_1
export ORACLE_SID=PROD
rman TARGET / <<EOF
RUN {
  CROSSCHECK BACKUP;
  DELETE NOPROMPT EXPIRED BACKUP;
  BACKUP DATABASE PLUS ARCHIVELOG TAG 'DAILY_BACKUP';
}
```



DELETE NOPROMPT OBSOLETE;

}

EXIT;

EOF

Schedule using **cron**:

0 2 \* \* \* /u01/scripts/rman\_daily.sh > /u01/scripts/rman\_daily.log 2>&1

### Example Windows Batch File

set ORACLE\_SID=PROD

rman TARGET / CMDfile=rman\_daily.cmd LOG=rman\_daily.log

Document each scheduled task and verify successful execution through log inspection and email or alerting tools.

## 4.13 Post-Backup Validation and Testing

A backup's value is proven only through successful restoration.

Perform routine test restores to an auxiliary instance or sandbox environment.

### Verification Example

RESTORE DATABASE VALIDATE;

For a full functional test:

1. Create a temporary instance.
2. Restore control files and datafiles from backup.
3. Recover until the last archived redo.
4. Open with RESETLOGS.
5. Confirm application access.

Conduct recovery drills at least quarterly to verify that all procedures, storage paths, and logs function as expected.



## 4.14 Monitoring and Housekeeping

Regular monitoring ensures backup infrastructure stability.

### Key Tasks

- Review RMAN logs daily.
- Monitor FRA usage and growth trends.
- Check for failed or hung backup jobs.
- Verify timestamps of latest control-file autobackups.
- Clean up obsolete and expired files to reclaim space.

Automated monitoring scripts can query V\$BACKUP\_SET and V\$BACKUP\_ASYNC\_IO to track backup performance and identify slow devices.

---

## 4.15 Summary

In this chapter, you executed practical RMAN backup operations that form the backbone of database protection:

- Performed full, incremental, and cumulative backups.
- Protected archived redo, control files, and parameter files.
- Implemented compression, encryption, and incrementally updated copies.
- Tuned parallelism for optimal throughput.
- Validated, scheduled, and tested backups to guarantee recoverability.

With backups now secured and validated, the next chapter moves into **database restoration and recovery**—the real-world application of these backups to rebuild an Oracle database after failure.



---

## PART III – RMAN RESTORE AND RECOVERY PROCEDURES

---

### Chapter 5 – Instance and Database Recovery

---

#### 5.1 Introduction

Database recovery restores the consistency of an Oracle database after a failure.

A failure can be caused by hardware malfunction, file loss, human error, or operating system crash.

Recovery ensures that committed transactions are preserved and that uncommitted changes are rolled back.

There are two principal recovery scenarios:

1. **Instance recovery** – performed automatically by Oracle at startup after an abnormal termination.
2. **Media recovery** – performed manually using RMAN or SQL\*Plus when datafiles, control files, or redo logs are damaged or missing.

This chapter explains the procedures, requirements, and verification steps for both.

Each example assumes that RMAN configuration is complete, backups exist on disk, and the database operates in **ARCHIVELOG** mode.

---

#### 5.2 Understanding Recovery Workflow

Recovery always follows a sequence of steps:

1. **Mount or open the instance** to a state where recovery can begin.
2. **Restore** the necessary datafiles or control files from backup.
3. **Recover** by applying archived redo and online redo logs to synchronize data.
4. **Open the database**, optionally using RESETLOGS if the redo stream has been reset.



RMAN automates these phases, ensuring the correct redo sequence is applied in the right order.

## 5.3 Instance Recovery

Instance recovery is the simplest form of recovery and occurs automatically when the database is restarted after an abnormal shutdown.

The background process SMON performs the following steps:

- **Roll forward** all committed changes from the redo logs that were not yet written to the datafiles.
- **Rollback** uncommitted transactions using undo data.

### Procedure

1. Attempt to start the instance normally:

SQL> STARTUP;

2. Oracle detects the need for instance recovery and performs it automatically.
3. Review the alert log or trace files for entries such as:

Instance recovery complete. Database opened.

4. No manual intervention is required unless redo or undo files are missing.

If the database fails to start due to missing redo or control files, perform media recovery as described in later sections.

## 5.4 Determining the Type of Recovery Required

Before initiating recovery, assess the situation carefully:

Symptom	Possible Cause	Required Action
Database fails to open with “file not found” error	Datafile or control file deleted or inaccessible.	Perform RMAN restore and recover.
Alert log reports “corrupt block detected”	Media corruption on disk.	Block-level media recovery.



Symptom	Possible Cause	Required Action
Database mounts but fails to open	Incomplete recovery required.	Recover until specific SCN or time.
Instance fails during startup	Instance recovery required.	Restart database; Oracle will perform recovery automatically.

Identifying the correct recovery path prevents unnecessary data loss or downtime.

---

## 5.5 Complete Database Recovery

A complete recovery restores all database files and applies all redo available, returning the database to its most recent consistent state.

### Procedure

1. Mount the database:

rman TARGET /

RMAN> STARTUP MOUNT;

2. Restore the database from backup:

RMAN> RESTORE DATABASE;

3. Recover the database by applying redo:

RMAN> RECOVER DATABASE;

4. Open the database normally:

RMAN> ALTER DATABASE OPEN;

5. Confirm successful recovery in the alert log.

### Notes

- If control-file autobackup is enabled, RMAN automatically restores it if needed.
- This operation preserves the existing redo sequence; no RESETLOGS is required.
- Use VALIDATE before recovery to ensure backup integrity.



## 5.6 Incomplete Database Recovery

Incomplete recovery (also known as **point-in-time recovery**) restores the database to a previous SCN, timestamp, or log sequence when complete recovery is not possible or desirable—for example, to reverse user error.

### Procedure

1. Mount the database but do not open it:

```
rman TARGET /
```

```
RMAN> STARTUP MOUNT;
```

2. Restore the necessary files:

```
RMAN> RESTORE DATABASE;
```

3. Recover until a specific SCN, time, or log sequence:

```
RMAN> RUN {
```

```
  SET UNTIL TIME "TO_DATE('2025-11-10 23:00:00','YYYY-MM-DD HH24:MI:SS')";
```

```
  RESTORE DATABASE;
```

```
  RECOVER DATABASE;
```

```
}
```

or

```
RMAN> SET UNTIL SCN 123456789;
```

```
RMAN> RESTORE DATABASE;
```

```
RMAN> RECOVER DATABASE;
```

4. Open the database with RESETLOGS:

```
RMAN> ALTER DATABASE OPEN RESETLOGS;
```

5. Perform an immediate new full backup to establish a new baseline.

### Notes

- Incomplete recovery creates a new database incarnation.
- Backups made before the resetlogs operation remain usable but must be referenced via the catalog if further recovery is needed.



- Always record the SCN or timestamp used in recovery documentation.

---

## 5.7 Control-File Recovery During Database Recovery

If the control file is missing or damaged, RMAN can restore it automatically from the control-file autobackup.

### Procedure

1. Start the instance in NOMOUNT mode:

RMAN> STARTUP NOMOUNT;

2. Restore the control file:

RMAN> RESTORE CONTROLFILE FROM AUTOBACKUP;

3. Mount the database:

RMAN> ALTER DATABASE MOUNT;

4. Continue with full or incomplete recovery:

RMAN> RESTORE DATABASE;

RMAN> RECOVER DATABASE;

RMAN> ALTER DATABASE OPEN RESETLOGS;

5. If autobackup is not configured, manually specify the backup piece name:

RMAN> RESTORE CONTROLFILE FROM '/backup/CTL\_12345\_abc.bkp';

## 5.8 Datafile Media Recovery

When a single datafile becomes corrupted or lost, it can be restored independently without affecting the rest of the database.

### Procedure

1. Identify the affected file:

SQL> SELECT FILE#, NAME FROM V\$DATAFILE WHERE STATUS != 'ONLINE';

2. Take the datafile offline:



```
ALTER DATABASE DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf' OFFLINE;
```

3. Restore and recover only the damaged file:

```
RMAN> RESTORE DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf';
```

```
RMAN> RECOVER DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf';
```

4. Bring the file back online:

```
ALTER DATABASE DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf' ONLINE;
```

This approach avoids downtime for unaffected datafiles.

---

## 5.9 Block Media Recovery

Block-level recovery corrects corruption in specific blocks without restoring an entire datafile.

### Procedure

1. Identify corrupted blocks using RMAN or the alert log:

```
RMAN> VALIDATE DATABASE;
```

or

```
SELECT * FROM V$DATABASE_BLOCK_CORRUPTION;
```

2. Recover corrupted blocks:

```
RMAN> RECOVER CORRUPTION LIST;
```

or specify a range:

```
RMAN> BLOCKRECOVER DATAFILE 5 BLOCK 120 TO 140;
```

3. Verify correction:

```
RMAN> VALIDATE DATAFILE 5;
```

Block media recovery is available only when valid backups exist that contain the uncorrupted version of the affected blocks.

---



## 5.10 Using Datafile Copies for Fast Recovery

If datafile copies exist on disk (for example, from incrementally updated backups), recovery can switch instantly to the copy.

### Procedure

```
RMAN> SWITCH DATAFILE ALL TO COPY;
```

```
RMAN> RECOVER DATABASE;
```

```
RMAN> ALTER DATABASE OPEN;
```

This technique eliminates restore time and is ideal for systems that use daily incremental merges.

---

## 5.11 Recovery Using Archived Logs Only

When a datafile is intact but inconsistent due to media loss or a crash during backup, apply archived logs to resynchronize it.

```
RMAN> RECOVER DATAFILE 7;
```

RMAN automatically locates and applies archived and online redo logs until the datafile is consistent with the control file checkpoint.

---

## 5.12 Recovering a Dropped Tablespace

If a tablespace was accidentally dropped, it can be recovered from backup by restoring and recovering the relevant datafiles.

### Procedure

1. Mount the database:

```
RMAN> STARTUP MOUNT;
```

2. Restore the tablespace:

```
RMAN> RESTORE TABLESPACE users;
```

3. Recover it:

```
RMAN> RECOVER TABLESPACE users;
```

4. Open the database with resetlogs if required:



RMAN> ALTER DATABASE OPEN RESETLOGS;

---

## 5.13 Testing Recovery Scenarios

Recovery procedures should be validated regularly.

Testing ensures that scripts, paths, and operators are prepared for real-world emergencies.

### Recommended Test Cases

1. Instance failure followed by automatic recovery.
2. Media failure simulated by deleting a datafile.
3. Incomplete recovery using an earlier SCN.
4. Control-file loss and restoration from autobackup.
5. Block corruption recovery using BLOCKRECOVER.

Each scenario should be documented with time to recover, files restored, and validation results.

Testing verifies both the procedure and the team's readiness.

---

## 5.14 Post-Recovery Verification

After recovery completes:

1. Query V\$DATABASE and V\$DATAFILE to confirm all files are ONLINE.
2. Run DBVERIFY or RMAN VALIDATE to confirm datafile integrity.
3. Review alert logs for “Media Recovery Complete” messages.
4. Check application connectivity and transaction consistency.
5. Perform a full backup to capture the post-recovery state.

Verification ensures the database is stable and that recovery objectives were met.

---

## 5.15 Preventing Future Recovery Events

Recovery time is reduced when the system is designed for resiliency.

Recommended preventive measures include:



- Multiplex control files across different disks.
- Place redo logs on mirrored or high-speed storage.
- Schedule frequent backups of archived logs.
- Monitor tablespace free space to prevent file auto-extend errors.
- Test restore scripts quarterly.
- Validate backups automatically and report exceptions.

The goal is not only to recover but to minimize the need for recovery.

---

## 5.16 Summary

This chapter detailed the processes for restoring and recovering databases after various failure scenarios.

Key takeaways include:

- **Instance recovery** is automatic and requires no manual action.
- **Media recovery** uses RMAN to restore and synchronize datafiles, control files, or logs.
- **Complete recovery** restores to the most recent consistent state, while **incomplete recovery** returns to an earlier point in time.
- RMAN provides fine-grained control for **block-level** and **tablespace-level** restoration.
- Testing, validation, and prevention complete the recovery cycle.

With these procedures mastered, the DBA can reliably recover any on-premise Oracle database to a consistent and operational state.

The next chapter focuses on **Control File and SPFILE Recovery**, extending recovery methods to situations where the database structure definition itself is lost.



---

## Chapter 6 – Control File and SPFILE Recovery

---

### 6.1 Introduction

The **control file** and **server parameter file (SPFILE)** are the foundation of every Oracle database instance.

Losing either can prevent startup, backup, or recovery operations.

Fortunately, both can be recreated or restored quickly when properly backed up.

This chapter provides comprehensive procedures for detecting, restoring, and verifying these critical files using RMAN and SQL\*Plus.

All examples assume the environment uses **disk-based backups, non-ASM storage**, and **autobackup** is enabled.

---

### 6.2 Understanding the Control File

The control file is a small binary file that records the physical structure and operational state of the database.

It contains essential metadata including:

- Database name and unique DBID
- Datafile and redo log locations
- Checkpoints and SCNs
- Archive log history
- Backup metadata for RMAN operations

Each database typically maintains **multiple copies** of the control file, specified by the **CONTROL\_FILES** initialization parameter.

These should be stored on **separate disks** to protect against physical device failure.

If all copies are lost or corrupted, the database cannot be mounted until a control file is restored or recreated.

---



## 6.3 Detecting Control File Failure

Symptoms of control file loss include:

- Startup fails with ORA-00205: error in identifying control file
- RMAN cannot connect to the target database because the control file cannot be read
- Database alert log reports corruption or missing files under \$ORACLE\_BASE/diag

### Verification

SQL> SELECT NAME FROM V\$CONTROLFILE;

If this query fails, or if the file paths returned are invalid or missing, recovery is required.

---

## 6.4 Restoring Control Files from Autobackup

When CONFIGURE CONTROLFILE AUTOBACKUP ON is enabled, RMAN automatically backs up the control file after every backup or structural change.

### Procedure

1. Start the instance in **NOMOUNT** mode:

rman TARGET /

RMAN> STARTUP NOMOUNT;

2. Restore the control file:

RMAN> RESTORE CONTROLFILE FROM AUTOBACKUP;

3. Mount the database:

RMAN> ALTER DATABASE MOUNT;

4. Verify success:

RMAN> LIST BACKUP OF CONTROLFILE;

5. Continue with full database recovery if other files are missing:

RMAN> RESTORE DATABASE;

RMAN> RECOVER DATABASE;

RMAN> ALTER DATABASE OPEN RESETLOGS;



## Notes

- RMAN automatically searches for autobackup files using the DBID.
- If multiple autobackups exist, RMAN selects the most recent valid one.
- After restoration, immediately perform a new backup to reestablish redundancy.

---

## 6.5 Restoring Control Files from a Known Backup Piece

If autobackup was not enabled or if the autobackup cannot be located, specify the backup file directly.

### Procedure

1. Identify a valid backup piece containing a control file copy.
2. Start RMAN in **NOMOUNT** state.
3. Execute:

RMAN> RESTORE CONTROLFILE FROM '/backup/CTL\_20250311\_01.bkp';

4. Mount the database and perform recovery as required:

RMAN> ALTER DATABASE MOUNT;

RMAN> RECOVER DATABASE;

RMAN> ALTER DATABASE OPEN RESETLOGS;

---

## 6.6 Recreating the Control File Manually

When no valid backup is available, you can recreate the control file using SQL commands.

This requires knowledge of the database structure (datafile names, redo log groups, and character set).

### Procedure

1. Create a text script with the control-file creation statement:

STARTUP NOMOUNT;

CREATE CONTROLFILE REUSE DATABASE "PROD" RESETLOGS ARCHIVELOG

MAXLOGFILES 16



MAXLOGMEMBERS 3

MAXDATAFILES 100

MAXINSTANCES 1

MAXLOGHISTORY 292

## LOGFILE

GROUP 1 '/u01/app/oracle/oradata/PROD/redo01.log' SIZE 100M,

GROUP 2 '/u01/app/oracle/oradata/PROD/redo02.log' SIZE 100M,

GROUP 3 '/u01/app/oracle/oradata/PROD/redo03.log' SIZE 100M

## DATAFILE

'/u01/app/oracle/oradata/PROD/system01.dbf',

'/u01/app/oracle/oradata/PROD/sysaux01.dbf',

'/u01/app/oracle/oradata/PROD/undo01.dbf',

'/u01/app/oracle/oradata/PROD/users01.dbf'

CHARACTER SET AL32UTF8;

2. Run the script in SQL\*Plus.
3. Recover the database using existing archived logs:

RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL;

4. Open with resetlogs:

ALTER DATABASE OPEN RESETLOGS;

### **Important:**

Immediately perform a new full backup to establish a new control-file baseline.

## 6.7 Multiplexing Control Files

Oracle recommends maintaining at least **two or three multiplexed copies** of the control file on separate physical devices.

### **Procedure**



1. Confirm current locations:

```
SHOW PARAMETER control_files;
```

2. Add a new control-file copy:

```
ALTER SYSTEM SET CONTROL_FILES =
```

```
'/u01/app/oracle/oradata/PROD/control01.ctl',
'/u02/app/oracle/oradata/PROD/control02.ctl',
'/u03/app/oracle/oradata/PROD/control03.ctl'
```

```
SCOPE=SPFILE;
```

3. Restart the instance:

```
SHUTDOWN IMMEDIATE;
```

```
STARTUP;
```

Multiplexed copies protect against single-device failures and ensure continuous recoverability.

---

## 6.8 Understanding the Server Parameter File (SPFILE)

The **SPFILE** stores initialization parameters that define database startup behavior.

Unlike the traditional text-based PFILE, the SPFILE is binary and can be modified dynamically.

### Key Benefits

- Allows persistent changes without manual editing.
- Supports parameter scope (MEMORY, SPFILE, BOTH).
- Ensures synchronization between memory and disk parameter sets.

If the SPFILE is missing or damaged, Oracle cannot start the instance in normal mode.

---

## 6.9 Creating a Backup of the SPFILE

Regular SPFILE backups ensure configuration recoverability.

### RMAN Backup

```
RMAN> BACKUP SPFILE FORMAT '/backup/SPFILE_%U.bkp';
```



## SQL\*Plus Text Copy

```
CREATE PFILE='/backup/init_backup.ora' FROM SPFILE;
```

Maintaining both binary and text copies provides flexibility for manual restoration.

---

## 6.10 Restoring the SPFILE from RMAN

If the SPFILE is lost or corrupted, restore it from the most recent backup.

### Procedure

1. Start the instance using a temporary PFILE:

```
STARTUP NOMOUNT PFILE='/backup/init_backup.ora';
```

2. Restore the SPFILE from backup:

```
RMAN> RESTORE SPFILE FROM AUTOBACKUP;
```

3. Restart the instance using the restored SPFILE:

```
SHUTDOWN IMMEDIATE;
```

```
STARTUP NOMOUNT;
```

4. Confirm active parameter usage:

```
SHOW PARAMETER spfile;
```

---

## 6.11 Recreating the SPFILE Manually

When no backup is available, you can recreate the SPFILE from a text initialization file.

### Procedure

1. Locate the text-based PFILE (usually \$ORACLE\_HOME/dbs/initSID.ora or %ORACLE\_HOME%\database\initSID.ora).
2. Modify as necessary to include valid control-file paths and memory parameters.
3. Create a new SPFILE:

```
CREATE SPFILE FROM PFILE='/u01/app/oracle/admin/initPROD.ora';
```

4. Restart the database:



```
SHUTDOWN IMMEDIATE;  
STARTUP;
```

---

## 6.12 Protecting the Control File and SPFILE

To avoid future recovery operations:

- Enable **CONTROLFILE AUTOBACKUP** and verify daily.
- Back up the SPFILE after any parameter change.
- Keep copies on both local and remote storage.
- Document control-file and parameter locations in the DBA runbook.
- Restrict OS permissions to prevent accidental deletion.
- After every database upgrade, manually copy both files to archival media.

These preventive measures reduce downtime and simplify future restores.

---

## 6.13 Post-Recovery Validation

After restoring control files or the SPFILE:

1. Validate database mount and open operations.
2. Run:

```
SELECT * FROM V$CONTROLFILE;
```

```
SHOW PARAMETER spfile;
```

3. Confirm that all expected control-file copies exist and that the correct SPFILE is in use.
4. Perform an immediate **full backup** of the entire database to establish a new recovery baseline.
5. Record recovery actions in your operational log for audit purposes.

---



## 6.14 Troubleshooting Common Errors

Error	Description	Resolution
<b>ORA-00205</b>	Control file not found or inaccessible.	Restore from autobackup or recreate manually.
<b>ORA-01507</b>	Database not mounted after restore.	Mount database using <code>ALTER DATABASE MOUNT;</code>
<b>ORA-01078 / LRM-00109</b>	Missing or corrupt SPFILE.	Recreate SPFILE from PFILE or restore from RMAN.
<b>ORA-19809</b>	FRA full during autobackup.	Delete obsolete backups or increase FRA size.
<b>ORA-00214</b>	Inconsistent control files detected.	Restore all control files from same backup.

Thorough log review in the alert file usually reveals the exact file path or parameter that caused the failure.

---

## 6.15 Recommended RMAN CONFIGURE Settings

The following commands can be run one time to set the RMAN configuration parameters.

```
-- Retention: keep full recoverability while not hoarding backups
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;

-- Don't re-backup files that haven't changed / already protected
CONFIGURE BACKUP OPTIMIZATION ON;

-- Control file & SPFILE autobackups (you already have this on)
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK
TO 'D:\DB_Backups\RMAN\DEVCDB\controlfile_%F';

-- Channels & format (keep your foldering; you can raise parallelism later)
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET;
CONFIGURE CHANNEL DEVICE TYPE DISK
FORMAT 'D:\DB_Backups\RMAN\DEVCDB\bk_%d_%T_%U';

-- Compression is useful for disk backups (CPU permitting)
```



```

CONFIGURE COMPRESSION ALGORITHM 'BASIC';
-- Only allow deletion of archivelogs that are already backed up at least
once
CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO DISK;

```

Why these settings? A **recovery-window** policy is the standard for reliable PITR while letting **DELETE OBSOLETE** do the right thing, and turning **backup optimization** on avoids redundant re-backups of unchanged files and already-captured archivelogs.

---

## 6.16 Nightly Backup Job (Windows to Disk)

```

-- Daily full + archivelog with post-backup cleanup (Windows paths)

RUN {
  SET COMMAND ID TO 'DAILY_FULL';

  -- Sanity: sync metadata so deletes are accurate
  CROSSCHECK BACKUP;
  CROSSCHECK ARCHIVELOG ALL;

  -- Back up SPFILE and CONTROLFILE explicitly (in addition to autobackup)
  BACKUP AS COMPRESSED BACKUPSET
    SPFILE TAG 'DAILY_SPFILE'
    FORMAT 'D:\DB_Backups\RMAN\DEVCDB\spfile_%d_%T_%U';

  BACKUP AS COMPRESSED BACKUPSET
    CURRENT CONTROLFILE TAG 'DAILY_CONTROLFILE'
    FORMAT 'D:\DB_Backups\RMAN\DEVCDB\controlfile_%F';

  -- Core: database + archivelogs in one run
  BACKUP AS COMPRESSED BACKUPSET
    DATABASE
    PLUS ARCHIVELOG
    TAG 'DAILY_FULL'
    FORMAT 'D:\DB_Backups\RMAN\DEVCDB\bk_%d_%T_%U';

  -- Force a fresh log switch so the very latest redo is captured
  SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';

  -- Make sure any new logs get picked up at least once
  BACKUP AS COMPRESSED BACKUPSET
    ARCHIVELOG ALL NOT BACKED UP 1 TIMES
    TAG 'DAILY_ARCH';

  -- Delete archivelogs older than 2 days *that have been backed up*.
  -- The ARCHIVELOG DELETION POLICY enforces the "must be backed up" rule.
  DELETE NOPROMPT ARCHIVELOG ALL COMPLETED BEFORE 'SYSDATE-2';

```



```
-- Clear out metadata for missing files, then purge per retention policy
DELETE NOPROMPT EXPIRED BACKUP;
DELETE NOPROMPT OBSOLETE; -- uses the 7-day recovery window

}
```

### Notes & small RMAN options you can toggle

- **Parallelism:** If disk can handle it, try CONFIGURE DEVICE TYPE DISK PARALLELISM 4 and measure throughput (keep FORMAT the same).
- **Compression:** If CPU is tight, you can switch to CONFIGURE COMPRESSION ALGORITHM 'LOW' for faster backups with less compression.
- **Retention:** If you prefer a shorter recovery window, set RECOVERY WINDOW OF 5 DAYS; the rest still works.
- **Why DELETE OBSOLETE and not just DELETE BACKUP?** It aligns with the retention policy so you never delete something still needed for your recovery window.

## 6.17 Summary

In this chapter, you learned how to detect, restore, and protect Oracle's two most critical metadata components:

- The **control file**, which records the physical and logical structure of the database.
- The **SPFILE**, which defines instance-level parameters and startup behavior.

You practiced restoring these components from RMAN autobackup, from specified backup pieces, and recreating them manually when no backup exists.

You also learned preventive measures—such as multiplexing control files and maintaining text copies of parameter files—to reduce recovery time and risk.

With the control file and SPFILE secured, the database foundation is complete.

The next chapter focuses on **Tablespace and Table-Level Recovery**, enabling you to recover specific portions of the database without restoring the entire system.



---

## Chapter 7 – Tablespace and Table-Level Recovery

---

### 7.1 Introduction

Tablespace and table-level recovery allow database administrators to restore only a portion of the database instead of performing a complete restore and recovery.

These targeted techniques minimize downtime, reduce resource usage, and provide flexibility for both planned and unplanned events.

This chapter covers the full range of partial recovery methods available for on-premise databases running Oracle 12c and above, including **Tablespace Point-in-Time Recovery (TSPITR)**, **individual datafile recovery**, and **table recovery** using RMAN.

---

### 7.2 Understanding Tablespace Recovery

A **tablespace** is a logical container of datafiles. When a tablespace or its datafiles are damaged or accidentally modified, RMAN can restore and recover them independently.

Tablespace recovery is typically required in the following situations:

- A tablespace is accidentally dropped or truncated.
- Datafiles within a tablespace become corrupted.
- A subset of the database must be restored to an earlier time without affecting other data.
- Testing or QA environments require rollback of a single functional area.

By recovering only the affected area, service for the rest of the database remains uninterrupted.

---

### 7.3 Offline Tablespace Recovery

When a datafile or tablespace becomes damaged but other areas are operational, the affected objects can be taken offline for recovery.

#### Procedure

1. Identify the damaged tablespace:



SELECT TABLESPACE\_NAME, STATUS FROM DBA\_TABLESPACES;

2. Take the tablespace offline:

ALTER TABLESPACE hr\_data OFFLINE IMMEDIATE;

3. Restore and recover the tablespace using RMAN:

RMAN> RESTORE TABLESPACE hr\_data;

RMAN> RECOVER TABLESPACE hr\_data;

4. Bring the tablespace back online:

ALTER TABLESPACE hr\_data ONLINE;

### Notes

- The database remains open during this process.
- If other tablespaces depend on the recovered one (e.g., foreign key relationships), validate integrity after completion.

## 7.4 Recovering a Dropped Tablespace

If a tablespace was dropped, its metadata is removed from the control file, but the physical datafiles may still exist in backups.

### Procedure

1. Mount the database:

RMAN> STARTUP MOUNT;

2. Restore and recover the tablespace:

RMAN> RESTORE TABLESPACE hr\_data;

RMAN> RECOVER TABLESPACE hr\_data;

3. Open the database:

RMAN> ALTER DATABASE OPEN RESETLOGS;

This operation reintroduces the tablespace into the control file based on the restored datafiles.



## 7.5 Partial Datafile Recovery within a Tablespace

Sometimes only one or two datafiles within a tablespace are damaged.

In this case, you can restore and recover those files individually.

### Procedure

1. Identify affected datafiles:

```
SELECT FILE#, NAME FROM V$DATAFILE WHERE TABLESPACE_NAME = 'HR_DATA';
```

2. Take only the damaged files offline:

```
ALTER DATABASE DATAFILE '/u01/oradata/hr_data02.dbf' OFFLINE;
```

3. Restore and recover:

```
RMAN> RESTORE DATAFILE '/u01/oradata/hr_data02.dbf';
```

```
RMAN> RECOVER DATAFILE '/u01/oradata/hr_data02.dbf';
```

4. Bring them back online:

```
ALTER DATABASE DATAFILE '/u01/oradata/hr_data02.dbf' ONLINE;
```

This method limits disruption to the specific datafiles that experienced failure.

## 7.6 Tablespace Point-in-Time Recovery (TSPITR)

**TSPITR** allows recovery of one or more tablespaces to an earlier time or SCN while leaving the rest of the database intact.

This is essential for reversing logical errors such as accidental data deletion in a specific schema or module.

TSPITR uses an **auxiliary instance** to isolate and recover the selected tablespaces before importing them back into the target database.

### Conceptual Steps

1. **Create an auxiliary instance** to host temporary recovery files.
2. **Restore datafiles** of the target tablespace from backup.
3. **Recover** them to the desired point in time using archived redo logs.



4. **Export the recovered tablespace's metadata** and data.
5. **Import** the recovered tablespace back into the target database.

---

## RMAN Automated TSPITR

Starting with Oracle 12c, RMAN automates most of the TSPITR process.

### Procedure

1. Identify the target tablespace and the recovery point:

```
SELECT TABLESPACE_NAME FROM DBA_TABLESPACES WHERE TABLESPACE_NAME='HR_DATA';
```

2. Connect to RMAN and run:

```
rman TARGET / AUXILIARY sys/password@AUXDB
```

```
RUN {
```

```
RECOVER TABLESPACE hr_data
UNTIL TIME "TO_DATE('2025-11-10 10:00:00','YYYY-MM-DD HH24:MI:SS')"
AUXILIARY DESTINATION '/u01/app/oracle/aux_tspir';
```

```
}
```

3. RMAN automatically:

- o Creates a temporary auxiliary instance.
- o Restores the tablespace's datafiles from backup.
- o Recovers them to the specified time.
- o Reimports the recovered data into the target database.

4. Confirm the operation:

```
SELECT TABLESPACE_NAME, STATUS FROM DBA_TABLESPACES;
```

5. Drop the auxiliary instance's temporary files when complete.

### Notes

- TSPITR cannot include SYSTEM, UNDO, or TEMP tablespaces.



- The auxiliary destination must have sufficient space for temporary datafiles and archived redo logs.
- Always perform a new backup after TSPITR to capture the updated database state.

---

## 7.7 Recovering a Single Table Using RMAN

Oracle 12c introduced **table-level recovery** using RMAN, which eliminates the need for full TSPITR when only specific tables are affected.

This feature recovers individual tables or partitions into an auxiliary instance and automatically imports them back into the target database.

### Procedure

1. Identify the table to recover and its owner:

```
SELECT owner, table_name FROM dba_tables WHERE table_name='EMPLOYEES';
```

2. Execute the recovery in RMAN:

```
rman TARGET / AUXILIARY sys/password@AUXDB
```

```
RUN {
```

```
RECOVER TABLE hr.employees
```

```
UNTIL TIME "TO_DATE('2025-11-10 11:00:00','YYYY-MM-DD HH24:MI:SS')"
```

```
AUXILIARY DESTINATION '/u01/app/oracle/aux_table'
```

```
DATAPUMP DESTINATION '/u01/app/oracle/dump'
```

```
DUMP FILE 'emp_recov_%U.dmp'
```

```
NOTABLEIMPORT;
```

```
}
```

3. Once the table is recovered, import it manually with Data Pump if NOTABLEIMPORT was specified:

```
impdp hr/password DIRECTORY=DPUMP DUMPFILE=emp_recov_01.dmp TABLES=hr.employees
TABLE_EXISTS_ACTION=REPLACE;
```



4. Verify that the table's data is restored:

```
SELECT COUNT(*) FROM hr.employees;
```

### Notes

- Requires a valid RMAN backup that includes the table's tablespace.
- Cannot recover tables residing in SYSTEM or UNDO tablespaces.
- The auxiliary instance must be configured with sufficient temporary storage.
- Always perform a new full backup after recovery.

## 7.8 Recovering Tables Across Schemas

You can recover a table to a different schema using the **REMAP SCHEMA** clause.

### Example

```
RECOVER TABLE hr.employees
UNTIL TIME "TO_DATE('2025-11-10 10:00:00','YYYY-MM-DD HH24:MI:SS')"
AUXILIARY DESTINATION '/u01/app/oracle/aux_table'
DATAPUMP DESTINATION '/u01/app/oracle/dump'
REMAP SCHEMA hr:hr_recovered;
```

This creates the restored table in the HR\_RECOVERED schema while keeping the original untouched.

This approach is useful for testing and data verification before replacing live data.

## 7.9 Recovering Tablespace or Table to a Different Database

RMAN allows recovered objects to be exported from an auxiliary instance and imported into another database entirely.

### Procedure

1. Perform RMAN recovery as before, specifying an auxiliary destination.
2. Use Data Pump to export the recovered table or tablespace dump file.
3. Import into the target database using:



```
impdp system/password DIRECTORY=DPUMP DUMPFILE='recover_dump.dmp' SCHEMAS=HR;
```

This technique is effective for cross-environment restorations (for example, restoring production data to a staging or QA system).

---

## 7.10 Validating Tablespace and Table Recovery

Validation ensures that restored data is usable and consistent.

### Verification Checklist

- Query DBA\_TABLESPACES to confirm online status.
- Run integrity checks:

```
ANALYZE TABLE hr.employees VALIDATE STRUCTURE CASCADE;
```

- Compare record counts or checksums between restored and source environments.
- Review alert logs for recovery completion messages.
- Perform RMAN VALIDATE TABLESPACE hr\_data; for physical integrity.

Only after validation should the restored objects be made accessible to users.

---

## 7.11 Common Restrictions and Considerations

- SYSTEM, UNDO, and TEMP tablespaces cannot be recovered individually.
- TSPITR and table recovery require **ARCHIVELOG** mode.
- The recovery point must exist within the range of available archived logs.
- Tablespace names recovered using TSPITR must not already exist in the target database.
- Always ensure consistent NLS and character set configurations when importing data across databases.

Planning and testing these constraints prevent failure during critical recovery events.

---

## 7.12 Example: End-to-End Table Recovery

The following example demonstrates complete table recovery after user error:



1. A user accidentally deletes records from the HR.EMPLOYEES table.
2. The DBA identifies the time before deletion: 2025-11-10 11:15:00.
3. The DBA runs:

```
rman TARGET / AUXILIARY sys/password@AUXDB
```

```
RUN {
```

```
RECOVER TABLE hr.employees
UNTIL TIME "TO_DATE('2025-11-10 11:10:00','YYYY-MM-DD HH24:MI:SS')"
AUXILIARY DESTINATION '/u01/app/oracle/aux_table'
DATAPUMP DESTINATION '/u01/app/oracle/dump'
DUMP FILE 'hr_emp_restore_%U.dmp';
```

```
}
```

4. RMAN creates a temporary auxiliary instance, recovers the table, and exports it as a Data Pump dump file.
5. The DBA imports it back into the target:

```
impdp hr/password DIRECTORY=DPUMP DUMPFILE=hr_emp_restore_01.dmp
TABLES=hr.employees TABLE_EXISTS_ACTION=REPLACE;
```

6. Verification confirms the table's contents have been restored successfully.
7. A new full backup is taken to capture the post-recovery state.

## 7.13 Performance Optimization During Recovery

- Store auxiliary and Data Pump destinations on fast, local disks.
- Use multiple channels to parallelize large tablespace restores.
- Temporarily disable flashback during recovery to improve throughput.
- Adjust memory settings (SGA\_TARGET, PGA\_AGGREGATE\_TARGET) for auxiliary instances.
- Monitor disk I/O performance using V\$SESSION\_LONGOPS.



Efficient configuration ensures that partial recoveries complete within operational windows.

---

## 7.14 Summary

This chapter demonstrated advanced, selective recovery techniques for Oracle databases:

- Offline and online **tablespace recovery**.
- **Tablespace Point-in-Time Recovery (TSPITR)** using an auxiliary instance.
- **Table-level recovery** introduced in Oracle 12c for restoring individual tables or partitions.
- Recovery to **alternate schemas** or **different databases** for controlled restoration.
- Comprehensive validation, restrictions, and optimization guidelines.

Partial recovery capabilities enable the DBA to respond to localized data loss without impacting the entire system.

With these tools mastered, the next chapter—**Flashback Technologies**—focuses on instant, undo-based recovery methods that can often resolve logical errors even faster than RMAN-based approaches.



## PART IV – FLASHBACK AND ADVANCED RECOVERY

### Chapter 8 – Flashback Technologies

#### 8.1 Introduction

Flashback technologies offer administrators a rapid and flexible means of recovering from **logical errors**—such as accidental data deletions, incorrect updates, or dropped tables—without the need to restore physical backups.

These features leverage **undo data**, **flashback logs**, and **restore points** to reverse changes while the database remains online.

Flashback complements RMAN rather than replacing it:

- RMAN is used for **physical** recovery (e.g., file loss, corruption).
- Flashback is used for **logical** recovery (e.g., bad transactions).

This chapter explains each flashback option, configuration requirements, and the procedures to implement them safely in production environments.

#### 8.2 Flashback Architecture Overview

Flashback relies on several Oracle components working together to provide time-based recovery:

Component	Description
<b>UNDO Tablespace</b>	Stores before-image data for active transactions, used by most flashback operations.
<b>Flashback Logs</b>	Stored in the Fast Recovery Area (FRA), they record historical block images for FLASHBACK DATABASE.
<b>Restore Points</b>	Named markers that record an SCN for quick flashback operations.



Component	Description
<b>System Change Number (SCN)</b>	A monotonically increasing counter representing the database timeline; every flashback target is tied to an SCN.

Flashback operations are efficient because they reuse existing redo and undo information rather than requiring datafile restoration.

---

## 8.3 Enabling Flashback Features

Before using any flashback operation, ensure that:

1. The database runs in **ARCHIVELOG** mode.
2. A **Fast Recovery Area (FRA)** is configured with sufficient space.
3. Flashback Database is explicitly enabled.

### Procedure

-- Confirm ARCHIVELOG mode

ARCHIVE LOG LIST;

-- Set up the FRA if not already configured

ALTER SYSTEM SET DB\_RECOVERY\_FILE\_DEST\_SIZE=200G SCOPE=BOTH;

ALTER SYSTEM SET DB\_RECOVERY\_FILE\_DEST='/u01/app/oracle/FRA' SCOPE=BOTH;

-- Enable flashback

ALTER DATABASE FLASHBACK ON;

Verify that flashback is enabled:

SELECT FLASHBACK\_ON FROM V\$DATABASE;

Output should read **YES**.

Flashback can be disabled when not required to conserve storage:




---

ALTER DATABASE FLASHBACK OFF;

---

## 8.4 Types of Flashback Operations

Oracle provides several flashback features, each addressing a specific recovery scenario:

Feature	Scope	Description
<b>Flashback Query</b>	Row-level	View past data for any table as of a previous SCN or timestamp.
<b>Flashback Table</b>	Table-level	Restore one or more tables to a past point in time.
<b>Flashback Drop</b>	Object-level	Restore a dropped table from the Recycle Bin.
<b>Flashback Transaction Query / Backout</b>	Transaction-level	Identify and undo changes from specific transactions.
<b>Flashback Database</b>	Database-level	Revert the entire database to a prior SCN or restore point.

Each technique can be used independently or combined with RMAN for layered protection.

---

## 8.5 Flashback Query

**Flashback Query** allows you to view data “as of” a previous time without actually modifying the current table.

This is particularly useful for investigating or comparing historical states.

### Example

SELECT \* FROM hr.employees AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '10' MINUTE)

WHERE employee\_id = 100;

or using an SCN:

SELECT \* FROM hr.employees AS OF SCN 1289564;

### Usage Notes



- Requires sufficient undo retention to reconstruct the requested version.
- The database must have the parameter UNDO\_RETENTION set to a duration covering expected query needs.
- Ideal for ad-hoc verification before performing Flashback Table operations.

---

## 8.6 Flashback Table

**Flashback Table** restores one or more tables to a previous point in time, undoing committed changes without affecting other database objects.

### Prerequisites

- Database in ARCHIVELOG mode.
- Flashback Database enabled.
- User must have FLASHBACK and ALTER privileges on the target table.

### Procedure

1. Identify the SCN or timestamp to which you wish to restore:

```
SELECT SCN_TO_TIMESTAMP(CURRENT_SCN - 500) FROM DUAL;
```

2. Perform the flashback:

```
FLASHBACK TABLE hr.employees
```

```
TO TIMESTAMP TO_TIMESTAMP('2025-11-10 10:00:00', 'YYYY-MM-DD HH24:MI:SS');
```

3. Verify the result:

```
SELECT COUNT(*) FROM hr.employees;
```

4. If constraints or foreign keys reference the table, disable them temporarily and re-enable after recovery.

### Notes

- Indexes, triggers, and grants are automatically restored.
- Tables involved in transactions spanning multiple objects may require coordinated recovery.

---



## 8.7 Flashback Drop (Recycle Bin)

The **Recycle Bin** stores dropped tables and their dependent objects, allowing fast recovery.

### View Recycle Bin Contents

```
SHOW RECYCLEBIN;
```

### Restore Dropped Table

```
FLASHBACK TABLE hr.employees TO BEFORE DROP;
```

### Permanently Remove Dropped Objects

```
PURGE RECYCLEBIN;
```

### Considerations

- The recycle bin must be enabled (it is on by default).
- If insufficient disk space exists, Oracle automatically purges the oldest objects.
- Dropped objects from SYSTEM tablespaces are not stored in the recycle bin.

## 8.8 Flashback Transaction Query and Backout

**Flashback Transaction Query** allows you to view the history of changes made by individual transactions.

**Flashback Transaction Backout** enables selective undo of committed transactions.

### Enable Supplemental Logging

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
```

### Identify Transactions

```
SELECT xid, operation, undo_sql
```

```
FROM flashback_transaction_query
```

```
WHERE table_name='EMPLOYEES' AND table_owner='HR';
```

### Undo a Specific Transaction

```
BEGIN
```

```
DBMS_FLASHBACK.TRANSACTION_BACKOUT(numtxns => 1,
```



```

xids => xid_array('05000A0019020000'),
options => dbms_flashback.cascade);
END;
/

```

## Notes

- Use caution; dependent transactions may require cascading undo.
- Suitable for correcting errors such as accidental salary updates or mass changes.
- Always back up the affected tablespaces before performing transaction backout in production.

## 8.9 Flashback Database

**Flashback Database** reverts the entire database to a previous SCN or restore point.

It is faster than full media recovery because it replays block changes from **flashback logs** instead of restoring backups.

### Enabling Flashback Database

```
ALTER SYSTEM SET DB_RECOVERY_FILE_DEST_SIZE=200G;
```

```
ALTER SYSTEM SET DB_RECOVERY_FILE_DEST='/u01/app/oracle/FRA';
```

```
ALTER DATABASE FLASHBACK ON;
```

Confirm activation:

```
SELECT FLASHBACK_ON FROM V$DATABASE;
```

### Creating a Restore Point

```
CREATE RESTORE POINT pre_patch GUARANTEE FLASHBACK DATABASE;
```

### Performing a Flashback

1. Mount but do not open the database:

```
SHUTDOWN IMMEDIATE;
```

```
STARTUP MOUNT;
```

2. Flash back to a restore point or SCN:



FLASHBACK DATABASE TO RESTORE POINT pre\_patch;

-- or --

FLASHBACK DATABASE TO SCN 1234567;

3. Open the database with resetlogs:

ALTER DATABASE OPEN RESETLOGS;

## Notes

- Guaranteed restore points preserve flashback logs until dropped.
- Non-guaranteed restore points are automatically purged when FRA space is needed.
- Flashback Database cannot recover from media corruption or physical file loss; use RMAN for that.

## 8.10 Managing Restore Points

Restore points provide convenient markers for quick rollback operations during planned maintenance or testing.

### Creating Restore Points

CREATE RESTORE POINT before\_upgrade GUARANTEE FLASHBACK DATABASE;

CREATE RESTORE POINT nightly\_snapshot;

### Listing Restore Points

SELECT NAME, GUARANTEE\_FLASHBACK\_DATABASE, SCN, TIME FROM V\$RESTORE\_POINT;

### Dropping Restore Points

DROP RESTORE POINT before\_upgrade;

### Best Practices

- Use guaranteed restore points for high-risk operations (patching, upgrades).
- Drop unused restore points to free space in the FRA.
- Maintain naming conventions (e.g., PRE\_UPGRADE\_YYYYMMDD) for traceability.



## 8.11 Verifying Flashback Success

After performing flashback operations:

1. Check the alert log for confirmation messages such as  
“Flashback complete, SCN adjusted to...”
2. Query the database time or SCN to verify rollback point:

```
SELECT CURRENT_SCN FROM V$DATABASE;
```

3. Validate object structure and data consistency:

```
ANALYZE TABLE hr.employees VALIDATE STRUCTURE CASCADE;
```

4. If performing Flashback Database, back up immediately after reopening to reset the baseline.

## 8.12 Troubleshooting Flashback Operations

Issue	Cause	Resolution
<b>ORA-38760: This database instance failed to turn on flashback</b>	Insufficient FRA size.	Increase FRA size and retry.
<b>ORA-38706: Cannot turn on flashback logging</b>	Database must be in ARCHIVELOG mode.	Enable ARCHIVELOG and reattempt.
<b>ORA-38754: Flashback database not started</b>	Flashback logs not available.	Enable flashback and generate restore point.
<b>ORA-38301: Object not in recycle bin</b>	Table purged or recycle bin disabled.	Restore from backup or Data Pump export.
<b>ORA-38714: Instance not mounted</b>	Flashback Database requires MOUNT mode.	Restart instance in MOUNT mode.

Consistent monitoring of the FRA and undo tablespace prevents most flashback errors.



## 8.13 Best Practices for Flashback Usage

- Enable **ARCHIVELOG** and **FLASHBACK DATABASE** on all production systems.
- Allocate sufficient **FRA** space to accommodate expected flashback retention.
- Create **guaranteed restore points** before software upgrades, schema changes, or patch installations.
- Maintain a documented flashback retention policy (e.g., 48 hours).
- Perform periodic validation by executing test flashbacks in non-production environments.
- Use flashback operations as a supplement—not replacement—for RMAN backups.

---

## 8.14 Summary

Flashback technologies provide immediate recovery from logical errors, often within seconds, without needing to restore from physical backups.

In this chapter, you learned to:

- Configure and enable Flashback Database.
- Query and restore historical data using Flashback Query, Table, and Drop.
- Undo transactions with Flashback Transaction Backout.
- Revert entire databases to restore points using Flashback Database.
- Manage and validate restore points for planned maintenance.

Flashback empowers DBAs to correct mistakes safely and efficiently, greatly reducing downtime and recovery complexity.

The next chapter, **Advanced Recovery and Duplication**, explores cross-platform restoration, cloning, and database duplication techniques using RMAN to expand recovery capabilities across environments.



# Chapter 9 – Advanced Recovery and Duplication

## 9.1 Introduction

Advanced recovery techniques extend the DBA's ability to not only restore a damaged database but also **clone**, **duplicate**, or **migrate** environments for testing, reporting, and standby purposes.

These methods use RMAN's DUPLICATE and CONVERT commands to rebuild a database on the same host or a different one, without impacting the source system.

Database duplication is especially valuable for:

- Creating QA or development environments identical to production.
- Performing recovery testing and validation without affecting live data.
- Moving databases between hosts or operating systems.
- Accelerating migration planning or failover readiness.

## 9.2 Duplication Methods Overview

RMAN supports several duplication strategies depending on connectivity, backup location, and system configuration.

Duplication Type	Description	Connectivity
<b>Active Database Duplication</b>	Copies data directly from the source database over Oracle Net without requiring backup files.	Direct connection between target and auxiliary.
<b>Backup-Based Duplication</b>	Uses existing disk backups to create a duplicate database.	No connection to target required.
<b>Duplicate from Backup Set</b>	Reads backup sets instead of image copies to improve performance.	Source and auxiliary share backup location.



Duplication Type	Description	Connectivity
Duplicate for Standby	Creates a physical standby database for Data Guard.	Not used in this guide's on-premise scope, but conceptually similar.

Each method performs a full copy of the source database, updates the DBID, and optionally renames the database.

---

## 9.3 Preparing for Duplication

Before starting duplication, verify the following prerequisites:

1. The target database is accessible or its backups are available.
2. The auxiliary database has an identical or compatible Oracle version.
3. The auxiliary instance can start in **NOMOUNT** mode.
4. Network connectivity (for active duplication) is tested using TNS or local connections.
5. Sufficient disk space exists to accommodate the full size of the source database.

### Auxiliary Initialization Parameters

Create a minimal initialization file on the auxiliary host:

DB\_NAME=PROD\_CLONE

DB\_UNIQUE\_NAME=PROD\_CLONE

CONTROL\_FILES='/u01/app/oracle/oradata/PROD\_CLONE/control01.ctl'

DB\_CREATE\_FILE\_DEST='/u01/app/oracle/oradata/PROD\_CLONE'

DB\_RECOVERY\_FILE\_DEST='/u01/app/oracle/FRA/PROD\_CLONE'

DB\_RECOVERY\_FILE\_DEST\_SIZE=100G

Start the instance:

STARTUP NOMOUNT PFILE='/u01/app/oracle/admin/initPROD\_CLONE.ora';

---



## 9.4 Active Database Duplication

Active duplication copies live data directly from the source database to the auxiliary instance via Oracle Net.

This method is fast and efficient for local or high-speed network environments.

### Procedure

1. Ensure both target and auxiliary databases are accessible:

```
rman TARGET sys/password@PROD AUXILIARY sys/password@PROD_CLONE;
```

2. Run the duplication command:

```
DUPLICATE TARGET DATABASE TO PROD_CLONE
  FROM ACTIVE DATABASE
  SPFILE
  PARAMETER_VALUE_CONVERT
  ('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
  SET DB_FILE_NAME_CONVERT =
  ('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
  SET LOG_FILE_NAME_CONVERT =
  ('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
  NOFILENAMECHECK;
```

3. RMAN automatically:

- Creates control files for the clone.
- Copies datafiles over the network.
- Applies archived redo to ensure consistency.
- Opens the new database with a new DBID.

4. Verify the clone:

```
SELECT NAME, DB_UNIQUE_NAME FROM V$DATABASE;
```

### Advantages

- No need for backup files.



- Real-time duplication of active production systems.

## Limitations

- Requires network bandwidth equal to or greater than backup throughput.
- Higher CPU load on the source system during duplication.

## 9.5 Backup-Based Duplication

Backup-based duplication creates a database from existing RMAN backup pieces.

It is ideal when a direct network connection to the source database is not possible.

### Procedure

1. Copy all necessary backups and archived redo logs from the source host to the auxiliary host.
2. Start the auxiliary instance in NOMOUNT mode:

STARTUP NOMOUNT;

3. Connect to RMAN:

```
rman TARGET sys/password@none CATALOG rmanuser/rmanpw@CATDB AUXILIARY
sys/password@PROD_CLONE;
```

4. Issue the duplication command:

DUPLICATE DATABASE TO PROD\_CLONE

BACKUP LOCATION '/backup/rman'

SPFILE

```
PARAMETER_VALUE_CONVERT
('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
```

```
SET DB_FILE_NAME_CONVERT =
('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
```

```
SET LOG_FILE_NAME_CONVERT =
('/u01/app/oracle/oradata/PROD','/u01/app/oracle/oradata/PROD_CLONE')
```

NOFILENAMECHECK;



5. RMAN restores control files, restores datafiles from backups, and recovers to the most recent SCN.
6. Verify the clone:

```
SELECT DBID, NAME FROM V$DATABASE;
```

### Advantages

- No load on production system.
- Suitable for offline or isolated networks.

### Limitations

- Requires current backups to be transferred manually.
- Slightly longer duration than active duplication.

## 9.6 Duplicating for Testing or QA

Creating duplicate databases for testing or quality assurance allows teams to experiment with schema changes, patches, or data analysis without affecting production.

### Steps

1. Perform backup-based duplication as described above.
2. Rename the duplicate database using the DB\_NEwid utility if needed:

```
nid TARGET=sys/password DBNAME=TESTDB SETNAME=YES
```

3. Update connection identifiers and TNS entries.
4. Modify application configuration to connect to the test environment.

### Post-Duplication Actions

- Disable jobs, replication, and scheduled backups on the clone.
- Change passwords or remove sensitive data if required.
- Document the duplication for compliance and audit tracking.



## 9.7 Creating Test Databases from Backups

For recurring test refreshes, a simple RMAN script can automate restoration and recovery to a specific point in time.

### Example Script

```
RUN {
  STARTUP NOMOUNT;
  RESTORE CONTROLFILE FROM AUTOBACKUP;
  ALTER DATABASE MOUNT;
  RESTORE DATABASE;
  RECOVER DATABASE UNTIL TIME "TO_DATE('2025-11-10 23:00:00','YYYY-MM-DD HH24:MI:SS')";
  ALTER DATABASE OPEN RESETLOGS;
}
```

This method provides a controlled refresh that can be repeated on demand, maintaining environment consistency across development and QA systems.

## 9.8 Encrypted and Compressed Backups in Duplication

When duplicating from encrypted or compressed backups:

- RMAN automatically decrypts backups if the same wallet or password is available.
- For password-encrypted backups, specify the password in the duplication command:

DUPLICATE DATABASE TO PROD\_CLONE

BACKUP LOCATION '/backup/encrypted'

DECRYPT USING PASSWORD 'StrongPassword'

NOFILENAMECHECK;

Compression reduces duplication time by minimizing disk I/O, though CPU usage may increase. Always test decryption prior to duplication to ensure compatibility.



## 9.9 Multitenant Database Duplication (CDB/PDB)

For container databases (CDB) and pluggable databases (PDB), RMAN supports duplication at multiple levels.

### Duplicate Entire CDB

```
DUPLICATE DATABASE TO CDBCLONE FROM ACTIVE DATABASE
```

```
USING COMPRESSED BACKUPSET
```

```
SPFILE
```

```
PARAMETER_VALUE_CONVERT ('CDB1','CDBCLONE')
```

```
NOFILENAMECHECK;
```

### Duplicate Individual PDB

```
DUPLICATE PLUGGABLE DATABASE pdb_sales
```

```
TO pdb_sales_clone
```

```
FROM ACTIVE DATABASE
```

```
FILE_NAME_CONVERT =
```

```
('/u01/app/oracle/oradata/CDB1/pdb_sales','/u01/app/oracle/oradata/CDB1/pdb_sales_clone');
```

This capability simplifies environment cloning for specific business units or modules in multitenant architectures.

## 9.10 Cross-Platform Duplication and Conversion

When moving a database between platforms with different endian formats (for example, from Linux x86-64 to Windows), use the RMAN CONVERT DATABASE command.

### Procedure

1. Confirm platform compatibility:

```
SELECT PLATFORM_ID, PLATFORM_NAME, ENDIAN_FORMAT FROM
V$TRANSPORTABLE_PLATFORM;
```

2. Backup the source database:

```
BACKUP DATABASE FORMAT '/backup/portable/%U' SECTION SIZE 10G;
```



3. Convert the backup to the target platform:

```
CONVERT DATABASE NEW DATABASE 'PROD_WIN'
```

```
FORMAT '/backup/portableConverted/%U'
```

```
TO PLATFORM 'Microsoft Windows x86 64-bit';
```

4. Transfer files to the destination and restore as usual.

Cross-platform duplication is especially useful for migration projects and disaster recovery testing.

---

## 9.11 Converting Backups into Data Pump Exports

For selective recovery or cross-environment migration, you can restore backups into an auxiliary instance, then use Data Pump to export specific schemas or tables.

### Procedure

1. Restore the database in an auxiliary environment.

2. Export required schemas:

```
expdp system/password DIRECTORY=DPUMP DUMPFILE=partial_recovery.dmp
SCHEMAS=HR,FINANCE;
```

3. Import into another environment:

```
impdp system/password DIRECTORY=DPUMP DUMPFILE=partial_recovery.dmp
SCHEMAS=HR,FINANCE TABLE_EXISTS_ACTION=REPLACE;
```

This hybrid method combines RMAN's physical recovery speed with Data Pump's flexibility.

---

## 9.12 Verifying and Securing Duplicated Databases

After duplication or conversion, always perform these checks:

1. Verify DBID difference to ensure the clone is unique:

```
SELECT DBID, NAME FROM V$DATABASE;
```

2. Run validation:

```
RMAN> VALIDATE DATABASE;
```



3. Disable replication jobs and database links to production.
4. Verify that backups and control files point to the clone's file paths.
5. Secure credentials and update listener configuration.

Perform an immediate full backup of the duplicated environment to establish its own recovery baseline.

---

## 9.13 Common Duplication Errors

Error	Cause	Resolution
<b>ORA-12514</b>	TNS listener not recognizing service name.	Verify tnsnames.ora and restart listener.
<b>ORA-19505</b>	File path invalid or permissions insufficient.	Confirm mount points and directory access.
<b>ORA-01100</b>	Database already mounted.	Shutdown auxiliary instance and retry.
<b>RMAN-06025</b>	No backup of the specified file found.	Verify backup catalog or backup location.
<b>RMAN-05537</b>	DUPLICATE requires target or catalog connection.	Ensure all connections are established correctly.

Maintaining consistent environment configuration files across servers minimizes these errors.

---

## 9.14 Best Practices for Advanced Recovery

- Document every duplication and record DBID, timestamp, and SCN.
- Keep consistent file naming conventions between source and target environments.
- Validate network throughput before performing active duplication.
- Protect encryption wallets and password files during migration.
- Immediately back up the cloned database after duplication.



- Use NOFILENAMECHECK cautiously—only when directories are isolated.
- Schedule routine duplication drills to validate disaster recovery readiness.

---

## 9.15 Summary

Advanced recovery operations extend the DBA's control beyond basic restoration, enabling efficient replication, migration, and testing of entire databases or containers.

In this chapter, you learned how to:

- Perform **active** and **backup-based duplication** using RMAN.
- Create **test and QA databases** from backups.
- Handle **encrypted and compressed backups** during duplication.
- Duplicate **multitenant CDB and PDB** databases.
- Execute **cross-platform conversions** and hybrid recoveries.

These techniques allow organizations to safeguard data while increasing agility for development, testing, and compliance.

The next chapter—**Recovery Troubleshooting and Diagnostics**—focuses on analyzing RMAN logs, interpreting common error codes, and applying diagnostic utilities to identify and resolve recovery failures efficiently.




---

## PART V – EXPERT TECHNIQUES AND TROUBLESHOOTING

---

### Chapter 10 – Recovery Troubleshooting and Diagnostics

---

#### 10.1 Introduction

Even the most carefully planned recovery process can encounter unexpected failures—missing files, corrupt backups, incomplete metadata, or environmental issues.

An Oracle DBA must be able to **detect, diagnose, and resolve** recovery problems quickly, minimizing downtime and maintaining data integrity.

This chapter explains how to analyze RMAN logs, alert logs, and system views to determine the cause of a failed recovery, how to interpret common error codes, and which utilities to use for verification and validation.

The goal is to make troubleshooting a **structured and repeatable process** rather than a series of guesses during a crisis.

---

#### 10.2 Common Causes of Recovery Failure

Most recovery issues fall into a few major categories:

Category	Description
<b>Missing Backups</b>	Backup files deleted, renamed, or expired before recovery.
<b>Corrupt Backup Pieces</b>	Storage corruption or incomplete transfer of backup files.
<b>Incorrect File Paths</b>	Mismatched directory structure between source and destination.
<b>Control File Inconsistency</b>	Control file references missing or outdated backups.
<b>Archivelog Gaps</b>	Missing archived redo prevents full recovery.



Category	Description
Parameter Mismatch	Wrong DBID, incompatible incarnation, or mismatched SPFILE settings.
Hardware or OS Constraints	Disk space exhaustion or permission errors during restore.

Identifying the category early allows you to target the appropriate corrective action.

---

## 10.3 Understanding RMAN Log Files

Every RMAN operation generates a detailed log that records command execution, backup and restore progress, and errors.

Reviewing the log is the first step in any troubleshooting process.

### Log Locations

- **Windows** – `\app\oracle\diag\rdbms\<DB_NAME>\trace\rman_<date>.log`
- **Linux/Unix** – `$ORACLE_BASE/diag/rdbms/<DB_NAME>/<SID>/trace/rman_<date>.log`

### Important Sections

1. **Channel Allocation** – Confirms devices and paths used.
2. **Piece Creation and Validation** – Displays backupset names, timestamps, and checksums.
3. **RMAN Errors (RMAN-#####)** – Contain specific failure messages.
4. **Oracle Errors (ORA-#####)** – Indicate underlying database or OS issues.

### Example Log Fragment

channel ORA\_DISK\_1: starting piece 1 at 2025-11-10:01:00:23

RMAN-03009: failure of restore command on ORA\_DISK\_1 channel at 11/10/2025 01:10:42

ORA-19505: failed to identify file /backup/DF\_01.bkp

ORA-27037: unable to obtain file status

**Diagnosis:** Backup piece not found or inaccessible. Verify directory path and file permissions.



## 10.4 Using the Database Alert Log

The **alert log** is a chronological record of database events and errors.

It complements the RMAN log by providing instance-level information, such as startup messages, checkpoint operations, and block corruption detection.

### View Alert Log via SQL

```
SELECT MESSAGE_TEXT FROM V$DIAG_ALERT_EXT
WHERE ORIGINATING_COMPONENT='rdbms'
ORDER BY ORIGINATING_TIMESTAMP DESC;
```

### Common Messages During Recovery

- “*Media recovery start*” / “*Media recovery complete*” – Recovery succeeded.
- “*Checkpoint not complete*” – Redo not fully applied; check missing logs.
- “*Control file sequential read error*” – Possible corruption or wrong control file.

Always review the alert log before reopening a recovered database.

---

## 10.5 Diagnostic Data Viewer (ADRCI)

The **Automatic Diagnostic Repository Command Interpreter (ADRCI)** is a command-line tool for managing diagnostic data.

### Key Commands

adrci> show homes

adrci> set home diag/rdbms/prod/prod

adrci> show alert -tail 50

adrci> show incident

adrci> ips pack incident 12345 in /u01/incident\_zips

ADRCI automatically captures incidents related to recovery errors.

Packaging these for Oracle Support accelerates problem resolution during escalation.



## 10.6 Troubleshooting Backup Failures

Backup failures often surface during validation or subsequent recovery.

The following checks identify most root causes:

### 1. Verify Channel Configuration

SHOW ALL;

### 2. Crosscheck Backup Metadata

CROSSCHECK BACKUP;

REPORT OBSOLETE;

### 3. Validate Backup Sets

RESTORE DATABASE VALIDATE;

### 4. Confirm FRA or Disk Capacity

SELECT SPACE\_LIMIT, SPACE\_USED FROM V\$RECOVERY\_FILE\_DEST;

### 5. Check Operating System Permissions

Ensure Oracle has read/write access to backup directories.

#### Common Error Example

ORA-19504: failed to create file /backup/DF\_01.bkp

ORA-27040: file create error, unable to create file

**Resolution:** Verify directory permissions and free space.

## 10.7 Troubleshooting Restore and Recovery Failures

### Scenario 1 – Backup Piece Not Found

#### Symptom

ORA-19505: failed to identify file

#### Resolution

- Check directory paths in the control file:  
LIST BACKUP OF DATABASE;



- If paths changed, use RMAN cataloging:

CATALOG START WITH '/new\_backup\_path/';

## Scenario 2 – Missing Archivelogs

### Symptom

ORA-00308: cannot open archived log

### Resolution

- Identify missing logs:  
LIST ARCHIVELOG ALL;
- Restore from backup or request regeneration from standby (if available).
- If unavailable, perform incomplete recovery using SET UNTIL SCN.

## Scenario 3 – Control File Mismatch

### Symptom

ORA-01110: data file mismatch with control file

### Resolution

- Ensure the control file corresponds to the restored backup set.
- Restore control file from same backup date as datafiles.

## Scenario 4 – Corrupted Backup Piece

### Symptom

ORA-19566: exceeded limit for datafile copies

or CRC mismatch during restore.

### Resolution

- Validate backup integrity before restore:  
VALIDATE BACKUPSET <handle>;
- Recreate or replace the corrupted backup.



## 10.8 Troubleshooting Flashback Operations

Error	Likely Cause	Corrective Action
<b>ORA-38760: Flashback database not started</b>	Flashback logs missing or FRA disabled.	Re-enable flashback and regenerate restore point.
<b>ORA-38301: Object not in recycle bin</b>	Object purged or recycle bin disabled.	Restore from RMAN backup or Data Pump export.
<b>ORA-38714: Instance not mounted</b>	Flashback Database requires MOUNT state.	Restart in MOUNT mode before issuing command.
<b>ORA-38706: Cannot turn on flashback</b>	Database not in ARCHIVELOG mode.	Enable ARCHIVELOG and retry.

Most flashback failures result from configuration gaps rather than software errors.

---

## 10.9 Using Dynamic Performance Views

Oracle exposes extensive recovery diagnostics through **V\$** and **RC\_** views.

### Essential Views for Recovery

View	Purpose
<b>V\$BACKUP_SET</b>	Lists backup sets and completion times.
<b>V\$BACKUP_PIECE_DETAILS</b>	Displays file handles and statuses.
<b>V\$DATABASE_BLOCK_CORRUPTION</b>	Identifies corrupted data blocks.
<b>V\$RECOVERY_LOG</b>	Shows redo logs needed for recovery.
<b>V\$RECOVER_FILE</b>	Lists datafiles requiring media recovery.
<b>V\$FLASH_RECOVERY_AREA_USAGE</b>	Reports FRA usage by file type.
<b>RC_BACKUP_SET, RC_DATABASE</b>	Catalog-based equivalents for centralized reporting.



## Example Query

```
SELECT file#, error, block#, blocks
```

```
FROM v$database_block_corruption;
```

If corruption is detected, perform BLOCKRECOVER or restore affected datafiles from backup.

---

## 10.10 Data Recovery Advisor (LIST, ADVISE, REPAIR)

The **Data Recovery Advisor** (DRA) provides automated analysis and recommendations for recovery issues through RMAN.

### Usage Example

```
RMAN> LIST FAILURE;
```

```
RMAN> ADVISE FAILURE;
```

```
RMAN> REPAIR FAILURE;
```

### Features

- Detects missing or corrupted files automatically.
- Suggests repair options (restore, recover, or block repair).
- Can execute repairs interactively or automatically.

### Example Output

Failure ID: 2

Type: Corruption

Status: Open

Repair Script: Restores datafile 5 from backup and recovers it.

Although many DBAs prefer manual recovery, DRA is an excellent diagnostic starting point.

---

## 10.11 Interpreting Common Error Codes

### RMAN Errors



Error	Description	Resolution
<b>RMAN-06004</b>	Invalid backup piece handle.	Re-catalog or re-copy backup.
<b>RMAN-06025</b>	No backup of the specified file found.	Verify backup retention and control file.
<b>RMAN-08120</b>	Warning: archived log not deleted, still needed for recovery.	Adjust retention policy.
<b>RMAN-10035</b>	Exception during job execution.	Review channel errors; check disk I/O.

## ORA Errors

Error	Description	Resolution
<b>ORA-01113</b>	File needs media recovery.	Execute RECOVER DATAFILE.
<b>ORA-01547</b>	Warning: RECOVER succeeded but RESETLOGS required.	Open database with RESETLOGS.
<b>ORA-19809</b>	FRA full.	Delete obsolete backups or increase FRA size.
<b>ORA-00205</b>	Error identifying control file.	Restore or recreate control file.

Maintaining a quick-reference list of these errors accelerates diagnosis during emergencies.

## 10.12 Verification and Post-Recovery Checks

After resolving recovery failures or completing successful restoration, verify integrity and log results.

### 1. Confirm Open Mode



SELECT OPEN\_MODE FROM V\$DATABASE;

## 2. Validate Datafiles

RMAN> VALIDATE DATABASE;

## 3. Check Redo Application

SELECT FILE#, ERROR FROM V\$RECOVER\_FILE;

## 4. Review FRA and Control File Status

SELECT SPACE\_USED\_PERCENT FROM V\$RECOVERY\_FILE\_DEST;

SHOW PARAMETER control\_files;

## 5. Back Up Immediately

Take a new full backup to establish a reliable post-recovery baseline.

## 10.13 Automation and Logging Standards

Consistent log management reduces troubleshooting time dramatically.

### Recommended Practices

- Direct RMAN output to timestamped log files in /logs/rman/.
- Implement automated alerting (e.g., email or script) for ORA- and RMAN- errors.
- Maintain a standard folder for alert, trace, and diagnostic logs.
- Archive recovery scripts with logs for future reference.
- Review RMAN logs daily, especially after scheduled jobs.

Automation ensures failures are detected proactively rather than during the next recovery attempt.

## 10.14 Escalation and Support Readiness

When internal troubleshooting is exhausted:

### 1. Collect Evidence

- RMAN log files.
- Alert logs.



- Diagnostic package (adrci ips pack).
- Output of LIST FAILURE and REPORT SCHEMA.

## 2. Validate Environment

- Confirm Oracle version and PSU level.
- Capture SHOW ALL RMAN configuration.

## 3. Engage Oracle Support

- Provide zipped diagnostic bundle.
- Include a concise description of the incident and steps already taken.

Well-prepared diagnostic data shortens resolution time and demonstrates professional incident handling.

---

## 10.15 Best Practices for Continuous Improvement

- Perform **post-mortem reviews** of every recovery event.
- Record elapsed times, issues, and corrective actions in a recovery logbook.
- Periodically **rehearse recovery scenarios** to maintain team readiness.
- Regularly validate RMAN backup paths and FRA space.
- Incorporate improvements into written procedures and automation scripts.

A culture of continuous improvement ensures that each incident strengthens the overall recovery strategy.

---

## 10.16 Summary

Troubleshooting and diagnostics are essential skills for every Oracle DBA.

This chapter provided the framework to analyze, interpret, and resolve recovery failures quickly and confidently:

- Reviewed RMAN and alert logs to identify root causes.
- Used ADRCI, Data Recovery Advisor, and diagnostic views for analysis.
- Interpreted common error messages and applied standard fixes.



- Emphasized post-recovery validation and automated monitoring.

Effective troubleshooting transforms recovery from a crisis event into a **controlled, documented process**.

The next chapter—**Performance, Optimization, and Validation**—will explore how to tune backup and recovery throughput, manage retention policies efficiently, and ensure recovery readiness through performance testing and proactive monitoring.



---

# Chapter 11 – Performance, Optimization, and Validation

---

## 11.1 Introduction

Performance optimization in Oracle backup and recovery is not solely about running backups faster—it's about achieving **predictable throughput**, **minimal downtime**, and **verified recoverability**.

A well-optimized recovery system balances CPU, memory, disk I/O, and network resources to ensure that both backup and restore operations consistently meet Recovery Time Objective (RTO) and Recovery Point Objective (RPO) targets.

This chapter explains how to analyze, tune, and validate RMAN performance on on-premise disk-based systems. It also covers how to confirm the validity and efficiency of your backup configurations through regular testing and measurement.

---

## 11.2 Key Performance Concepts

### Throughput

The rate at which data is backed up or restored, usually measured in megabytes per second (MB/s).

Backup throughput depends on disk I/O speed, channel parallelism, CPU availability, and compression level.

### Concurrency

The number of channels or processes working simultaneously. Excessive concurrency can lead to resource contention, while insufficient concurrency underutilizes hardware.

### Checkpoint Frequency

Affects recovery time. More frequent checkpoints reduce redo application during recovery but can slow runtime performance.

### Recovery Validation Time

The time required to test restore operations. This is an often-overlooked but critical metric for verifying recoverability.



Understanding these metrics allows the DBA to tune both backup and restore strategies effectively.

---

## 11.3 Factors Influencing Backup Performance

1. **Disk I/O Bandwidth** – The primary determinant of speed. Use high-performance storage for backup destinations.
2. **RMAN Channels** – Each channel performs a single I/O stream; more channels usually mean faster backups.
3. **Parallelism Configuration** – Balance channel count against CPU and I/O saturation.
4. **Compression** – Reduces file size but consumes CPU; test to find the optimal balance.
5. **Network Latency** – Significant only for remote backups or recovery over NFS.
6. **File Layout** – Separate datafiles, redo, and FRA across different disks or mount points.
7. **Block Size and Section Size** – Affect read and write efficiency; large sections can improve performance on sequential I/O.

---

## 11.4 Measuring RMAN Performance

Oracle provides several dynamic performance views and commands for measuring backup efficiency.

### RMAN Timing Output

RMAN automatically reports time elapsed for each step in the backup or restore sequence.

Example:

channel ORA\_DISK\_1: backup set complete, elapsed time: 00:02:47

### V\$SESSION\_LONGOPS

Displays ongoing backup and recovery tasks:

```
SELECT SID, OPNAME, SOFAR, TOTALWORK,
       ROUND(SOFAR/TOTALWORK*100,2) AS "%COMPLETE"
  FROM V$SESSION_LONGOPS
```



WHERE OPNAME LIKE 'RMAN%';

### **V\$BACKUP\_ASYNC\_IO**

Shows I/O throughput per file and channel:

```
SELECT FILE#, IO_COUNT, IO_BYTES/1024/1024 MB_PROCESSED,
```

```
OPEN_TIME, CLOSE_TIME FROM V$BACKUP_ASYNC_IO;
```

By tracking these metrics across multiple runs, the DBA can quantify performance improvements and detect bottlenecks.

## 11.5 Parallelism and Channel Optimization

RMAN can perform parallel operations by allocating multiple disk channels.

Each channel corresponds to a server session that performs read/write activity independently.

### **Configuring Parallel Channels**

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 4;
```

```
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/%U';
```

### **Tuning Guidelines**

- Begin with one channel per CPU core or storage path.
- Measure performance gain for each additional channel.
- Avoid setting more channels than the system can handle concurrently.
- When restoring, use the same number of channels as used during backup for maximum efficiency.

### **Example: Optimized Backup Block**

```
RUN {
```

```
ALLOCATE CHANNEL c1 DEVICE TYPE DISK;
```

```
ALLOCATE CHANNEL c2 DEVICE TYPE DISK;
```

```
ALLOCATE CHANNEL c3 DEVICE TYPE DISK;
```

```
BACKUP DATABASE PLUS ARCHIVELOG;
```

```
RELEASE CHANNEL c1;
```



```

RELEASE CHANNEL c2;
RELEASE CHANNEL c3;
}

```

Parallel execution reduces total elapsed time significantly but increases CPU and I/O load; monitor system utilization closely during tuning.

---

## 11.6 Compression and Deduplication

### Compression

RMAN supports three main compression algorithms:

Algorithm	Description	Typical Use
<b>BASIC</b>	Default; moderate compression and low CPU usage.	General-purpose disk backups.
<b>LOW</b>	Minimal compression for fastest throughput.	Systems with ample disk space.
<b>HIGH</b>	Maximum compression, high CPU usage.	Archival backups, constrained storage environments.

Enable compression globally:

```
CONFIGURE COMPRESSION ALGORITHM 'BASIC';
```

```
CONFIGURE BACKUP OPTIMIZATION ON;
```

### Deduplication

When using network shares or deduplicating storage appliances, store RMAN backups in backup sets (not image copies).

This format enables file-level deduplication at the storage layer.

---



## 11.7 Incremental Merge Optimization

Incrementally updated backups provide a rolling full backup without requiring regular full restore operations.

### Process Overview

1. Perform a Level 0 image copy.
2. Apply Level 1 incremental backups nightly using the RECOVER COPY OF DATABASE command.

Example:

```
RUN {
  BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE TAG 'ROLLING_FULL';
  BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG 'ROLLING_FULL' DATABASE;
}
```

### Benefits

- Faster recovery time (copy already exists as full image).
- Smaller nightly backup window.
- Reduced I/O load compared to full backup schedules.

## 11.8 Block Change Tracking

Block Change Tracking (BCT) records which blocks have changed since the last backup, dramatically improving incremental backup speed.

### Enable BCT

```
ALTER DATABASE ENABLE BLOCK CHANGE TRACKING
  USING FILE '/u01/app/oracle/oradata/PROD/bct.trk';
```

### Verify BCT Status

```
SELECT STATUS, FILENAME FROM V$BLOCK_CHANGE_TRACKING;
```

### Best Practices

- Place the tracking file on fast, dedicated storage.



- Monitor file size growth; typically a few hundred MB for large databases.
- Re-enable BCT after control file recreation or restore.

---

## 11.9 Restore and Recovery Performance

The real test of backup optimization is **restore speed**, not just backup speed.

### Tips for Faster Restores

- Match parallelism between backup and restore operations.
- Keep control-file autobackups on local storage for quick access.
- Use multiple channels to restore in parallel.
- Validate restoration paths before recovery to reduce downtime.
- Ensure FRA and backup locations are on separate spindles.
- For partial recoveries, restore only necessary datafiles.

Example restore:

```
RUN {
  ALLOCATE CHANNEL c1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL c2 DEVICE TYPE DISK;
  RESTORE DATABASE;
  RECOVER DATABASE;
}
```

The goal is not only recovery success but **predictable and measurable recovery time**.

---

## 11.10 Monitoring and Analyzing Backup History

Oracle provides several views for analyzing backup statistics and identifying long-running jobs.

### V\$RMAN\_BACKUP\_JOB\_DETAILS

Displays start and end times for each backup job.

SELECT SESSION\_KEY, INPUT\_TYPE, STATUS,



```
START_TIME, END_TIME,
OUTPUT_BYTES/1024/1024 "MB_WRITTEN"
FROM V$RMAN_BACKUP_JOB_DETAILS
ORDER BY START_TIME DESC;
```

### **V\$BACKUP\_SET\_DETAILS**

Provides size and elapsed time of backup sets.

```
SELECT BACKUP_TYPE, INCREMENTAL_LEVEL, BYTES/1024/1024 MB, ELAPSED_SECONDS
FROM V$BACKUP_SET_DETAILS;
```

Analyzing these metrics helps establish baselines and detect performance degradation over time.

## 11.11 Validating Backups

Backup validation confirms that backups can be read and restored successfully without actually writing the datafiles.

### **Validation Commands**

- Validate the entire database:

```
VALIDATE DATABASE;
```

- Validate a specific tablespace or datafile:

```
VALIDATE TABLESPACE users;
```

```
VALIDATE DATAFILE 5;
```

- Validate backup pieces:

```
VALIDATE BACKUPSET 45;
```

Validation should be run weekly or after major configuration changes.

Errors indicate missing or corrupted files that must be replaced before they are needed in a real recovery.



## 11.12 Periodic Test Restores

A verified backup is only as good as its ability to restore.

Test restores ensure that your recovery process works end-to-end under realistic conditions.

### Recommended Frequency

- Quarterly for small systems.
- Monthly for mission-critical environments.
- After any major OS, patch, or configuration change.

### Test Restore Process

1. Restore control file to a test location.
2. Restore database to a temporary instance.
3. Recover using archived logs.
4. Open database with RESETLOGS.
5. Verify application connectivity.

This process validates not only RMAN backups but also storage performance, network speed, and operational readiness.

## 11.13 Proactive Recovery Readiness Tests

To ensure recoverability, implement routine checks in your DBA maintenance cycle.

### Checklist

- Verify that all required files are backed up.
- Confirm no “EXPIRED” backups in RMAN.
- Ensure latest control file backup is available.
- Monitor redo and archive log generation rates.
- Confirm FRA free space > 20% of configured size.
- Validate alert logs for RMAN or flashback errors.

Automation of these checks through SQL or shell scripts enhances operational reliability.



## 11.14 Retention Policy Optimization

Retention policies balance storage consumption with recoverability.

### Common Configurations

- **Recovery Window:**

CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 14 DAYS;

- **Redundancy:**

CONFIGURE RETENTION POLICY TO REDUNDANCY 2;

### Optimization Tips

- Use the recovery window for production environments to guarantee a time-based recovery point.
- Apply redundancy for development or testing systems where frequent refreshes occur.
- Regularly purge obsolete backups:

DELETE NOPROMPT OBSOLETE;

Efficient retention management ensures consistent FRA usage and prevents backup failures due to insufficient space.

---

## 11.15 Automating Performance Reports

Automation ensures that backup and recovery performance remains visible to DBAs and management.

### Example Daily Report Script

```
sqlplus -s / as sysdba <<EOF
SET LINES 200
SET PAGESIZE 200
SPOOL /logs/rman_performance_$(date +%F).txt
SELECT SESSION_KEY, INPUT_TYPE, STATUS,
ELAPSED_SECONDS, OUTPUT_BYTES/1024/1024 MB_WRITTEN
```



```
FROM V$RMAN_BACKUP_JOB_DETAILS
```

```
WHERE START_TIME > SYSDATE-1;
```

```
SPOOL OFF
```

```
EXIT
```

```
EOF
```

Email or archive the generated reports for trending analysis.

Graphing throughput over time helps visualize performance improvements and capacity planning.

## 11.16 Best Practices for Optimal Backup and Recovery Performance

1. Separate datafiles, redo logs, and backups onto different disks or storage paths.
2. Use **incremental merge** and **block change tracking** for large databases.
3. Maintain at least two control file copies and test RESTORE CONTROLFILE regularly.
4. Schedule backups during off-peak hours to reduce resource contention.
5. Use compressed backups for long-term storage and uncompressed for fast recovery.
6. Validate backup pieces weekly and after patching.
7. Automate reporting and alerting on RMAN job status.
8. Document observed throughput, RTO, and RPO performance metrics.

## 11.17 Summary

Optimized backup and recovery are cornerstones of high-availability Oracle environments.

This chapter demonstrated how to:

- Measure and improve backup and restore performance.
- Use parallelism, compression, and block change tracking effectively.
- Validate backups and perform regular test restores.
- Automate monitoring and retention management for reliability and space efficiency.

By applying these techniques, DBAs can ensure their recovery processes not only work—but work fast, predictably, and verifiably.



The next chapter, **Periodic Testing and Lab Drills**, concludes the recovery guide by providing structured exercises and simulations to test every aspect of the recovery process under controlled conditions.



## Chapter 12 – Periodic Testing and Lab Drills

### 12.1 Introduction

Backups without testing are **assumptions, not guarantees**.

A database recovery strategy is only as strong as its last successful test.

Routine, documented recovery tests confirm that backups are complete, procedures are accurate, and personnel are prepared to act confidently in an emergency.

This chapter provides a framework for **structured recovery testing**, also called *lab drills*, that simulate real-world failure events in a controlled environment.

The goal is to make database recovery predictable, repeatable, and measurable.

### 12.2 Objectives of Recovery Testing

Recovery testing validates:

- **Backup Integrity** – Ensures backups are readable and complete.
- **Procedural Accuracy** – Confirms recovery scripts and documentation are current.
- **Performance Readiness** – Measures restore and recovery speed against RTO goals.
- **Staff Competence** – Verifies that the DBA team can execute recovery without hesitation.
- **Configuration Reliability** – Confirms all required parameters, paths, and privileges exist in the recovery environment.

Each successful test increases confidence that the production system can be restored within defined business continuity targets.

### 12.3 Recommended Testing Frequency

Environment Type	Test Frequency	Example Scope
Production	Quarterly (minimum)	Simulated datafile loss, incomplete recovery, flashback restore.



Environment Type	Test Frequency	Example Scope
Development / QA	Monthly	Full restore from backup, duplication to test environment.
Disaster Recovery Site	Semi-Annually	Full site restore using off-site backups.
After Major Change	Immediately after patch, upgrade, or storage migration.	Verify control file and SPFILE integrity.

Testing frequency should be documented in the organization's data protection policy and reviewed annually.

---

## 12.4 Preparing a Test Environment

A well-prepared test environment isolates the recovery process from production operations while maintaining data fidelity.

### Requirements

- Access to valid and recent RMAN backups.
- Adequate disk space to restore the full database.
- A separate Oracle Home and instance name (e.g., TESTRECOV).
- Optional dedicated network segment for safety.
- Properly configured initialization parameters, including unique DB\_NAME and CONTROL\_FILES paths.

### Best Practices

- Mirror production directory structures for consistency.
- Use the same Oracle version and patch level as production.
- Disable automated jobs or replication features to prevent interference.
- Record all test parameters (SCN, timestamp, backup tag, host name) before starting.

---



## 12.5 Categories of Recovery Tests

Recovery testing should cover the full range of potential failures:

1. **Instance Recovery Test** – Validate automatic recovery after simulated crash.
2. **Media Recovery Test** – Restore and recover a lost datafile or tablespace.
3. **Control File Loss Test** – Restore from autobackup or manual backup.
4. **SPFILE Recovery Test** – Recreate the server parameter file from backup.
5. **Incomplete Recovery Test** – Recover to a specific SCN or timestamp.
6. **Flashback Recovery Test** – Perform Flashback Database or Table.
7. **Block Media Recovery Test** – Repair corrupted data blocks using RMAN.
8. **Full Database Restore Test** – Restore the entire database from a recent backup.
9. **Duplicate Database Test** – Rebuild a test or QA environment from production backups.

Each test type exercises a different portion of the recovery framework, ensuring complete coverage.

---

## 12.6 Simulating Common Failure Scenarios

### Scenario 1 – Lost Datafile

1. Delete or rename a datafile at the OS level.
2. Attempt to start or open the database (expect failure).
3. Restore and recover using RMAN:

RESTORE DATAFILE 6;

RECOVER DATAFILE 6;

ALTER DATABASE OPEN;

4. Validate recovered data.

### Scenario 2 – Corrupted Backup Piece

1. Intentionally remove a backup file.
2. Run a simulated restore to detect missing pieces:



RESTORE DATABASE VALIDATE;

3. Observe error handling and re-catalog missing backups.

### Scenario 3 – User Error

1. Delete or modify a table's data.
2. Recover using RMAN table recovery or Flashback Table.
3. Validate consistency post-recovery.

### Scenario 4 – Control File Loss

1. Remove all control files.
2. Attempt startup (expect ORA-00205).
3. Restore from autobackup:

STARTUP NOMOUNT;

RESTORE CONTROLFILE FROM AUTOBACKUP;

ALTER DATABASE MOUNT;

RECOVER DATABASE;

ALTER DATABASE OPEN RESETLOGS;

### Scenario 5 – Total Media Failure

1. Simulate full disk loss.
2. Rebuild entire database from RMAN backups.
3. Measure restore and recovery time to verify RTO compliance.

## 12.7 Measuring Recovery Performance

Document measurable metrics for each test:

Metric	Description	Example
<b>RTO (Recovery Time Objective)</b>	Maximum acceptable downtime.	2 hours



Metric	Description	Example
<b>Actual Recovery Time</b>	Time measured from failure detection to database open.	1 hr 42 min
<b>RPO (Recovery Point Objective)</b>	Maximum acceptable data loss (usually in minutes).	15 min
<b>Data Volume Restored</b>	Total MB restored from backup.	350,000 MB
<b>Throughput</b>	Average MB/sec during restore.	285 MB/s
<b>Operator / DBA</b>	Person executing the recovery test.	J. Smith

Regularly updating these statistics provides a historical record of improvement and consistency.

---

## 12.8 Establishing a Test Procedure Template

Each recovery test should be documented using a standardized format.

### Sample Template

#### 1. Objective

Describe the failure scenario being tested.

#### 2. Preconditions

Identify the backup sets, timestamps, and environment setup.

#### 3. Test Steps

List every command executed (SQL, RMAN, OS).

#### 4. Observations

Record system messages, elapsed times, and log locations.

#### 5. Results

State whether the test met success criteria.

#### 6. Lessons Learned

List procedural or configuration improvements identified.



## 7. Approval

Document review by lead DBA or manager.

By enforcing standard templates, the organization ensures that recovery tests are reproducible and auditable.

---

## 12.9 Testing Flashback Operations

Flashback operations can often recover from logical errors faster than full restores, but they must be validated regularly.

### Example Flashback Test

1. Create a guaranteed restore point:

```
CREATE RESTORE POINT pre_update GUARANTEE FLASHBACK DATABASE;
```

2. Simulate a schema change:

```
ALTER TABLE hr.employees ADD (temp_col NUMBER);
```

3. Flash back the database:

```
SHUTDOWN IMMEDIATE;
```

```
STARTUP MOUNT;
```

```
FLASHBACK DATABASE TO RESTORE POINT pre_update;
```

```
ALTER DATABASE OPEN RESETLOGS;
```

4. Verify the schema no longer includes temp\_col.

5. Drop the restore point:

```
DROP RESTORE POINT pre_update;
```

This verifies that Flashback Database can safely restore to a consistent earlier state without restoring from backup.

---

## 12.10 Recovery Validation with Auxiliary Instances

Testing recovery in an **auxiliary instance** ensures the procedure can be executed independently of the production database.



## Procedure

1. Create a new instance using a minimal PFILE.
2. Restore control file and datafiles from backup to alternate paths.
3. Recover until the last archived log.
4. Open with RESETLOGS.
5. Verify recovered database consistency and size.

This method is particularly useful for validating disaster recovery backups without interrupting live operations.

---

## 12.11 Recovery Checklist and Audit Log

Each recovery test should produce a written audit trail.

Use a standardized checklist for every test event.

### Recovery Checklist

Task	Completed	Verified By
Backups validated prior to test	<input checked="" type="checkbox"/>	DBA Lead
Required backups accessible	<input checked="" type="checkbox"/>	
RMAN scripts updated and version-controlled	<input checked="" type="checkbox"/>	
Control file autobackup verified	<input checked="" type="checkbox"/>	
Restore executed successfully	<input checked="" type="checkbox"/>	
Database opened with RESETLOGS	<input checked="" type="checkbox"/>	
Post-recovery backup created	<input checked="" type="checkbox"/>	
Results documented and archived	<input checked="" type="checkbox"/>	

Maintaining signed checklists ensures accountability and demonstrates compliance with audit or regulatory standards.



## 12.12 Post-Test Review

After every recovery exercise, conduct a **post-test review** to identify improvements.

This review should include:

1. **Timing Analysis** – Compare actual recovery time to RTO goals.
2. **Procedure Review** – Identify unclear steps or missing details in documentation.
3. **Error Summary** – List encountered errors and their resolutions.
4. **Configuration Adjustments** – Modify RMAN, FRA, or retention settings based on findings.
5. **Personnel Assessment** – Evaluate operator readiness and training needs.
6. **Corrective Actions** – Schedule changes and re-tests for any failed components.

Documenting reviews turns each exercise into a continuous improvement opportunity.

---

## 12.13 Reporting and Continuous Improvement

All testing results should be summarized in a recurring **Recovery Readiness Report**.

The report should include:

- Date and type of test performed.
- Duration of recovery and overall RTO performance.
- Validation results and verification method.
- Identified deficiencies or risks.
- Recommended corrective actions and target completion dates.

These reports are essential for executive review and demonstrate operational reliability to auditors or regulatory bodies.

---

## 12.14 Integrating Recovery Testing into Change Management

To ensure recovery preparedness is always current:

- Require a **successful recovery test** before major upgrades or patch applications.
- Link recovery test documentation to official **change tickets**.



- Mandate **DBA sign-off** confirming that backups are valid prior to go-live events.
- Archive test records in the organization's configuration management system.

Embedding recovery drills into the change lifecycle institutionalizes recoverability as a mandatory component of production readiness.

---

## 12.15 Example Lab Exercise: Full Database Recovery Simulation

**Objective:** Validate the ability to restore the entire database after catastrophic storage loss.

### Steps

1. Simulate data loss by unmounting or renaming datafile directories.
2. Mount a test instance with identical parameters.
3. Restore control file and datafiles from RMAN backup.
4. Apply all archived redo logs.
5. Open database with RESETLOGS.
6. Validate by connecting as application user and running transaction tests.
7. Measure recovery time and compare to RTO threshold.
8. Perform a full backup of the recovered database.
9. Complete documentation and sign-off checklist.

**Expected Outcome:** Database restored successfully with no data loss beyond RPO target.

---

## 12.16 Common Mistakes During Recovery Testing

- Using outdated RMAN scripts or hard-coded paths.
- Forgetting to re-enable block change tracking or flashback after testing.
- Overwriting live database directories due to incorrect environment variables.
- Failing to reset archive log destinations after simulation.
- Skipping validation of restored data or indexes.
- Neglecting to perform a new full backup after successful recovery.



Avoiding these mistakes prevents accidental disruptions and ensures that testing reflects real production behavior.

---

## 12.17 Continuous Learning and Team Preparedness

Recovery success depends as much on people as on technology.

Developing recovery expertise should be part of every DBA's professional growth plan.

### Recommendations

- Rotate team members through hands-on recovery drills.
- Maintain up-to-date runbooks and checklists.
- Conduct internal "recovery contests" or timed exercises to reinforce proficiency.
- Document lessons learned and incorporate them into the next revision of the Recovery Guide.
- Encourage peer reviews of recovery documentation.

A confident, well-trained DBA team is the best guarantee of database resilience.

---

## 12.18 Summary

Recovery testing transforms backup management from a passive safety net into an active, proven discipline.

In this chapter, you learned to:

- Plan and schedule recurring recovery drills.
- Prepare test environments that replicate production conditions.
- Simulate a variety of failures to verify full and partial recovery.
- Measure and document recovery performance.
- Integrate testing into organizational change management and continuous improvement processes.

With these practices in place, your organization can ensure that no recovery scenario—no matter how critical—will ever depend on improvisation.

Instead, it will be guided by tested procedures, skilled personnel, and validated backups.



The appendices that follow contain **command references**, **automation scripts**, and **lab scenario templates** for implementing these exercises in both Windows and Linux environments.



---

## PART VI – APPENDICES

---

### Appendix A – RMAN Command Reference

---

#### A.1 Introduction

This appendix serves as a quick reference to the most frequently used Recovery Manager (RMAN) commands for on-premise Oracle 12c–21c environments.

All examples assume disk-based backups, local file systems, and databases operating in ARCHIVELOG mode.

---

#### A.2 Connection Commands

```
rman TARGET /
```

```
rman TARGET sys/password@ORCL
```

```
rman TARGET sys/password@ORCL CATALOG rmanuser/rmanpw@CATDB
```

```
rman TARGET / AUXILIARY sys/password@AUXDB
```

---

#### A.3 Configuration Commands

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;
```

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;
```

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 4;
```

```
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/%U';
```

```
CONFIGURE ENCRYPTION FOR DATABASE ON;
```

```
CONFIGURE COMPRESSION ALGORITHM 'BASIC';
```

```
SHOW ALL;
```



---

## A.4 Backup Commands

### **Full Backup**

```
BACKUP DATABASE TAG 'FULL_DB_BACKUP';
```

### **Incremental Backup**

```
BACKUP INCREMENTAL LEVEL 1 DATABASE TAG 'INCR_L1';
```

### **Archived Redo Backup**

```
BACKUP ARCHIVELOG ALL DELETE INPUT TAG 'ARCHIVE_BACKUP';
```

### **Control File and SPFILE Backup**

```
BACKUP CURRENT CONTROLFILE;
```

```
BACKUP SPFILE;
```

### **Validate Backup**

```
VALIDATE DATABASE;
```

---

## A.5 Restore and Recovery Commands

### **Full Database Recovery**

```
STARTUP MOUNT;
```

```
RESTORE DATABASE;
```

```
RECOVER DATABASE;
```

```
ALTER DATABASE OPEN;
```

### **Incomplete Recovery**

```
RUN {
```

```
  SET UNTIL TIME "TO_DATE('2025-11-11 10:00:00','YYYY-MM-DD HH24:MI:SS')";
```

```
  RESTORE DATABASE;
```

```
  RECOVER DATABASE;
```

```
}
```



ALTER DATABASE OPEN RESETLOGS;

### **Control File Restoration**

STARTUP NOMOUNT;

RESTORE CONTROLFILE FROM AUTOBACKUP;

ALTER DATABASE MOUNT;

### **Block Media Recovery**

BLOCKRECOVER DATAFILE 6 BLOCK 100 TO 120;

## **A.6 Maintenance and Validation**

CROSSCHECK BACKUP;

DELETE NOPROMPT OBSOLETE;

LIST BACKUP SUMMARY;

REPORT NEED BACKUP;

VALIDATE BACKUPSET ALL;

## **A.7 Duplication and Cloning**

### **Active Duplication**

DUPLICATE TARGET DATABASE TO TESTDB FROM ACTIVE DATABASE NOFILENAMECHECK;

### **Backup-Based Duplication**

DUPLICATE DATABASE TO TESTDB BACKUP LOCATION '/backup/rman' NOFILENAMECHECK;

## **A.8 Flashback Commands**

FLASHBACK DATABASE TO RESTORE POINT pre\_upgrade;

FLASHBACK TABLE hr.employees TO TIMESTAMP (SYSTIMESTAMP - INTERVAL '5' MINUTE);

FLASHBACK TABLE hr.employees TO BEFORE DROP;

CREATE RESTORE POINT pre\_patch GUARANTEE FLASHBACK DATABASE;



---

## A.9 Diagnostic Commands

```
LIST FAILURE;  
ADVISE FAILURE;  
REPAIR FAILURE;
```

### Useful Views

```
SELECT * FROM V$DATABASE_BLOCK_CORRUPTION;  
SELECT * FROM V$BACKUP_SET_DETAILS;  
SELECT * FROM V$FLASH_RECOVERY_AREA_USAGE;
```

---

## A.10 Notes

- Always perform a full backup immediately after an OPEN RESETLOGS.
- RMAN commands are case-insensitive but keywords should be in uppercase for readability.
- Use scripts and logs to ensure repeatable and auditable backup operations.



## Appendix B – Recovery Catalog Views and Metadata

### B.1 Overview

The RMAN **recovery catalog** stores metadata about all backups, copies, and database incarnations.

This information extends beyond the retention window of the control file and is essential for long-term audit and multi-database administration.

### B.2 Common Catalog Views

View	Description
RC_DATABASE	Lists all databases registered in the catalog.
RC_BACKUP_SET	Displays summary of backup sets.
RC_BACKUP_PIECE	Provides details for each backup piece.
RC_DATAFILE_COPY	Lists all datafile copies.
RC_RESTORE_POINT	Shows restore points available for each database.
RC_RMAN_STATUS	Records job status and messages for each RMAN session.
RC_UNUSABLE_BACKUP	Identifies expired or inaccessible backup files.

### B.3 Example Metadata Queries

#### List Registered Databases

```
SELECT DB_KEY, DB_NAME, DBID FROM RC_DATABASE;
```

#### List Backup Sets for a Database

```
SELECT BACKUP_KEY, INCREMENTAL_LEVEL, STATUS, COMPLETION_TIME
FROM RC_BACKUP_SET
```



```
WHERE DB_KEY = (SELECT DB_KEY FROM RC_DATABASE WHERE DB_NAME='PROD');
```

### **Report Unusable Backups**

```
SELECT HANDLE, DEVICE_TYPE FROM RC_UNUSABLE_BACKUP;
```

### **Identify Restore Points**

```
SELECT NAME, TIME, GUARANTEE_FLASHBACK_DATABASE FROM RC_RESTORE_POINT;
```

---

## **B.4 Catalog Maintenance**

### **Resynchronize**

```
RMAN> RESYNC CATALOG;
```

### **Upgrade Schema**

```
RMAN> UPGRADE CATALOG;
```

### **Unregister Database**

```
RMAN> UNREGISTER DATABASE;
```

Perform catalog backups regularly and secure the catalog schema with minimal privileges.



---

## Appendix C – Automation and Script Samples

---

### C.1 Purpose

Automating RMAN operations ensures consistent, error-free backups and recoveries. This appendix provides sample scripts for Linux and Windows environments, including daily backups, validations, and retention maintenance.

---

### C.2 Daily Backup Script (Linux)

```
#!/bin/bash

export ORACLE_SID=PROD
export ORACLE_HOME=/u01/app/oracle/product/19c/dbhome_1
export PATH=$ORACLE_HOME/bin:$PATH

rman TARGET / <<EOF
RUN {
  CROSSCHECK BACKUP;
  DELETE NOPROMPT EXPIRED BACKUP;
  BACKUP DATABASE PLUS ARCHIVELOG TAG 'DAILY_BACKUP';
  DELETE NOPROMPT OBSOLETE;
}
EXIT;
EOF

Schedule with cron:
0 2 * * * /u01/scripts/rman_daily.sh > /u01/scripts/rman_daily.log 2>&1
```

---



### C.3 Weekly Validation Script

```
rman TARGET / <<EOF

RUN {

  CROSSCHECK BACKUP;

  VALIDATE DATABASE;

  LIST BACKUP SUMMARY;

}

EXIT;

EOF
```

---

### C.4 Recovery Drill Script

```
rman TARGET / <<EOF

RUN {

  STARTUP NOMOUNT;

  RESTORE CONTROLFILE FROM AUTOBACKUP;

  ALTER DATABASE MOUNT;

  RESTORE DATABASE VALIDATE;

}

EXIT;

EOF
```

---

### C.5 Automated Log Rotation

```
find /logs/rman -type f -mtime +30 -delete
```

Run weekly to maintain a clean logging environment.

---



## C.6 Windows RMAN Batch Example

```
set ORACLE_SID=PROD
set ORACLE_HOME=C:\app\oracle\product\19c\dbhome_1
rman TARGET / CMDfile=backup_full.cmd LOG=backup_full.log

backup_full.cmd

RUN {
    BACKUP DATABASE PLUS ARCHIVELOG;
}
```

---

## C.7 Scheduled Job Verification

After scheduling RMAN jobs, verify with:

```
SELECT JOB_NAME, STATUS, LAST_START_DATE, LAST_RUN_DURATION
FROM DBA_SCHEDULER_JOBS;
```



## Appendix D – Troubleshooting Reference

### D.1 Purpose

This appendix consolidates common Oracle and RMAN errors encountered during backup and recovery, along with their corrective actions.

It serves as a quick-access diagnostic guide.

### D.2 Common RMAN Errors

Error	Description	Resolution
<b>RMAN-06004</b>	Invalid backup piece handle.	Verify path and re-catalog missing files.
<b>RMAN-06025</b>	No backup found for specified file.	Check retention policy; revalidate backups.
<b>RMAN-10035</b>	Error during job execution.	Check I/O subsystem and network paths.
<b>RMAN-08120</b>	Archived log not deleted; still needed.	Review recovery window and log dependencies.
<b>RMAN-20005</b>	Invalid backup or recovery object.	Crosscheck and re-catalog backups.

### D.3 Common ORA Errors

Error	Meaning	Resolution
<b>ORA-01113</b>	File needs media recovery.	Run RECOVER DATAFILE or RECOVER DATABASE.



Error	Meaning	Resolution
<b>ORA-01110</b>	Datafile mismatch with control file.	Restore consistent control file and redo logs.
<b>ORA-00205</b>	Error identifying control file.	Restore from autobackup.
<b>ORA-01547</b>	Resetlogs required after recovery.	Open with ALTER DATABASE OPEN RESETLOGS;;
<b>ORA-19809</b>	FRA full.	Delete obsolete backups or increase FRA size.

## D.4 Troubleshooting Sequence

1. Review RMAN log for first error occurrence.
2. Check alert log for corresponding ORA messages.
3. Validate FRA, disk space, and permissions.
4. Crosscheck and re-catalog missing backup sets.
5. Validate and re-test restore in a controlled environment.
6. Record corrective action for audit purposes.



---

## Appendix E – Lab Scenarios

---

### E.1 Introduction

This appendix contains practical recovery exercises designed to reinforce the procedures covered in the main chapters.

Each scenario can be executed in a non-production lab or virtualized test environment.

---

### E.2 Scenario 1 – Control File Loss

**Objective:** Restore database after control file deletion.

**Steps:**

1. Remove control files from the OS.
2. Start database in NOMOUNT mode.
3. Restore control file from autobackup.
4. Mount and recover database.
5. Open with RESETLOGS.

---

### E.3 Scenario 2 – Datafile Corruption

**Objective:** Repair a damaged datafile using RMAN.

**Steps:**

1. Identify corrupted file via alert log or V\$DATABASE\_BLOCK\_CORRUPTION.
2. Execute BLOCKRECOVER or RESTORE DATAFILE.
3. Recover database and verify integrity.

---

### E.4 Scenario 3 – Redo Log Loss

**Objective:** Recover database after redo log failure.

**Steps:**



1. Identify damaged redo group.
2. Drop and recreate logs using:

ALTER DATABASE ADD LOGFILE GROUP 4 ('redo04.log') SIZE 100M;

3. Recover missing transactions from archived logs.

---

## E.5 Scenario 4 – Tablespace Point-in-Time Recovery

**Objective:** Recover a specific tablespace to an earlier time.

**Steps:**

1. Use RMAN to create auxiliary instance.
2. Perform RECOVER TABLESPACE hr\_data UNTIL TIME '...'
3. Validate and bring tablespace online.

---

## E.6 Scenario 5 – Flashback and Restore Point Test

**Objective:** Validate Flashback Database functionality.

**Steps:**

1. Create guaranteed restore point.
2. Perform test DML changes.
3. Flash back to restore point.
4. Validate results and drop restore point.

---

## E.7 Scenario 6 – RMAN Validation Drill

**Objective:** Confirm backup integrity.

**Steps:**

CROSSCHECK BACKUP;

VALIDATE DATABASE;

LIST BACKUP SUMMARY;

Document findings and note any missing or corrupted backups.



---

## E.8 Documentation

Each lab scenario should be followed by a written summary:

- Date and time of test
- RMAN logs and results
- Observations and elapsed times
- DBA signature and approval



## Appendix F – Version Compatibility and Feature Map

### F.1 Purpose

This appendix provides a concise comparison of Oracle backup and recovery features across major database releases from 11g through 21c.

It helps DBAs identify version-specific capabilities and deprecations when supporting mixed environments.

### F.2 Feature Evolution Summary

Feature	11g	12c	18c	19c	21c
RMAN Active Duplication	✓	✓	✓	✓	✓
Block Change Tracking	✓	✓	✓	✓	✓
Incremental Merge	✓	✓	✓	✓	✓
Table-Level Recovery	–	✓	✓	✓	✓
Multitenant (CDB/PDB) Recovery	–	✓	✓	✓	✓
Unified Auditing for RMAN	–	✓	✓	✓	✓
Flashback Transaction Backout	✓	✓	✓	✓	✓
Automatic Block Repair	✓	✓	✓	✓	✓
Recovery Advisor Integration	✓	✓	✓	✓	✓
Persistent Memory / AutoOptimize	–	–	✓	✓	✓



## F.3 Deprecated or Replaced Features

Deprecated Feature	Replacement
Legacy dbms_backup_restore scripts	RMAN automated recovery
Standalone Flash Recovery Area (pre-11g)	Fast Recovery Area (FRA)
BACKUP CONTROLFILE TO TRACE (manual only)	RMAN control file autobackup
Logical Standby-only Duplication	RMAN Active Duplication
Recovery Manager without Catalog	Catalog or control file repository

## F.4 Operating System Considerations

OS	Special Notes
Windows	Ensure consistent drive mapping; use rman.cmd scripts for automation.
Linux/Unix	Use cron scheduling and shell logging for performance reporting.
Mixed Environments	Verify endianness before cross-platform duplication.

## F.5 Version-Specific Enhancements

- **Oracle 12c:** Introduction of table-level recovery and multitenant backup management.
- **Oracle 18c:** Enhanced RMAN transportable tablespace and unified audit integration.
- **Oracle 19c:** Performance improvements in incremental merge and Flashback Database.
- **Oracle 21c:** Support for hybrid database duplication and automatic storage tiering.

## F.6 Summary

Understanding version differences ensures that recovery procedures are aligned with each environment's capabilities.



Before implementing advanced RMAN or Flashback features, confirm database version and licensing options.

Consistent documentation of supported features across systems minimizes confusion and ensures reliability during multi-version recovery scenarios.



## Appendix G – Linux RMAN Backup Script and Shell Example

Scheduling an RMAN backup on Linux typically involves creating a shell script containing the RMAN commands and then scheduling this script to run using cron.

### 1. Create the RMAN Backup Script:

Create a shell script, for example, rman\_backup.sh, that includes the necessary environment variables and RMAN commands.

```
#!/bin/bash

# Set Oracle environment variables
ORACLE_SID=ORCL  # Replace with your Oracle SID
export ORACLE_SID
ORACLE_HOME=/u01/app/oracle/product/19.0.0/dbhome_1 # Replace with your
Oracle Home
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:$PATH
export PATH
LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export LD_LIBRARY_PATH

# Define backup directory and log file
BACKUP_DIR=/backup/rman
LOG_FILE=$BACKUP_DIR/rman_backup_$(date +%Y%m%d%H%M%S).log

# Ensure backup directory exists
mkdir -p $BACKUP_DIR

# RMAN commands
rman target / log=$LOG_FILE <<EOF
run {
  allocate channel d1 device type disk;
  backup database plus archivelog format '$BACKUP_DIR/full_backup_%U';
  release channel d1;
  delete noprompt obsolete;
  delete noprompt archivelog all backed up 1 times to disk;
}
EOF

# Add any post-backup actions here, e.g., email notifications
```



## 2. Make the Script Executable:

Grant execute permissions to the script:

```
chmod +x /path/to/your/rman_backup.sh
```

## 3. Schedule with Cron:

Use crontab -e to edit the cron table for the user that will run the backup (e.g., the Oracle user).

```
crontab -e
```

Add an entry to schedule the script. For example, to run the backup daily at 2:00 AM:

```
0 2 * * * /path/to/your/rman_backup.sh
```

### Explanation of Cron Entry:

- 0: Minute (0-59)
- 2: Hour (0-23)
- \*: Day of the month (1-31)
- \*: Month (1-12)
- \*: Day of the week (0-7, where 0 and 7 are Sunday)

Replace /path/to/your/rman\_backup.sh with the actual path to your script. Adjust the time and frequency as needed.

## Important Considerations:

- **Permissions:**

Ensure the user running the cron job has the necessary permissions to execute the RMAN commands and write to the backup directory.

- **Environment Variables:**

Verify that all necessary Oracle environment variables are correctly set within the script.

- **Error Handling:**

Implement robust error handling and logging within your script to monitor backup success or failure.

- **Backup Strategy:**

Tailor the RMAN commands in your script to your specific backup strategy (e.g., full, incremental, archivelog-only).



- **Recovery Catalog:**

If using a recovery catalog, connect to it in your RMAN script.

- **Retention Policy:**

Configure your RMAN retention policy to manage backup sets and obsolete backups.



## Appendix H – RMAN Backup Script Example for Veeam on a Linux Operating System

Shown below is a sample RMAN backup script using Veeam on a Linux Operating System. You will need to use your Veeam values for the VEEAM\_AGENT\_ID and the VEEAM\_REPOSITORY\_ID. Default values are pulled from the default RMAN settings.

### Default RMAN Settings

```
DEFAULT DEVICE TYPE TO 'SBT_TAPE';
DEVICE TYPE 'SBT_TAPE' PARALLELISM 4 BACKUP TYPE TO BACKUPSET;
CONTROLFILE AUTOBACKUP ON;
CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO
'%F_RMAN_AUTOBACKUP.vab';
CHANNEL DEVICE TYPE 'SBT_TAPE' PARMs
'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMANPlugin.so' FORMAT
'c49f22f2-892f-4d97-8259-cb7a161dccc5/RMAN_%I_%d_%T_%U.vab';
ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
DATAFILE BACKUP COPIES FOR DEVICE TYPE 'SBT_TAPE' TO 1;
```

### RMAN Veeam Script Example

```
run {
  ALLOCATE CHANNEL CH1 DEVICE TYPE SBT_TAPE
  PARMs 'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMANPlugin.so,
  ENV=(VEEAM_AGENT_ID=44454c4c-3100-1035-8054-b5c04f353634,
  VEEAM_REPOSITORY_ID=c49f22f2-892f-4d97-8259-cb7a161dccc5))';

  ALLOCATE CHANNEL CH2 DEVICE TYPE SBT_TAPE
  PARMs 'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMANPlugin.so,
  ENV=(VEEAM_AGENT_ID=44454c4c-3100-1035-8054-b5c04f353634,
  VEEAM_REPOSITORY_ID=c49f22f2-892f-4d97-8259-cb7a161dccc5))';

  ALLOCATE CHANNEL CH3 DEVICE TYPE SBT_TAPE
  PARMs 'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMANPlugin.so,
  ENV=(VEEAM_AGENT_ID=44454c4c-3100-1035-8054-b5c04f353634,
  VEEAM_REPOSITORY_ID=c49f22f2-892f-4d97-8259-cb7a161dccc5))';

  ALLOCATE CHANNEL CH4 DEVICE TYPE SBT_TAPE
  PARMs 'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMANPlugin.so,
  ENV=(VEEAM_AGENT_ID=44454c4c-3100-1035-8054-b5c04f353634,
  VEEAM_REPOSITORY_ID=c49f22f2-892f-4d97-8259-cb7a161dccc5))';

  BACKUP AS COMPRESSED BACKUPSET
    DATABASE
    INCLUDE CURRENT CONTROLFILE;

  BACKUP AS COMPRESSED BACKUPSET
    ARCHIVELOG ALL DELETE INPUT;

  BACKUP SPFILE;
```



```
RELEASE CHANNEL CH1;  
RELEASE CHANNEL CH2;  
RELEASE CHANNEL CH3;  
RELEASE CHANNEL CH4;  
}
```



---

## PART VII – QUICK REFERENCE

---

### Quick Reference Overview

This section provides **immediate-access command blocks** for the most common recovery scenarios.

They are designed to be executed directly in a terminal session during an emergency, with minimal modification.

The goal is rapid, error-free recovery when time and precision are critical.

Each block assumes:

- On-premise Oracle Database 12c–21c
- Disk-based backups (no ASM or cloud storage)
- Database running in ARCHIVELOG mode
- Proper permissions and environment variables (ORACLE\_SID, ORACLE\_HOME) are set

---

### RMAN Emergency Command Blocks

These are short, self-contained procedures that can restore a database to operational status in a crisis.

---

#### A. Full Database Restore and Recovery

**Use when:** The entire database has been lost or datafiles are inaccessible.

```
# Connect to RMAN
```

```
rman TARGET /
```

```
# Startup and mount
```



STARTUP MOUNT;

# Restore all datafiles from the latest backup

RESTORE DATABASE;

# Recover database using all available archived and online redo logs

RECOVER DATABASE;

# Open database normally

ALTER DATABASE OPEN;

### Notes

- Use RESTORE DATABASE VALIDATE; first if unsure of backup integrity.
- If recovery fails due to missing redo logs, use incomplete recovery (next section).

## B. Incomplete Recovery to a Specific Time or SCN

**Use when:** You need to restore to a prior point (e.g., before user error).

# Connect and mount

rman TARGET /

STARTUP MOUNT;

# Set recovery point

RUN {

SET UNTIL TIME "TO\_DATE('2025-11-11 10:00:00','YYYY-MM-DD HH24:MI:SS')";

RESTORE DATABASE;

RECOVER DATABASE;

}



```
# Open with RESETLOGS (required after incomplete recovery)
```

```
ALTER DATABASE OPEN RESETLOGS;
```

**Alternate:**

```
SET UNTIL SCN 123456789;
```

---

## C. Restore Control File from Autobackup

**Use when:** Control file is missing or corrupted.

```
rman TARGET /
```

```
# Start instance without mounting
```

```
STARTUP NOMOUNT;
```

```
# Restore control file
```

```
RESTORE CONTROLFILE FROM AUTOBACKUP;
```

```
# Mount and recover
```

```
ALTER DATABASE MOUNT;
```

```
RESTORE DATABASE;
```

```
RECOVER DATABASE;
```

```
ALTER DATABASE OPEN RESETLOGS;
```

**Tip:** Use SHOW ALL; before failure to confirm autobackup is enabled.

---

## D. Restore a Single Datafile

**Use when:** One datafile is damaged or missing.

```
rman TARGET /
```



```
# Identify and restore only the missing file

RESTORE DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf';
RECOVER DATAFILE '/u01/app/oracle/oradata/PROD/users01.dbf';

# Bring file online

SQL 'ALTER DATABASE DATAFILE "/u01/app/oracle/oradata/PROD/users01.dbf" ONLINE';

Alternate (by number):

RESTORE DATAFILE 5;
RECOVER DATAFILE 5;
```

---

## E. Restore and Recover Specific Tablespace

**Use when:** A single functional area (e.g., HR\_DATA) needs recovery.

```
rman TARGET /
SQL 'ALTER TABLESPACE hr_data OFFLINE IMMEDIATE';
RESTORE TABLESPACE hr_data;
RECOVER TABLESPACE hr_data;
SQL 'ALTER TABLESPACE hr_data ONLINE';
```

---

## F. Restore SPFILE from Backup

**Use when:** The server parameter file is missing or corrupted.

```
rman TARGET /
STARTUP NOMOUNT PFILE='/backup/init_temp.ora';
RESTORE SPFILE FROM AUTOBACKUP;
SHUTDOWN IMMEDIATE;
```



```
STARTUP NOMOUNT;
```

---

## G. Block Media Recovery

**Use when:** Corruption affects specific blocks only.

```
rman TARGET /
```

```
BLOCKRECOVER DATAFILE 7 BLOCK 1024 TO 1056;
```

**Optional:** Identify corrupt blocks first:

```
SELECT * FROM V$DATABASE_BLOCK_CORRUPTION;
```

---

## H. Recreate Control File Manually (No Backup Available)

**Use when:** All control file backups are missing, but datafiles and redo logs exist.

```
STARTUP NOMOUNT;
```

```
CREATE CONTROLFILE REUSE DATABASE "PROD" RESETLOGS ARCHIVELOG
```

```
LOGFILE
```

```
GROUP 1 '/u01/app/oracle/oradata/redo01.log' SIZE 100M,
```

```
GROUP 2 '/u01/app/oracle/oradata/redo02.log' SIZE 100M
```

```
DATAFILE
```

```
'/u01/app/oracle/oradata/system01.dbf',
```

```
'/u01/app/oracle/oradata/sysaux01.dbf',
```

```
'/u01/app/oracle/oradata/users01.dbf'
```

```
CHARACTER SET AL32UTF8;
```

```
RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL;
```

```
ALTER DATABASE OPEN RESETLOGS;
```



---

## RESTORE DATABASE and RECOVER DATABASE Templates

These templates serve as “drop-in” recovery scripts for different scenarios.  
They can be pre-customized and stored in /scripts/recovery for immediate use.

---

### Template 1 – Standard Full Restore and Recovery

```
RUN {  
  STARTUP MOUNT;  
  RESTORE DATABASE;  
  RECOVER DATABASE;  
  ALTER DATABASE OPEN;  
}
```

---

### Template 2 – Full Restore with Validation and Log

```
SPOOL LOG TO '/logs/full_restore_&DATE..log';  
RUN {  
  STARTUP MOUNT;  
  RESTORE DATABASE VALIDATE;  
  RESTORE DATABASE;  
  RECOVER DATABASE;  
  ALTER DATABASE OPEN;  
}  
SPOOL LOG OFF;
```

---



## Template 3 – Point-in-Time (Incomplete) Recovery

```
RUN {  
  STARTUP MOUNT;  
  SET UNTIL TIME "TO_DATE('2025-11-10 22:45:00','YYYY-MM-DD HH24:MI:SS')";  
  RESTORE DATABASE;  
  RECOVER DATABASE;  
  ALTER DATABASE OPEN RESETLOGS;  
}
```

---

## Template 4 – Tablespace Recovery

```
RUN {  
  SQL 'ALTER TABLESPACE hr_data OFFLINE IMMEDIATE';  
  RESTORE TABLESPACE hr_data;  
  RECOVER TABLESPACE hr_data;  
  SQL 'ALTER TABLESPACE hr_data ONLINE';  
}
```

---

## Template 5 – Datafile-Level Restore

```
RUN {  
  RESTORE DATAFILE 5;  
  RECOVER DATAFILE 5;  
  SQL 'ALTER DATABASE DATAFILE 5 ONLINE';  
}
```

---

## Template 6 – Control File Restore and Database Recovery

```
RUN {
```



```
STARTUP NOMOUNT;  
RESTORE CONTROLFILE FROM AUTOBACKUP;  
ALTER DATABASE MOUNT;  
RESTORE DATABASE;  
RECOVER DATABASE;  
ALTER DATABASE OPEN RESETLOGS;  
}
```

---

### Template 7 – Validation Only (No Restore)

```
RUN {  
  CROSSCHECK BACKUP;  
  VALIDATE DATABASE;  
  REPORT OBSOLETE;  
}
```



## Flashback and Table Recovery Quick Commands

These commands offer instant recovery from logical or user errors without using backups. They should be part of every DBA's daily quick-recovery toolkit.

---

### A. Flashback Table

**Use when:** Rows were accidentally deleted or updated.

```
FLASHBACK TABLE hr.employees
```

```
TO TIMESTAMP TO_TIMESTAMP('2025-11-10 11:00:00', 'YYYY-MM-DD HH24:MI:SS');
```

**Check SCN before use:**

```
SELECT SCN_TO_TIMESTAMP(CURRENT_SCN - 1000) FROM DUAL;
```

---

### B. Flashback Table to Before Drop

**Use when:** A table was dropped unintentionally.

```
FLASHBACK TABLE hr.employees TO BEFORE DROP;
```

**View available objects in recycle bin:**

```
SHOW RECYCLEBIN;
```

---

### C. Flashback Query

**Use when:** You need to view historical data without changing current rows.

```
SELECT * FROM hr.employees AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '15' MINUTE)  
WHERE employee_id = 100;
```

---

### D. Flashback Transaction Backout

**Use when:** A specific transaction needs to be undone.



```

BEGIN

DBMS_FLASHBACK.TRANSACTION_BACKOUT (
    numtxns => 1,
    xids => xid_array('08001F001C020000'),
    options => dbms_flashback.cascade);

END;
/

```

**View recent transactions:**

```
SELECT xid, operation, undo_sql FROM flashback_transaction_query;
```

---

## E. Flashback Database

**Use when:** A major logical error has affected the entire database.

```

SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
FLASHBACK DATABASE TO RESTORE POINT pre_patch;
ALTER DATABASE OPEN RESETLOGS;

```

**Or flashback to SCN:**

```
FLASHBACK DATABASE TO SCN 123456789;
```

**Create restore point before risky operation:**

```
CREATE RESTORE POINT pre_patch GUARANTEE FLASHBACK DATABASE;
```

---

## F. Drop Restore Point After Success

```
DROP RESTORE POINT pre_patch;
```



## Emergency Recovery Decision Matrix

Situation	Recommended Action
Control file lost	Use “Restore Control File from Autobackup” command block.
One or more datafiles missing	Run “Restore Datafile” block.
Entire database lost	Run “Full Database Restore and Recovery.”
User dropped or truncated table	Use Flashback Table or Table Recovery.
Data corruption in a few blocks	Use Block Media Recovery.
Parameter file (SPFILE) missing	Use “Restore SPFILE from Backup.”
System crash or power loss	Simply restart; Oracle performs instance recovery automatically.



## Emergency Recovery Notes

- Always **mount** the database before attempting recovery.
- Use **VALIDATE DATABASE** if unsure about backup integrity.
- Perform a **new full backup** immediately after any recovery using **RESETLOGS**.
- Keep a printed copy of this quick reference near production consoles.
- In multi-DB environments, verify **ORACLE\_SID** and **DB\_NAME** before executing any command.
- Document every recovery event and store the logs in a secured audit directory.