

Oracle® Database

Backup and Recovery Advanced User's Guide

10g Release 2 (10.2)

B14191-03

August 2008

A guide to advanced backup and recovery of Oracle databases and advanced uses of Recovery Manager (RMAN), including advanced RMAN database backup and recovery concepts and scenarios, using RMAN for data migration, transport and duplication, and user-managed backup and recovery without RMAN.

Oracle Database Backup and Recovery Advanced User's Guide, 10g Release 2 (10.2)

B14191-03

Copyright © 2003, 2008, Oracle. All rights reserved.

Primary Author: Antonio Romero

Contributing Author: Lance Ashdown

Contributors: Kat Weill, Tammy Bednar, Anand Beldalker, Timothy Chien, Raymond Guzman, Alex Hwang, Ashok Joshi, J. William Lee, Valarie Moore, Muthu Olagappan, Samitha Samaranyake, Francisco Sanchez, Steven Wertheimer, Wanli Yang

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	xxiii
Audience	xxiii
Documentation Accessibility	xxiii
Related Documentation	xxiv
Conventions	xxiv
What's New in Backup and Recovery?	xxv
Oracle Database 10g Release 2 New Features in Backup and Recovery	xxv
Part I Recovery Manager Advanced Architecture and Concepts	
1 Recovery Manager Architecture	
About the RMAN Environment	1-1
RMAN Session Architecture	1-2
RMAN Command Line Client	1-2
How RMAN Compiles and Executes Commands	1-2
Compilation Phase	1-2
Execution Phase	1-3
Issuing RMAN Commands	1-3
Entering Commands at the RMAN Prompt	1-3
Using RMAN with Command Files	1-3
Stored Scripts	1-3
Commands Valid Only in RUN Blocks	1-4
Commands Not Valid in RUN Blocks	1-4
Directing RMAN Output to a Log File	1-4
RMAN Repository	1-4
Storage of the RMAN Repository in the Control File	1-5
Types of Records in the Control File	1-5
Recovery Without a Recovery catalog	1-6
Storage of the RMAN Repository in the Recovery Catalog	1-6
Registration of Databases in the Recovery Catalog	1-6
Contents of the Recovery Catalog	1-6
Resynchronization of the Recovery Catalog	1-7
Backups of the Recovery Catalog	1-7
Compatibility of the Recovery Catalog	1-7

Media Management	1-7
Performing Backup and Restore with a Media Manager	1-8
Backup Solutions Program	1-8

2 RMAN Backups Concepts

About RMAN Channels	2-1
Automatic and Manual Channel Allocation	2-2
Automatic Channel Device Configuration and Parallelism	2-3
Automatic Channel Default Device Types	2-4
Automatic Channel Naming Conventions	2-5
Automatic Channel Generic Configurations.....	2-5
Automatic Channel-Specific Configurations	2-6
Clearing Automatic Channel Settings.....	2-6
Determining Channel Parallelism to Match Hardware Devices.....	2-7
Channel Control Options for Manual and Automatic Channels	2-8
Channel Failover	2-9
About RMAN Backups	2-9
About Image Copies	2-9
Using RMAN-Created Image Copies.....	2-10
User-Managed Image Copies.....	2-10
About Proxy Copies.....	2-10
Storage of Backups on Disk and Tape.....	2-11
Backups of Archived Logs.....	2-11
Deletion of Archived Logs After Backups.....	2-11
Backup Failover for Archived Redo Logs	2-11
Multiplexed Backup Sets.....	2-12
Algorithm for Multiplexed Backups	2-12
Multiplexing by the Media Manager	2-13
Manual Parallelization of Backups	2-13
Multiple Copies of RMAN Backups	2-14
Duplexed Backup Sets.....	2-15
Backups of Backup Sets.....	2-16
Uses for Backups of Backup Sets	2-17
Backup Optimization When Backing Up Backup Sets	2-17
Backup Failover When Backing Up Backup Sets	2-17
Backups of Image Copies	2-18
RMAN Backup Options: Naming, Sizing, and Speed	2-18
Generating Filenames for RMAN Backup Pieces.....	2-18
Generating Filenames for RMAN Image Copies.....	2-19
Using FORMAT with RMAN BACKUP AS COPY	2-19
Using DB_FILE_NAME_CONVERT With BACKUP AS COPY.....	2-19
Using Tags to Identify RMAN Backups	2-19
Default RMAN Backup Tag Format.....	2-20
How Tags Are Applied	2-20
Uniqueness of Backup Tags.....	2-20
Size of Backup Pieces.....	2-21
Number and Size of Backup Sets	2-22

Factors Affecting the Number and Size of Backup Sets	2-22
Overview of the MAXSETSIZE Parameter	2-23
I/O Read Rate of Backups	2-23
RMAN Backup Types	2-23
Incremental Backups.....	2-24
Incremental Backup Algorithm	2-25
Multilevel Incremental Backups	2-25
Differential Incremental Backups	2-25
Cumulative Incremental Backups	2-26
Planning an Incremental Backup Strategy	2-27
Control File and Server Parameter File Autobackups.....	2-28
How RMAN Performs Control File Autobackups	2-28
When RMAN Performs Control File Autobackups	2-29
Control File Autobackups After Backup Activities.....	2-29
Control File Autobackups After Database Structural Changes	2-29
Backup Retention Policies	2-30
Recovery Window	2-31
Backup Redundancy	2-33
Batch Deletes of Obsolete Backups.....	2-33
Exempting Backups from the Retention Policy	2-34
Relationship Between Retention Policy and Flash Recovery Area Rules	2-35
Backup Optimization	2-35
Rules for Identifying Identical Files for Backup Optimization	2-35
Backup Optimization Algorithm	2-35
Requirements for Backup Optimization	2-37
Overriding and Disabling Backup Optimization	2-37
Effect of Retention Policies on Backup Optimization for SBT Backups	2-37
Backup Optimization for SBT Backups with Recovery Window Retention Policy	2-38
Backup Optimization for SBT Backups With Redundancy Retention Policy	2-38
Restartable Backups.....	2-39
Managing Backup Windows and Performance: BACKUP.. DURATION	2-40
Controlling RMAN Behavior when Backup Window Ends with PARTIAL	2-40
Managing Backup Performance with MINIMIZE TIME and MINIMIZE LOAD	2-41
How RMAN Responds to Backup Errors	2-41
Handling I/O Errors in RMAN Backup: NOT BACKED UP SINCE.....	2-41
Handling Corrupt Datafile Blocks in RMAN Backup: MAXCORRUPT	2-42
Tests and Integrity Checks for Backups.....	2-42
Detecting Physical Block Corruption With RMAN BACKUP.. VALIDATE.....	2-43
Detection of Logical Block Corruption	2-43
Detection of Fractured Blocks During Open Backups.....	2-44
Backup Validation with RMAN.....	2-44

3 RMAN Recovery Concepts

Restoring Files with RMAN.....	3-1
Mechanics of Datafile Restore Operations	3-1
How RMAN Selects Backups for Use in Restore Operations.....	3-2
Restore Failover	3-3

Restore Optimization.....	3-3
Datafile Media Recovery with RMAN	3-4
RMAN Media Recovery: Basic Steps	3-4
Mechanics of Recovery: Incremental Backups and Redo Logs	3-5
How RMAN Searches for Archived Redo Logs During Recovery.....	3-5
RMAN Behavior When the Repository Is Not Synchronized	3-6
Incomplete or Point-In-Time Recovery.....	3-7
Tablespace Point-in-Time Recovery	3-7
Block Media Recovery with RMAN	3-7
When to Use Block Media Recovery	3-8
Block Media Recovery When Redo Is Missing	3-9
Database Duplication with RMAN	3-9
Physical Standby Database Creation with RMAN	3-11

4 Connecting to Databases with RMAN

Starting RMAN Without Connecting to a Database	4-1
Connecting to a Target Database and a Recovery Catalog	4-1
Connecting to the Target Database and Recovery Catalog from the Command Line.....	4-2
Connecting to the Target Database and Recovery Catalog from the RMAN Prompt.....	4-2
Connecting to an Auxiliary Database	4-2
Connecting to an Auxiliary Database from the Command Line	4-3
Connecting to an Auxiliary Database from the RMAN Prompt.....	4-3
Diagnosing Connection Problems	4-3
Diagnosing Target and Auxiliary Database Connection Problems.....	4-4
Diagnosing Recovery Catalog Connection Problems	4-4
Hiding Passwords When Connecting to Databases	4-4
Sending RMAN Output Simultaneously to the Terminal and a Log File	4-5
Using the RMAN Pipe Interface	4-5
Executing Multiple RMAN Commands In Succession Through a Pipe: Example.....	4-6
Executing RMAN Commands In a Single Job Through a Pipe: Example	4-6

Part II Advanced RMAN Backup and Recovery Topics

5 Configuring the RMAN Environment: Advanced Topics

Configuring the Flash Recovery Area: Advanced Topics	5-1
Configuring Online Redo Log Creation in the Flash Recovery Area.....	5-1
Configuring Control File Creation in the Flash Recovery Area	5-2
Archived Redo Log Creation in the Flash Recovery Area	5-2
Rules for Initialization Parameters Affecting Redo Log File Destinations.....	5-2
Filenames for Archived Redo Log Files in the Flash Recovery Area	5-3
RMAN File Creation in the Flash Recovery Area	5-3
Configuring RMAN to Make Backups to a Media Manager	5-4
Prerequisites for Using a Media Manager with RMAN.....	5-4
Locating the Media Management Library: The SBT_LIBRARY Parameter.....	5-4
Testing Whether the Media Manager Library Is Integrated Correctly	5-5
Configuring Media Management Software for RMAN Backups	5-5

Configuring Backup Piece Names and Sizes for a Media Manager	5-6
Testing ALLOCATE CHANNEL on the Media Manager	5-6
Testing a Backup to the Media Manager	5-7
Configuring SBT Channels for Use with a Media Manager	5-8
Configuring Channels	5-9
Configuring Channel Parallelism	5-9
Configuring Channel Settings for a Device Type.....	5-9
Showing the Configured Channel Settings	5-10
Showing the Currently Configured Channel Settings.....	5-10
Showing the Configured Device Types	5-11
Showing the Default Device Type	5-11
Manually Overriding Configured Channels.....	5-11
Configuring a Specific Channel for a Device Type	5-12
Configuring Specific Channels: Examples	5-12
Mixing Generic and Specific Channels	5-13
Relationship Between CONFIGURE CHANNEL and Parallelism Setting	5-14
Clearing Channel and Device Settings.....	5-14
Configuring the Maximum Size of Backup Sets and Pieces	5-15
Showing the Default Maximum Size of Backup Sets: SHOW MAXSETSIZE	5-15
Configuring Backup Optimization	5-16
Displaying Backup Optimization Setting: SHOW BACKUP OPTIMIZATION	5-16
Configuring Backup Duplexing: CONFIGURE... BACKUP COPIES	5-17
Showing the Configured Degree of Duplexing: SHOW... BACKUP COPIES	5-18
Configuring Tablespaces for Exclusion from Whole Database Backups	5-18
Showing the Tablespaces Excluded from Backups	5-19
Configuring Auxiliary Instance Datafile Names: CONFIGURE AUXNAME	5-19
Showing the Default Filenames Configured for Auxiliary Channels	5-20
Setting the Snapshot Control File Location	5-20
Default Location of the Snapshot Control File.....	5-20
Viewing the Configured Location of the Snapshot Control File.....	5-20
Setting the Location of the Snapshot Control File	5-21
Showing the Current Snapshot Control File Name	5-21
Setting Up RMAN for Use with a Shared Server	5-22

6 Making Backups with RMAN: Advanced Topics

Configuring and Allocating Channels for Use in Backups	6-1
Duplexing Backup Sets	6-2
Duplexing Backup Sets with CONFIGURE BACKUP COPIES	6-2
Duplexing Backupsets with BACKUP... COPIES.....	6-3
Making Split Mirror Backups with RMAN	6-3
Backing Up Backup Sets with RMAN	6-5
BACKUP BACKUPSET and Backup Retention Policy	6-5
Backing Up Existing Image Copy Backups with RMAN	6-6
RMAN Encrypted Backups	6-7
RMAN Backup Encryption Modes.....	6-7
Transparent Encryption of Backups.....	6-7
Password Encryption of Backups.....	6-8

Dual Mode Encryption of Backups	6-8
Using CONFIGURE and SET to Control RMAN Backup Encryption	6-9
Creating Encrypted Backups	6-9
Restoring Data from Encrypted Backups	6-10
Encryption of Archived Log Backups	6-10
Performance Impact of Encrypting RMAN Backups.....	6-10
Restarting and Optimizing RMAN Backups	6-10
Backing Up Files Using Backup Optimization	6-11
Restarting a Backup After It Partially Completes	6-11
Backups to CD, DVD and Other Disk Devices with Large Block Sizes	6-11
Validating Backups with RMAN	6-11
RMAN Backup Examples	6-12
Skipping Tablespaces when Backing Up a Database: Example	6-13
Restarting a Backup: Example.....	6-13
Spreading a Backup Across Multiple Disk Drives: Example	6-14
Specifying the Size of Backup Sets: Example	6-14
Limiting the Size of Backup Pieces: Example	6-15
Backing Up Archived Redo Logs in a Failover Scenario: Example	6-15
Backing Up Archived Logs Needed to Recover an Online Backup: Example	6-16
Backing Up and Deleting Multiple Copies of an Archived Redo Log: Example.....	6-17
Determining How Channels Distribute a Backup Workload: Example	6-17
Backing Up in NOARCHIVELOG Mode: Example	6-18
Keeping a Long-Term Backup: Example	6-19
Using Backup Optimization: Examples	6-19
Optimizing a Database Backup: Example	6-19
Optimizing a Daily Archived Log Backup to a Single Tape: Example	6-20
Optimizing a Daily Archived Log Backup to Multiple Tapes: Example	6-20
Creating a Weekly Secondary Backup of Archived Logs: Example.....	6-21
Handling Corruption During Backups: Example	6-21

7 Advanced RMAN Recovery Techniques

Restore and Recovery of NOARCHIVELOG Databases.....	7-1
Recovering a Database in NOARCHIVELOG Mode Using Incremental Backups	7-1
Restore and Recovery of the Database on a New Host.....	7-2
Preparing for Restore of a Database to a New Host	7-3
Testing the Restore of a Database to a New Host: Scenario	7-3
Performing Recovery with a Backup Control File	7-7
Performing Recovery with a Backup Control File and No Recovery Catalog: Scenario.....	7-8
Performing Disaster Recovery	7-10
Performing Block Media Recovery with RMAN	7-12
Block Media Recovery Using All Available Backups	7-12
Block Media Recovery Using Specific Backups	7-13
Block Media Recovery of Blocks Listed in V\$DATABASE_BLOCK_CORRUPTION	7-14
RMAN Restore and Recovery Examples	7-14
Restoring Datafile Copies to a New Host: Example	7-14
Restoring Control File When Databases in the Catalog Have the Same Name: Example ...	7-15
Restoring a Backup Control File By Using the DBID	7-15

8 RMAN Tablespace Point-in-Time Recovery (TSPITR)

Understanding RMAN TSPITR	8-1
RMAN TSPITR Concepts.....	8-1
How TSPITR Works With an RMAN-Managed Auxiliary Instance.....	8-3
Deciding When to Use TSPITR	8-3
Limitations of TSPITR	8-4
Planning and Preparing for TSPITR	8-5
Choosing the Right Target Time for TSPITR	8-5
Determining the Recovery Set: Analyzing Data Relationships.....	8-6
Identifying and Resolving Dependencies on the Primary Database	8-6
Identifying and Preserving Objects That Will Be Lost After TSPITR.....	8-7
Performing Basic RMAN TSPITR	8-8
Fully Automated RMAN TSPITR.....	8-8
Using an Auxiliary Destination for Automated RMAN TSPITR.....	8-9
Performing Fully Automated RMAN TSPITR	8-9
Tasks to Perform After Successful TSPITR	8-9
Handling Errors in Automated TSPITR	8-10
Performing Customized RMAN TSPITR with an RMAN-Managed Auxiliary Instance	8-10
Renaming TSPITR Recovery Set Datafiles with SET NEWNAME	8-10
Renaming TSPITR Auxiliary Set Datafiles	8-11
Renaming TSPITR Auxiliary Set Datafiles with SET NEWNAME	8-12
Using DB_FILE_NAME_CONVERT to Name Auxiliary Set Datafiles	8-12
Order of Precedence Among File Renaming Methods in TSPITR.....	8-13
Specifying Auxiliary Instance Control File Location in TSPITR.....	8-14
Specifying Auxiliary Instance Online Log Location in TSPITR.....	8-14
Using Image Copies for Faster RMAN TSPITR Performance	8-15
Using CONFIGURE AUXNAME With Recovery Set Image Copies in TSPITR	8-15
SET NEWNAME and CONFIGURE AUXNAME With Auxiliary Set Image Copies ...	8-16
Performing TSPITR With CONFIGURE AUXNAME and Image Copies: Scenario	8-16
Customizing Initialization Parameters for the Automatic Auxiliary Instance in TSPITR ...	8-17
Performing RMAN TSPITR Using Your Own Auxiliary Instance	8-18
Preparing Your Own Auxiliary Instance for RMAN TSPITR	8-18
Step 1: Create an Oracle Password File for the Auxiliary Instance	8-18
Step 2: Create an Initialization Parameter File for the Auxiliary Instance	8-19
Step 3: Check Oracle Net Connectivity to the Auxiliary Instance	8-20
Preparing RMAN Commands for TSPITR with Your Own Auxiliary Instance.....	8-20
Planning Channels for TSPITR with Your Own Auxiliary Instance.....	8-21
Planning Datafile Names with Your Own Auxiliary Instance: SET NEWNAME	8-21
Executing TSPITR with Your Own Auxiliary Instance	8-21
Step 1: Start the Auxiliary Instance in NOMOUNT Mode	8-21
Step 2: Connect the RMAN Client to Target and Auxiliary Instances.....	8-21
Step 3: Execute the RECOVER TABLESPACE Command.....	8-21
Executing TSPITR With Your Own Auxiliary Instance: Scenario.....	8-22
Troubleshooting RMAN TSPITR	8-23
Troubleshooting TSPITR: Filename Conflicts.....	8-24

Troubleshooting TSPITR: Insufficient Sort Space during Export	8-24
Troubleshooting TSPITR: RMAN Cannot Identify Tablespaces with Undo Segments.....	8-24
Troubleshooting: Restarting Manual Auxiliary Instance After TSPITR Failure.....	8-24

9 RMAN Backup and Repository Maintenance

RMAN Reporting	9-1
Using the RMAN LIST Command.....	9-1
RMAN Reports	9-2
Reports of Obsolete Backups.....	9-2
Crosschecks of RMAN Backups	9-4
Deleting RMAN Backups	9-5
Summary of RMAN Methods for Deleting Backups	9-5
Removal of Backups with the DELETE Command.....	9-6
Advantage of Using DELETE Instead of Operating System Commands.....	9-6
Deletion of Obsolete Backups	9-6
Deletion of Expired Backups	9-6
Deletion of Archived Redo Logs That Are Already Backed Up	9-7
Deletion of Files After Backup Using BACKUP ... DELETE INPUT	9-7
BACKUP... DELETE INPUT With Corrupt Archive Logs.....	9-8
How RMAN Deletes Backup Records from the RMAN Repository	9-8
Behavior of DELETE Command When the Repository and Media Do Not Correspond	9-9
CHANGE... AVAILABLE and UNAVAILABLE with RMAN Backups	9-9
Changing Retention Policy Status of RMAN Backups	9-10
Monitoring RMAN Through V\$ Views	9-10
Correlating Server Sessions with RMAN Channels.....	9-11
Matching Server Sessions with Channels When One RMAN Session Is Active	9-11
Matching Server Sessions with Channels in Multiple RMAN Sessions	9-11
Monitoring RMAN Job Progress.....	9-13
Monitoring RMAN Interaction with the Media Manager	9-15
Monitoring RMAN Job Performance	9-16

10 Managing the Recovery Catalog

Creating a Recovery Catalog	10-1
Configuring the Recovery Catalog Database.....	10-1
Planning the Size of the Recovery Catalog Schema.....	10-2
Allocating Disk Space for the Recovery Catalog Database.....	10-2
Creating the Recovery Catalog Owner	10-3
Creating the Recovery Catalog.....	10-4
Managing Target Database Records in the Recovery Catalog	10-4
Registering a Database in the Recovery Catalog	10-4
Cataloging Older Files in the Recovery Catalog	10-5
Registering Multiple Databases in a Recovery Catalog	10-6
Unregistering a Target Database from the Recovery Catalog	10-7
Resetting the Database Incarnation in the Recovery Catalog.....	10-8
Removing DELETED Records From the Recovery Catalog After Upgrade	10-9
Resynchronizing the Recovery Catalog	10-9
Types of Records Updated When Recovery Catalog is Resynchronized	10-10

Full and Partial Resynchronization	10-10
When to Resynchronize the Recovery Catalog.....	10-11
Resynchronizing After the Recovery Catalog is Unavailable	10-11
Resynchronizing in ARCHIVELOG Mode When You Back Up Infrequently.....	10-11
Resynchronizing After Physical Database Changes	10-11
Forcing a Full Resynchronization of the Recovery Catalog.....	10-12
Resynchronizing the Recovery Catalog and CONTROL_FILE_RECORD_KEEP_TIME...	10-12
Managing the Control File When You Use a Recovery Catalog	10-12
Working with RMAN Stored Scripts in the Recovery Catalog	10-13
Creating Stored Scripts: CREATE SCRIPT	10-13
Running Stored Scripts: EXECUTE SCRIPT	10-14
Displaying a Stored Script: PRINT SCRIPT	10-15
Listing Stored Scripts: LIST SCRIPT NAMES.....	10-15
Updating Stored Scripts: REPLACE SCRIPT	10-16
Deleting Stored Scripts: DELETE SCRIPT	10-16
Starting the RMAN Client and Running a Stored Script	10-16
Restrictions on Stored Script Names	10-17
Backing Up and Recovering the Recovery Catalog	10-17
Backing Up the Recovery Catalog	10-17
Back Up the Recovery Catalog Often.....	10-17
Choosing the Appropriate Method for Physical Backups	10-18
Safe Storage of the Recovery Catalog.....	10-19
Exporting the Recovery Catalog Data for Logical Backups.....	10-19
Restoring and Recovering the Recovery Catalog from Backup	10-19
Re-Creating the Recovery Catalog.....	10-19
Exporting and Importing the Recovery Catalog	10-20
Considerations When Moving Catalog Data	10-20
Exporting the Recovery Catalog	10-21
Importing the Recovery Catalog.....	10-21
Increasing Availability of the Recovery Catalog	10-22
Querying Recovery Catalog Views	10-22
Identifying Rows for a Database in the Catalog Views.....	10-23
Identifying Rows for a Database Object in the Catalog Views.....	10-23
Querying Catalog Views for the Target DB_KEY or DBID Values	10-23
Using RC_BACKUP_FILES and DBMS_RCVMAN.SETDATABASE.....	10-24
Determining the Schema Version of the Recovery Catalog	10-25
Upgrading the Recovery Catalog	10-25
Dropping the Recovery Catalog	10-26

11 Tuning Backup and Recovery

Tuning Recovery Manager: Overview	11-1
I/O Buffer Allocation	11-2
Allocation for Disk Buffers	11-2
Allocation of Tape Buffers	11-2
Synchronous and Asynchronous I/O	11-3
Synchronous I/O: Example	11-3
Factors Affecting Backup Speed to Tape	11-4

Native Transfer Rate.....	11-4
Tape Compression	11-5
Tape Streaming.....	11-5
Physical Tape Block Size	11-5
Features and Options Used to Tune RMAN Performance	11-5
Using the RATE Parameter to Control Disk Bandwidth Usage.....	11-5
Tuning RMAN Backup Performance: Procedure	11-6
Step 1: Remove RATE Parameters from Configured and Allocated Channels	11-6
Step 2: If You Use Synchronous Disk I/O, Set DBWR_IO_SLAVES.....	11-6
Step 3: If You Fail to Allocate Shared Memory, Set LARGE_POOL_SIZE	11-6
Step 4: Tune RMAN Tape Streaming Performance Bottlenecks	11-7
Using BACKUP... VALIDATE To Distinguish Between Tape and Disk Bottlenecks....	11-7
Using Multiplexing to Improve Tape Streaming with Disk Bottlenecks.....	11-7
Using Incremental Backups to Improve Backup Performance With Tape Bottlenecks	11-7
Step 5: Query V\$ Views to Identify Bottlenecks.....	11-8
Identifying Bottlenecks with Synchronous I/O	11-8
Identifying Bottlenecks with Asynchronous I/O.....	11-8
Instance Recovery Performance Tuning: Fast-Start Fault Recovery	11-9
Understanding Instance Recovery.....	11-9
Cache Recovery (Rolling Forward)	11-9
Transaction Recovery (Rolling Back)	11-9
Checkpointing and Cache Recovery	11-9
How Checkpoints Affect Performance	11-10
Configuring the Duration of Cache Recovery: FAST_START_MTTR_TARGET	11-10
Practical Values for FAST_START_MTTR_TARGET	11-10
Reducing Checkpoint Frequency to Optimize Runtime Performance	11-11
Monitoring Cache Recovery with V\$INSTANCE_RECOVERY	11-11
Tuning FAST_START_MTTR_TARGET and Using MTTR Advisor	11-12
Calibrate the FAST_START_MTTR_TARGET.....	11-12
Determine the Practical Range for FAST_START_MTTR_TARGET.....	11-13
Evaluate Different Target Values with MTTR Advisor.....	11-14
Determine Optimal Size for Redo Logs.....	11-15

12 Recovery Manager Troubleshooting

Interpreting RMAN Message Output	12-1
Identifying Types of Message Output	12-1
Recognizing RMAN Error Message Stacks	12-2
Identifying Error Codes	12-2
RMAN Error Message Numbers	12-3
ORA-19511: Media Manager Errors	12-3
Interpreting RMAN Error Stacks.....	12-5
Interpreting RMAN Errors: Example.....	12-5
Interpreting Server Errors: Example	12-6
Interpreting SBT 2.0 Media Management Errors: Example.....	12-6
Interpreting SBT 1.1 Media Management Errors: Example.....	12-6
Identifying RMAN Return Codes.....	12-7
Testing the Media Management API.....	12-7

Obtaining the sbttest Utility	12-7
Obtaining Online Documentation for the sbttest Utility	12-8
Using the sbttest Utility	12-8
Terminating an RMAN Command	12-9
Terminating the Session with ALTER SYSTEM KILL SESSION	12-9
Terminating the Session at the Operating System Level	12-10
Terminating an RMAN Session That Is Hung in the Media Manager	12-10
Components of an RMAN Session	12-10
Process Behavior During a Hung Job	12-10
Terminating an RMAN Session: Basic Steps	12-11
RMAN Troubleshooting Scenarios	12-12
After Installation of Media Manager, RMAN Channel Allocation Fails: Scenario	12-12
After Installation of Media Manager, RMAN Channel Allocation Fails: Diagnosis	12-12
After Installation of Media Manager, RMAN Channel Allocation Fails: Solution	12-13
Backup Job Is Hanging: Scenario	12-13
Backup Job Is Hanging: Diagnosis	12-14
Backup Job Is Hanging: Solution	12-14
RMAN Fails to Start RPC Call: Scenario	12-15
RMAN Fails to Start RPC Call: Diagnosis	12-15
Backup Fails with Invalid RECID Error: Scenario	12-15
Backup Fails with Invalid RECID Error: Diagnosis	12-15
Backup Fails with Invalid RECID Error: Solution 1	12-16
Backup Fails with Invalid RECID Error: Solution 2	12-17
Backup Fails Because of Control File Enqueue: Scenario	12-18
Backup Fails Because of Control File Enqueue: Diagnosis	12-18
Backup Fails Because of Control File Enqueue: Solution	12-19
RMAN Fails to Delete All Archived Logs: Scenario	12-19
RMAN Fails to Delete All Archived Logs: Diagnosis	12-20
RMAN Fails to Delete All Archived Logs: Solution	12-20
Backup Fails Because RMAN Cannot Locate an Archived Log: Scenario	12-20
Backup Fails Because RMAN Cannot Locate an Archived Log: Diagnosis	12-20
Backup Fails Because RMAN Cannot Locate an Archived Log: Solution	12-20
RMAN Does Not Recognize Character Set Name: Scenario	12-20
RMAN Does Not Recognize Character Set Name: Diagnosis	12-21
RMAN Does Not Recognize Character Set Name: Solution	12-21
RMAN Denies Logon to Target Database: Scenario	12-21
RMAN Denies Logon to Target Database: Diagnosis	12-21
RMAN Denies Logon to Target Database: Solution	12-22
Database Duplication Fails Because of Missing Log: Scenario	12-22
Database Duplication Fails Because of Missing Log: Diagnosis	12-22
Database Duplication Fails Because of Missing Log: Solution	12-22
Duplication Fails with Multiple RMAN-06023 Errors: Scenario	12-22
Duplication Fails with Multiple RMAN-06023 Errors: Diagnosis	12-23
Duplication Fails with Multiple RMAN-06023 Errors: Solution	12-23
UNKNOWN Database Name Appears in Recovery Catalog: Scenario	12-23
UNKNOWN Database Name Appears in Recovery Catalog: Diagnosis	12-23
UNKNOWN Database Name Appears in Recovery Catalog: Solution	12-23

Part III Using RMAN for Database Transport, Duplication and Migration

13 Creating and Updating Duplicate Databases with RMAN

Creating a Duplicate Database with RMAN: Overview	13-1
How Recovery Manager Duplicates a Database	13-2
RMAN DUPLICATE DATABASE: Options	13-3
RMAN DUPLICATE DATABASE: Prerequisites and Restrictions	13-4
Renaming Database Files in RMAN Duplicate Database	13-4
Renaming Control Files in RMAN DUPLICATE DATABASE	13-4
Renaming Online Logs in RMAN DUPLICATE DATABASE	13-4
Renaming Datafiles in RMAN DUPLICATE DATABASE	13-5
Preventing Filename Checking when Duplicating a Database	13-6
Renaming Tempfiles in RMAN DUPLICATE DATABASE	13-6
Skipping Tablespaces in RMAN DUPLICATE DATABASE	13-7
Skipping Read-Only Tablespaces When Duplicating a Database	13-7
Skipping OFFLINE NORMAL Tablespaces When Duplicating a Database	13-7
Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps	13-8
Task 1: Create an Oracle Password File for the Auxiliary Instance	13-8
Task 2: Establish Oracle Net Connectivity to the Auxiliary Instance	13-8
Task 3: Create an Initialization Parameter File for the Auxiliary Instance	13-8
Task 4: Start the Auxiliary Instance	13-10
Task 5: Mount or Open the Target Database	13-10
Task 6: Make Sure You Have the Necessary Backups and Archived Redo Logs	13-10
Task 7: Allocate Auxiliary Channels if Automatic Channels Are Not Configured	13-11
Creating a Duplicate Database on a Local or Remote Host	13-11
Duplicating a Database on a Remote Host with the Same Directory Structure	13-12
Duplicating a Database on a Remote Host with a Different Directory Structure	13-12
Converting Filenames with Only Initialization Parameters	13-12
Converting Filenames with Only DUPLICATE Parameters	13-13
Using RMAN DUPLICATE With SET NEWNAME: Example	13-14
Using RMAN DUPLICATE With CONFIGURE AUXNAME: Example	13-15
Creating a Duplicate Database on the Local Host	13-16
Using RMAN DUPLICATE with OMF and ASM	13-16
Initialization Parameters for RMAN DUPLICATE to OMF Storage	13-17
Duplicating a Database to ASM Storage	13-18
Using SET NEWNAME with RMAN DUPLICATE to OMF	13-18
Using RMAN DUPLICATE DATABASE: Examples	13-19
Duplicating When the Datafiles Use Inconsistent Paths: Example	13-19
RMAN DUPLICATE DATABASE From File System to ASM: Example	13-20
RMAN DUPLICATE DATABASE From ASM to ASM: Example	13-21
RMAN DUPLICATE DATABASE at a Past Point in Time: Example	13-22
Duplicating with a Client-Side Parameter File: Example	13-22
Running RMAN Client on host_dup	13-22
Running RMAN from host_src	13-22
Using DUPLICATE DATABASE to Resynchronize a Duplicate Database: Example	13-23
Using CONFIGURE AUXNAME to Name Datafiles at the Duplicate	13-23
Synchronizing the Duplicate With the Target Using Auxnames	13-24

Using RMAN Incremental Backups to Refresh a Standby Database	13-24
Using BACKUP INCREMENTAL... FROM SCN	13-24
Refreshing a Standby Database With INCREMENTAL FROM SCN Backups: Example ..	13-25
Step 1: Create the Incremental Backup	13-25
Step 2: Make the Incremental Backup Accessible at the Standby Database.....	13-25
Step 3: Catalog the Incremental Backup Files at the Standby Database	13-26
Step 4: Apply the Incremental Backup to the Standby Database.....	13-26
14 Creating Transportable Tablespace Sets from Backup with RMAN	
About Creating Transportable Tablespace Sets from Backup with RMAN	14-1
Transportable Tablespace Sets from Backup: Concepts	14-2
When to Use RMAN to Create Transportable Tablespace Sets.....	14-2
How RMAN Creates Transportable Tablespace Sets from Backup	14-3
Steps RMAN Performs in Transporting Tablespaces from Backup	14-4
Limitations of RMAN TRANSPORT TABLESPACE Command	14-6
Creating a Transportable Tablespace Set with RMAN: Procedure	14-7
Using RMAN TRANSPORT TABLESPACE: Basic Scenario	14-7
RMAN TRANSPORT TABLESPACE with UNTIL Time or SCN.....	14-8
RMAN TRANSPORT TABLESPACE: Specifying Locations for Data Pump Files	14-9
RMAN TRANSPORT TABLESPACE with Customized Initialization Parameters	14-10
RMAN TRANSPORT TABLESPACE Default Auxiliary Instance Parameters	14-10
Location of the Auxiliary Instance Parameter File.....	14-11
Using an Auxiliary Instance Parameter File: Example.....	14-11
Customize Shared Pool Size in RMAN TRANSPORT TABLESPACE.....	14-12
Customize Auxiliary Control File Location in TRANSPORT TABLESPACE.....	14-12
Customize Other Auxiliary File Locations in TRANSPORT TABLESPACE	14-12
Transport Tablespace with SET NEWNAME for Auxiliary Datafiles	14-12
Transport Tablespace with CONFIGURE AUXNAME for Auxiliary Datafiles.....	14-13
Transport Tablespace with AUXILIARY DESTINATION Parameter	14-13
Transport Tablespace and Naming Auxiliary Files with Initialization Parameters	14-14
Troubleshooting RMAN TRANSPORT TABLESPACE.....	14-14
Troubleshooting RMAN TRANSPORT TABLESPACE: Insufficient Shared Pool.....	14-14
Troubleshooting RMAN TRANSPORT TABLESPACE: Filename Conflicts	14-15
15 RMAN Cross-Platform Transportable Databases and Tablespaces	
Cross-Platform Transportable Tablespace: CONVERT DATAFILE or TABLESPACE	15-1
Using CONVERT TABLESPACE... TO PLATFORM on the Source Platform	15-2
Rules for Renaming Files with CONVERT TABLESPACE	15-3
Converting Tablespaces on the Source Host: Example	15-3
Using CONVERT DATAFILE... FROM PLATFORM on the Destination Host	15-4
Converting Datafiles on the Destination Platform: Example	15-5
Restrictions on CONVERT TABLESPACE and CONVERT DATAFILE	15-6
Cross-Platform Transportable Database: RMAN CONVERT DATABASE	15-7
Restrictions on Cross-Platform Transportable Database	15-8
Performing Cross-Platform Database Transport	15-8
Preparing for CONVERT DATABASE: Using the DBMS_TDB Package	15-8

Using DBMS_TDB.CHECK_DB to Check Database State	15-9
Using DBMS_TDB .CHECK_EXTERNAL to Identify External Objects	15-10
Using the RMAN CONVERT DATABASE Command	15-11
CONVERT DATABASE, Converting Datafiles on the Source Platform.....	15-11
CONVERT DATABASE. Converting Datafiles on the Destination Host.....	15-13
Using RMAN CONVERT to Copy Files Between ASM and Non-ASM Storage	15-16
Copying Datafiles To ASM Using CONVERT DATAFILE: Example	15-16
Copying Datafiles From ASM Using CONVERT DATAFILE: Example	15-17
Copying Tablespaces From ASM With CONVERT TABLESPACE: Example.....	15-17

16 Migrating Databases To and From ASM with Recovery Manager

Migrating a Database into ASM.....	16-1
Limitation on ASM Migration with Transportable Tablespaces.....	16-1
Preparing to Migrate a Database to ASM.....	16-1
Determine Names of Database Files	16-2
Generate RMAN Command File to Undo ASM Migration.....	16-2
Disk-Based Migration of a Database to ASM.....	16-2
Migrating the Flash Recovery Area to ASM	16-5
Setting Initialization Parameters for Flash Recovery Area in ASM.....	16-5
Migrating the Control File to an ASM Flash Recovery Area	16-6
Changing Flashback Log Location to ASM Flash Recovery Area	16-8
Migrating Online Logs to ASM Flash Recovery Area	16-8
Migrating Existing Backups to ASM Flash Recovery Area	16-9
Migrating a Database from ASM to Non-ASM Storage	16-9
PL/SQL Scripts Used in Migrating to ASM Storage	16-10
Generating ASM-to-Non-ASM Storage Migration Script	16-10
Migrating Online Logs of Primary Database to ASM.....	16-10

Part IV Performing User-Managed Backup and Recovery

17 Making User-Managed Backups

Querying V\$ Views to Obtain Backup Information	17-1
Listing Database Files Before a Backup	17-1
Determining Datafile Status for Online Tablespace Backups.....	17-2
Making User-Managed Backups of the Whole Database.....	17-3
Making Consistent Whole Database Backups	17-3
Making User-Managed Backups of Offline Tablespaces and Datafiles	17-4
Making User-Managed Backups of Online Tablespaces and Datafiles	17-5
Making User-Managed Backups of Online Read/Write Tablespaces	17-5
Making Multiple User-Managed Backups of Online Read/Write Tablespaces	17-6
Backing Up Online Tablespaces in Parallel.....	17-6
Backing Up Online Tablespaces Serially	17-7
Ending a Backup After an Instance Failure or SHUTDOWN ABORT.....	17-7
Ending Backup Mode with the ALTER DATABASE END BACKUP Statement.....	17-8
Ending Backup Mode with the SQL*Plus RECOVER Command	17-9
Making User-Managed Backups of Read-Only Tablespaces.....	17-9

Making User-Managed Backups of the Control File	17-10
Backing Up the Control File to a Binary File.....	17-10
Backing Up the Control File to a Trace File.....	17-11
Backing Up the Control File to a Trace File: Example	17-11
Making User-Managed Backups of Archived Redo Logs	17-12
Making User-Managed Backups in SUSPEND Mode	17-13
About the Suspend/Resume Feature.....	17-13
Making Backups in a Suspended Database.....	17-13
Making User-Managed Backups to Raw Devices	17-15
Backing Up to Raw Devices on UNIX.....	17-15
Backing Up with the dd utility on UNIX: Examples	17-16
Backing Up to Raw Devices on Windows	17-17
Backing Up with OCOPY: Example	17-17
Specifying the -b and -r Options for OCOPY: Example	17-18
Verifying User-Managed Backups	17-18
Testing the Restore of Backups	17-18
Running the DBVERIFY Utility	17-18
Making Logical Backups with Oracle Export Utilities	17-19
Making User-Managed Backups of Miscellaneous Oracle Files	17-19
Keeping Records of Current and Backup Database Files	17-20
Recording the Locations of Datafiles, Control Files, and Online Redo Logs	17-20
Recording the Locations of Archived Redo Logs.....	17-20
Recording the Locations and Dates of Backup Files	17-21

18 Performing User-Managed Database Flashback and Recovery

User-Managed Flashback Features of Oracle	18-1
Performing Flashback Database with SQL*Plus.....	18-2
About User-Managed Restore Operations	18-2
Determining Which Datafiles Require Recovery	18-3
Restoring Datafiles and Archived Redo Logs	18-4
Restoring Datafiles with Operating System Utilities	18-4
Restoring Archived Redo Logs with Operating System Utilities	18-5
Restoring Control Files	18-6
Restore Lost Copy of a Multiplexed Control File.....	18-6
Copying a Multiplexed Control File to a Default Location	18-6
Copying a Multiplexed Control File to a Nondefault Location	18-6
Restore Control File from Backup After Loss of All Current Control Files.....	18-7
Restoring a Backup Control File to the Default Location	18-8
Restoring a Backup Control File to a Nondefault Location	18-9
Create New Control File After Losing All Current and Backup Control Files	18-9
About User-Managed Media Recovery	18-11
Preconditions of Performing User-Managed Recovery.....	18-11
Applying Logs Automatically with the RECOVER Command	18-11
Automating Recovery with SET AUTORECOVERY.....	18-11
Automating Recovery with the AUTOMATIC Option of the RECOVER Command.	18-12
Recovering When Archived Logs Are in the Default Location.....	18-13
Recovering When Archived Logs Are in a Nondefault Location	18-13

Resetting the Archived Log Destination.....	18-14
Overriding the Archived Log Destination	18-14
Responding to Unsuccessful Application of Redo Logs	18-15
Interrupting User-Managed Media Recovery.....	18-15
Performing Complete User-Managed Media Recovery.....	18-15
Performing Closed Database Recovery	18-16
Preparing for Closed Database Recovery.....	18-16
Restoring Backups of the Damaged or Missing Files	18-16
Recovering the Database.....	18-17
Performing Datafile Recovery in an Open Database	18-18
Preparing for Open Database Recovery	18-19
Restoring Backups of the Inaccessible Datafiles.....	18-19
Recovering Offline Tablespaces in an Open Database	18-19
Performing User-Managed Database Point-in-Time Recovery	18-20
Preparing for Incomplete Recovery.....	18-20
Restoring Datafiles Before Performing Incomplete Recovery	18-21
Performing Cancel-Based Incomplete Recovery	18-22
Performing Time-Based or Change-Based Incomplete Recovery.....	18-23
Opening the Database with the RESETLOGS Option	18-24
About Opening with the RESETLOGS Option.....	18-24
Executing the ALTER DATABASE OPEN Statements.....	18-26
Checking the Alert Log After a RESETLOGS Operation	18-27
Recovering a Database in NOARCHIVELOG Mode.....	18-27
Restoring a NOARCHIVELOG Database to its Default Location	18-27
Restoring a NOARCHIVELOG Database to a New Location	18-28
Controlling Parallel Media Recovery.....	18-29

19 Advanced User-Managed Recovery Scenarios

Recovering After the Loss of Datafiles: Scenarios.....	19-1
Losing Datafiles in NOARCHIVELOG Mode	19-1
Losing Datafiles in ARCHIVELOG Mode.....	19-2
Recovering Through an Added Datafile with a Backup Control File: Scenario.....	19-2
Re-Creating Datafiles When Backups Are Unavailable: Scenario	19-3
Recovering Through RESETLOGS with Created Control File: Scenario	19-4
Recovering NOLOGGING Tables and Indexes: Scenario	19-4
Recovering Read-Only Tablespaces with a Backup Control File: Scenario.....	19-5
Recovery of Read-Only or Slow Media with a Backup Control File	19-5
Recovery of Read-Only Files with a Re-Created Control File	19-5
Media Recovery of Transportable Tablespaces: Scenario.....	19-6
Recovering After the Loss of Online Redo Log Files: Scenarios	19-6
Recovering After Losing a Member of a Multiplexed Online Redo Log Group	19-7
Recovering After the Loss of All Members of an Online Redo Log Group.....	19-8
Losing an Inactive Online Redo Log Group	19-9
Losing an Active Online Redo Log Group	19-10
Loss of Multiple Redo Log Groups	19-11
Recovering After the Loss of Archived Redo Log Files: Scenario	19-11
Recovering from a Dropped Table: Scenario	19-12

Performing Media Recovery in a Distributed Environment: Scenario	19-13
Coordinating Time-Based and Change-Based Distributed Database Recovery.....	19-13
Dropping a Database with SQL*Plus.....	19-14

20 Performing User-Managed TSPITR

Introduction to User-Managed Tablespace Point-in-Time Recovery	20-1
TSPITR Terminology	20-2
TSPITR Methods.....	20-2
Preparing for User-Managed Tablespace Point-in-Time Recovery: Basic Steps	20-3
Step 1: Review TSPITR Requirements	20-3
Step 2: Identify All of the Files in the Recovery and Auxiliary Set Tablespaces	20-4
Step 3: Determine Whether Objects Will Be Lost	20-5
Step 4: Choose a Method for Connecting to the Auxiliary Instance.....	20-5
Step 5: Create an Oracle Password File for the Auxiliary Instance	20-5
Step 6: Create the Initialization Parameter File for the Auxiliary Instance	20-6
Restoring and Recovering the Auxiliary Database in User-Managed TSPITR: Basic Steps..	20-7
Restoring and Recovering the Auxiliary Database on the Same Host	20-8
Restoring the Auxiliary Database on a Different Host with the Same Path Names	20-9
Restoring the Auxiliary Database on a Different Host with Different Path Names	20-10
Performing User-Managed TSPITR with Transportable Tablespaces	20-10
Step 1: Unplugging the Tablespaces from the Auxiliary Database	20-11
Step 2: Transporting the Tablespaces into the Primary Database.....	20-11
Performing Partial TSPITR of Partitioned Tables	20-12
Step 1: Create a Table on the Primary Database for Each Partition Being Recovered	20-13
Step 2: Drop the Indexes on the Partition Being Recovered	20-13
Step 3: Exchange Partitions with Standalone Tables	20-13
Step 4: Drop the Recovery Set Tablespace	20-13
Step 5: Create Tables at Auxiliary Database.....	20-13
Step 6: Drop Indexes on Partitions Being Recovered	20-13
Step 7: Exchange Partitions with Standalone Tables on the Auxiliary Database	20-14
Step 8: Transport the Recovery Set Tablespaces.....	20-14
Step 9: Exchange Partitions with Standalone Tables on the Primary Database.....	20-14
Step 10: Back Up the Recovered Tablespaces in the Primary Database	20-14
Performing User-Managed TSPITR of Partitioned Tables With a Dropped Partition	20-14
Step 1: Find the Low and High Range of the Partition that Was Dropped	20-15
Step 2: Create a Temporary Table.....	20-15
Step 3: Delete Records From the Partitioned Table.....	20-15
Step 4: Drop the Recovery Set Tablespace	20-15
Step 5: Create Tables at the Auxiliary Database	20-15
Step 6: Drop Indexes on Partitions Being Recovered	20-15
Step 7: Exchange Partitions with Standalone Tables	20-15
Step 8: Transport the Recovery Set Tablespaces	20-16
Step 9: Insert Standalone Tables into Partitioned Tables	20-16
Step 10: Back Up the Recovered Tablespaces in the Primary Database	20-16
Performing User-Managed TSPITR of Partitioned Tables When a Partition Has Split.....	20-16
Step 1: Drop the Lower of the Two Partitions at the Primary Database	20-16
Step 2: Follow Same Procedure as for Partial TSPITR of Partitioned Tablespaces.....	20-17

21 Troubleshooting User-Managed Media Recovery

About User-Managed Media Recovery Problems.....	21-1
Investigating the Media Recovery Problem: Phase 1.....	21-3
Trying to Fix the Recovery Problem Without Corrupting Blocks: Phase 2.....	21-4
Deciding Whether to Allow Recovery to Corrupt Blocks: Phase 3.....	21-5
Allowing Recovery to Corrupt Blocks: Phase 4.....	21-6
Performing Trial Recovery.....	21-6
How Trial Recovery Works.....	21-7
Executing the RECOVER ... TEST Statement.....	21-7

Index

Preface

This preface contains these topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documentation](#)
- [Conventions](#)

Audience

Backup and Recovery Advanced User's Guide is intended for database administrators who perform the following tasks:

- Back up, restore, and recover Oracle databases
- Perform maintenance on backups of database files

To use this document, you need to know the following:

- Relational database concepts and basic database administration as described in *Oracle Database Concepts* and the *Oracle Database Administrator's Guide*
- Basic backup and recovery concepts and strategies as described in *Oracle Database Backup and Recovery Basics*
- The operating system environment under which you are running the database

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an

otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documentation

For more information, see these Oracle resources:

- *Oracle Database Backup and Recovery Basics*
- *Oracle Database Backup and Recovery Reference*
- *Oracle Database Utilities*

You can access information about the Backup Solutions Program at

<http://www.oracle.com/technology/deploy/availability>

Many books in the documentation set use the sample schemas of the seed database, which is installed by default when you install Oracle. Refer to *Oracle Database Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Oracle error message documentation is only available in HTML. If you only have access to the Oracle Documentation CD, you can browse the error messages by range. Once you find the specific range, use your browser's "find in page" feature to locate the specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

What's New in Backup and Recovery?

This section describes new features of Recovery Manager in Oracle Database 10g Release 2, and provides pointers to additional information on each feature. For information on new features in Oracle Database 10g Release 1 and previous releases, refer to the documentation for those releases.

Oracle Database 10g Release 2 New Features in Backup and Recovery

The new features for this release greatly increase the manageability of backup and recovery with RMAN. Backup and recovery contains improvements to security, ease of use and performance. In addition, transporting data across platforms is more automated and simplified.

These features include:

- RMAN Encrypted Backups

RMAN now supports several forms of encryption for backups created as backup sets, whether on disk or on tape. Encryption can be based upon passwords provided through RMAN or transparent encryption capabilities based upon the Oracle Encryption Wallet. Once configured, existing RMAN backup procedures take advantage of encryption features with no change.

See Also: ["RMAN Encrypted Backups"](#) on page 6-7 for information on using encrypted backups

- Flashback Database Enhancements

Flashback Database can now reverse the effects of `OPEN RESETLOGS` operations, returning a database to points in time in ancestor or even sibling incarnations. This allows its use in many more data recovery scenarios. It also integrates with guaranteed restore points.

See Also: *Oracle Database Backup and Recovery Basics* for more details on configuring Flashback Database

- Restore Points

Restore Points are aliases for SCNs, which eliminate the need to manually research and record SCNs or timestamps to use for Flashback Database and Flashback Table operations.

See Also: *Oracle Database Backup and Recovery Basics* for details about restore points

- **Guaranteed Restore Points**

Guaranteed restore points ensure that RMAN `FLASHBACK DATABASE` can be used to return your database to a specific point in time. Using guaranteed restore points instead of regular logging for Flashback Database uses disk space more efficiently and reduces performance impact of flashback logging when the only requirement is return to a specific point in time. Used in this way, guaranteed restore points provide an efficient alternative to a storage snapshot.

Guaranteed restore points can also be used with normal Flashback Database logging, to guarantee `FLASHBACK DATABASE` works to any time as far back as the guaranteed restore point.

See Also: *Oracle Database Backup and Recovery Basics* for details about guaranteed restore points

- **Incremental Roll Forward of Database Copy**

RMAN incremental backups can now be used to update a standby database with changes from a primary since a given SCN.

See Also: "[Using RMAN Incremental Backups to Refresh a Standby Database](#)" on page 13-24 for details on using incremental backups to roll forward a standby database

- **Easy Conversion of Physical Standby Database to a Reporting Database**

Easy conversion of a physical standby database to a reporting database and back to a standby is now possible, because Flashback Database can now reverse the activation of a standby database. A guaranteed restore point retains the state of the standby before activation, and after reporting the DBA can flash back the standby to that guaranteed restore point, use incremental backups to update the standby with changes from the primary during reporting, and resume managed recovery.

See Also: *Oracle Data Guard Concepts and Administration* for more details on this scenario

- **Database Transport Across Same Endian Platforms**

RMAN now supports the `CONVERT DATABASE` command, which can prepare a whole database for transport to a new platform that uses the same endian format. Database transport across platforms provides a faster and easier way to move databases from one platform to another than previous solutions requiring the use of Data Pump.

See Also: "[Performing Cross-Platform Database Transport](#)" on page 15-8 for details on transporting databases across platforms using RMAN

- **Transportable Tablespaces from Backup**

RMAN now automates the creation of transportable tablespace sets using backups instead of the datafiles of the running database. With a single RMAN command, you can now create transportable sets without making the source datafiles read-only.

See Also: [Chapter 14, "Creating Transportable Tablespace Sets from Backup with RMAN"](#) for information on creating transportable tablespace sets from backup

- **Unused Block Compression**

RMAN now creates more compact backups of datafiles, by skipping datafile blocks that are not currently used to store data. In previous releases, RMAN only supported NULL compression, which skipped space in datafiles that had never been allocated. No extra action is required on the part of the DBA to use this feature.

See Also: *Oracle Database Backup and Recovery Reference* for details about when RMAN uses Unused Block Compression for backups

- **Temporary Datafiles Are Re-Created on RMAN Recovery**

Temporary datafiles that belong to locally managed temporary tablespaces are automatically re-created during database recovery. This eliminates the need to manually create temporary tablespaces after recovery.

See Also: *Oracle Database Backup and Recovery Basics* for details about how temporary files are re-created during RMAN recovery

- **Support for Backup Vaulting in Media Managers**

When used with a media manager that supports backup vaulting, RMAN RESTORE... PREVIEW now reports any backups that are currently stored remotely, and RMAN RESTORE... PREVIEW RECALL now initiates retrieval of vaulted backups for use in an actual RESTORE operation.

See Also: *Oracle Database Backup and Recovery Basics* for details on RESTORE... PREVIEW support for backup vaulting

- **Backup and Recovery Enhancements in Enterprise Manager**

Enterprise Manager now includes backup validation, enhanced backup reporting and scheduling, and automated creation and management of recovery catalog databases.

See Also: *Oracle Database 2 Day DBA* for details about Enterprise Manager enhancements related to backup and recovery

Part I

Recovery Manager Advanced Architecture and Concepts

Part I describes the architecture of the RMAN environment and introduces basic concepts. This part contains these chapters:

- [Chapter 1, "Recovery Manager Architecture"](#)
- [Chapter 2, "RMAN Backups Concepts"](#)
- [Chapter 3, "RMAN Recovery Concepts"](#)

Recovery Manager Architecture

This chapter describes the Recovery Manager (RMAN) interface and the basic components of the RMAN environment.

This chapter contains these topics:

- [About the RMAN Environment](#)
- [RMAN Command Line Client](#)
- [RMAN Repository](#)
- [Media Management](#)

About the RMAN Environment

Recovery Manager (RMAN) is a client application that performs backup and recovery operations. The **Recovery Manager environment** consists of the various applications and databases that play a role in a backup and recovery strategy.

[Table 1–1](#) lists possible components of the RMAN environment.

Table 1–1 Components of the RMAN Environment

Component	Description	Required?
Target database	The control files, datafiles, and optional archived redo logs that RMAN is in charge of backing up or restoring. RMAN uses the target database control file to gather metadata about the target database and to store information about its own operations. The work of backup and recovery is performed by server sessions running on the target database.	Yes
RMAN client	The client application that manages backup and recovery operations for a target database. The RMAN client can use Oracle Net to connect to a target database, so it can be located on any host that is connected to the target host through Oracle Net.	Yes
Recovery catalog database	A database containing the recovery catalog schema, which contains the metadata that RMAN uses to perform its backup and recovery operations.	No
Recovery catalog schema	The user within the recovery catalog database that owns the metadata tables maintained by RMAN. RMAN periodically propagates metadata from the target database control file into the recovery catalog.	No

Table 1–1 (Cont.) Components of the RMAN Environment

Component	Description	Required?
Standby database	A copy of the primary database that is updated using archived logs created by the primary database. RMAN can create or back up a standby database. You can fail over to the standby database if the primary database goes down.	No
Duplicate database	A copy of the primary database that you can use for testing purposes.	No
Media management application	A vendor-specific application that allows RMAN to back up to a storage system such as tape.	No
Media management catalog	A vendor-specific repository of information about a media management application.	No
Enterprise Manager	A browser-based interface to the database, including backup and recovery through RMAN.	No

The only required components in an RMAN environment are the target database and the RMAN client, but most real-world configurations are more complicated. One might use an RMAN client connecting to multiple media managers and multiple target, recovery catalog, and auxiliary databases, all accessed through Enterprise Manager.

RMAN Session Architecture

The RMAN client application directs database server sessions to perform all backup and recovery tasks. The meaning of "session" in this sense depends on the operating system. For example, on UNIX, a server session corresponds to a server process, while on Windows it corresponds to a thread within the database service.

The RMAN client itself does not perform backup, restore, or recovery operations. When you connect the RMAN client to a target database, RMAN allocates server sessions on the target instance and directs them to perform the operations. The RMAN client uses internal, undocumented PL/SQL packages to communicate with the target database and recovery catalog.

RMAN Command Line Client

Use the RMAN command line client to enter commands that you can use to manage all aspects of backup and recovery operations.

Even when you use the backup and recovery features in Enterprise Manager that are built on top of RMAN, an RMAN client executes behind the scenes.

Note: All RMAN commands for Oracle release 8.1 and higher also work in Oracle Database 10g Release 2.

How RMAN Compiles and Executes Commands

RMAN processes most commands in the two phases discussed in this section:

Compilation Phase

During the compilation phase, RMAN determines which objects the command will access (for example, resolving a tablespace name into its component datafiles). Then,

RMAN constructs a sequence of remote procedure calls (RPCs) that instruct the server sessions on the target database to perform the desired operation.

Execution Phase

During the execution phase, RMAN sends the RPC calls to the target database, monitors their progress, and collects the results. If more than one channel is allocated, then RMAN can execute certain commands in parallel so that all of the channels' target database sessions are concurrently executing an RPC call.

Issuing RMAN Commands

RMAN uses a command language interpreter (CLI) that can execute commands in interactive or batch mode.

Entering Commands at the RMAN Prompt

To run RMAN commands interactively, start RMAN and then type commands into the command-line interface. For example, you can start RMAN from the UNIX command shell and then execute interactive commands as follows:

```
% rman TARGET SYS/oracle@trgt CATALOG rman/cat@catdb
```

After the RMAN prompt is displayed, you can enter commands such as the following:

```
RMAN> BACKUP DATABASE;
```

Using RMAN with Command Files

A **command file** is a text file which contains RMAN commands as you would enter them at the command line. You can run the a command file by specifying its name on the command line. The contents of the command file will be interpreted as if entered at the command line. If the LOG command line argument is specified, RMAN directs output to the named log file. Command files are one way to automate scheduled backups through an operating system job control facility.

In this example, a sample RMAN script is placed into a command file called `commandfile.rcv`. You can run this file from the operating system command line and write the output into the log file `outfile.txt` as follows:

```
% rman TARGET / CATALOG rman/cat@catdb CMDFILE commandfile.rcv LOG outfile.txt
```

See Also: *Oracle Database Backup and Recovery Reference* for more information about RMAN command line options

Stored Scripts

A **stored script** is a block of RMAN job commands that is stored in the recovery catalog. By storing scripts in the recovery catalog, the script is available to any RMAN client that connects to the catalog and the target database. Stored scripts can be associated with a single database in the catalog, or they can be global stored scripts, which can be executed against any target database in the catalog.

See Also: ["Working with RMAN Stored Scripts in the Recovery Catalog"](#) on page 10-13 for more on stored scripts. Also refer to the sample scripts in the `?/rdms/demo` directory.

Commands Valid Only in RUN Blocks

There are RMAN commands which are only valid in RUN blocks. These typically involve setting up the environment within which the statements in the RUN block will execute. Typical of this group are `ALLOCATE CHANNEL` and `SET NEWNAME FOR DATAFILE`. Using these commands outside of a RUN block will generate an error.

See Also: The syntax diagrams for the RUN command in *Oracle Database Backup and Recovery Reference* regarding which commands are valid in RUN blocks

Commands Not Valid in RUN Blocks

There are a number of RMAN commands which cannot be used in RUN blocks. Typically these are used to control the RMAN environment (connecting to different databases, or configuring RMAN defaults), or to manage or query the recovery catalog (including creating and using stored scripts). Here are some examples:

- `CONNECT`
- `CONFIGURE`
- `CREATE CATALOG`, `DROP CATALOG`, `UPGRADE CATALOG`
- `CREATE SCRIPT`, `DELETE SCRIPT`, `REPLACE SCRIPT`
- `LIST`
- `REPORT`

You can include these commands inside command files, as long as they are not wrapped inside a RUN block. You cannot use them inside a stored script from the recovery catalog, because the contents of a stored script are executed within a RUN block.

See Also: The syntax diagrams for the RUN command in *Oracle Database Backup and Recovery Reference* regarding which commands are valid in RUN blocks

Directing RMAN Output to a Log File

When you run RMAN in command line mode, it sends the output to the terminal. If you specify the `LOG` option, then RMAN writes the output to a specified log file instead.

Output for currently executing RMAN jobs is also stored in the `V$RMAN_OUTPUT` view, which reads only from memory (that is, the information is not stored in the control file). The `V$RMAN_STATUS` view contains metadata about jobs in progress as well as completed jobs. The metadata for completed jobs is stored in the control file.

RMAN Repository

The RMAN repository is the collection of metadata about the target databases that RMAN uses for backup, recovery, and maintenance. RMAN always stores this information in records in the control file. The version of this information in the control file is the authoritative record of RMAN's backups of your database. This is one reason why protecting your control file is an important part of your backup strategy. RMAN can conduct all necessary backup and recovery operations using just the control file to store the RMAN repository information, and maintains all records necessary to meet your configured retention policy.

You can also create a **recovery catalog**, an external Oracle database in which to store this information. The control file has finite space for records of backup activities, while a recovery catalog can store a much longer history. The added complexity of operating a recovery catalog database can be offset by the convenience of having the extended backup history available if you have to do a recovery that goes further back in time than the history in the control file.

There are also a few features of RMAN that only function when you use a recovery catalog. For example, RMAN stored scripts are stored in the recovery catalog, so commands related to them require the use of a recovery catalog. Other RMAN commands are specifically related to managing the recovery catalog and so are not available (and not needed) if RMAN is not connected to a recovery catalog.

The recovery catalog's version of the RMAN repository is maintained solely by RMAN. The target instance never accesses it directly. RMAN propagates information about the database structure, archived redo logs, backup sets, and datafile copies into the recovery catalog from the target database's control file after any operation that updates the repository, and also before certain operations.

See Also: *Oracle Database Backup and Recovery Basics* for details on how to manage the RMAN repository, and [Chapter 10, "Managing the Recovery Catalog"](#) to learn more about features specific to the recovery catalog

Storage of the RMAN Repository in the Control File

Because most information in the recovery catalog is also available in the target database's control file, RMAN supports an operational mode in which it uses the target database control file instead of a recovery catalog. This mode is especially appropriate for small databases where installation and administration of a separate recovery catalog database is burdensome. The only RMAN feature that is not supported in NOCATALOG mode is stored scripts.

Types of Records in the Control File

When you do not use a recovery catalog, the control file is the exclusive source of information about backups and copies as well as other relevant information. The control file contains two types of records: **circular reuse records** and **noncircular reuse records**.

Circular Reuse Records Circular reuse records contain noncritical information that is eligible to be overwritten if the need arises. These records contain information that is continually generated by the database. Circular reuse records are arranged in a logical ring. When all available record slots are full, the database either expands the control file to make room for a new record or overwrites the oldest record. The `CONTROL_FILE_RECORD_KEEP_TIME` initialization parameter specifies the minimum age in days of a record before it can be reused.

See Also: *Oracle Database Backup and Recovery Basics* to learn how to manage the handling of circular reuse records

Noncircular Reuse Records Noncircular reuse records contain critical information that does not change often and cannot be overwritten. Some examples of information in noncircular reuse records include datafiles, online redo logs, and redo threads.

Recovery Without a Recovery catalog

To make it easier to restore and recover the database without using a recovery catalog, Oracle recommends that you:

- Enable the **control file autobackup** feature, which causes RMAN to automatically back up the control file, and also enables RMAN to restore the control file autobackup without access to a repository
- **Keep a record of your DBID**, which you may need to recover your database in the event that you lose your control file
- Use a minimum of two multiplexed or mirrored control files **on separate disks**
- Keep all Recovery Manager backup logs.

If you lose the current control files, then you can restore a control file autobackup even if you do not use a recovery catalog.

See Also: ["Control File and Server Parameter File Autobackups"](#) on page 2-28 to learn about disaster recovery using control file autobackups

Storage of the RMAN Repository in the Recovery Catalog

It is recommended that you store the recovery catalog in a dedicated database. If you store the recovery catalog alongside other data in some other database, then if you lose that other database you will lose your recovery catalog as well. This will make your recovery more difficult.

Registration of Databases in the Recovery Catalog

The enrolling of a database in a recovery catalog is called **registration**. You can register more than one target database in the same recovery catalog. For example, you can register databases `prod1`, `prod2`, and `prod3` in a single catalog owned by `catowner` in the database `catdb`. Because RMAN distinguishes databases by unique database identifier (DBID), each database registered in a given catalog must have a unique DBID.

See Also: ["Registering a Database in the Recovery Catalog"](#) on page 10-4, and *Oracle Database Utilities* to learn how to use the `DBNEWID` utility to change the DBID of a database

Contents of the Recovery Catalog

The recovery catalog contains information about RMAN operations, including:

- Datafile and archived redo log backup sets and backup pieces
- Datafile copies
- Archived redo logs and their copies
- Tablespaces and datafiles on the target database
- Stored scripts, which are named user-created sequences of RMAN commands
- Persistent RMAN configuration settings

Resynchronization of the Recovery Catalog

The recovery catalog obtains crucial RMAN metadata from the target database control file. Resynchronization of the recovery catalog ensures that the metadata that RMAN obtains from the control file stays current. Resynchronizations can be full or partial.

See Also: "[Types of Records in the Control File](#)" on page 1-5 for more information about control file records, and "[When to Resynchronize the Recovery Catalog](#)" on page 10-11 for rules on when to resynchronize

Snapshot Control File RMAN creates a **snapshot control file**, which is a temporary backup control file, in an operating system specific location each time it performs a full resynchronization. This snapshot control file ensures that RMAN has a consistent view of the control file. Because the snapshot control file is intended for RMAN's short-term use, it is not registered in the recovery catalog. RMAN records the snapshot control file checkpoint in the recovery catalog to indicate the currency of the recovery catalog.

The database server ensures that only one RMAN session accesses a snapshot control file at any point in time. This safeguard is necessary to ensure that two RMAN sessions do not interfere with each other's use of the snapshot control file.

Note: You can specify the name and location of the snapshot control file. For instructions, refer to "[Setting the Snapshot Control File Location](#)" on page 5-20.

See Also: "[Managing the Control File When You Use a Recovery Catalog](#)" on page 10-12 to learn how to resynchronize the recovery catalog, and *Oracle Database Backup and Recovery Reference* for syntax

Backups of the Recovery Catalog

A single recovery catalog is able to store information for multiple target databases. Consequently, loss of the recovery catalog can be disastrous. You should back up the recovery catalog frequently.

If the recovery catalog is destroyed and no backups of it are available, then you can partially reconstruct the catalog from the current control file or control file backups. Nevertheless, you should always aim to have a valid, recent backup of the catalog.

See Also: "[Backing Up the Recovery Catalog](#)" on page 10-17 to learn how to back up the recovery catalog

Compatibility of the Recovery Catalog

When you use RMAN with a recovery catalog in an environment where you have run past versions of the database, you can wind up with versions of the RMAN client, recovery catalog database, recovery catalog schema, and target database that all originated in different releases of the database. You will find a compatibility matrix in *Oracle Database Backup and Recovery Reference* that describes supported interoperability scenarios.

Media Management

Oracle's Media Management Layer (MML) API lets third-party vendors build a **media manager**, software that works with RMAN and the vendor's hardware to allow

backups to sequential media devices such as tape drives. The media manager handles loading, unloading and labeling of sequential media such as tapes. You must install media manager software to use RMAN with sequential media devices.

When backing up or restoring, the RMAN client connects to the target instance and directs the instance to send requests to its media manager. No direct communication occurs between the RMAN client and media manager.

Performing Backup and Restore with a Media Manager

Before performing backup or restore to a media manager, you must allocate one or more channels to handle the communication with the media manager. You can also configure default channels for use with the media manager, which will be applied for all backup and recovery tasks that use the media manager where you do not explicitly allocate channels. For a conceptual overview of channels, see "[About RMAN Channels](#)" on page 2-1. Configuring or allocating channels for backups is discussed further in "[Configuring and Allocating Channels for Use in Backups](#)" on page 6-1.

For example, this sequence of commands would configure channels for the media manager and back up the database to the media manager:

```
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 1;  
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;  
RMAN> CONFIGURE CHANNEL DEVICE TYPE sbt PARMS 'ENV=(NSR_SERVER=bksvr1)';  
RMAN> BACKUP DATABASE;
```

When RMAN executes the `BACKUP DATABASE` command, it sends the backup request to the database server session performing the backup. The database server session identifies the output channel as a media management device and makes a request to the media manager to write the output.

RMAN does not issue specific commands to load, label, or unload tapes. When backing up, RMAN gives the media manager a stream of bytes and associates a unique name with that stream. When RMAN needs to restore the backup, it asks the media manager to retrieve the byte stream. All details of how and where that stream is stored are handled entirely by the media manager.

The media manager labels and keeps track of the tape and names of files on each tape, and automatically loads and unloads tapes, or signals an operator to do so.

Some media managers support **proxy copy** functionality, in which they handle the entire data movement between datafiles and the backup devices. Such products may use technologies such as high-speed connections between storage and media subsystems to reduce load on the primary database server. RMAN provides a list of files requiring backup or restore to the media manager, which in turn makes all decisions regarding how and when to move the data.

Backup Solutions Program

The Oracle Backup Solutions Program (BSP), part of the Oracle Partner Program, is a group of leading media manager vendors whose products are compliant with Oracle's MML specification. Several products may be available for your platform from media management vendors. For more information, you can contact your Oracle representative for a list of available products, contact individual vendors to ask them if they participate, or access the Backup Solutions Program Web site at:

<http://www.oracle.com/technology/deploy/availability>

Note that Oracle does not certify media manager vendors for compatibility with RMAN. Questions about availability, version compatibility, and functionality can only be answered by the media manager vendor, not Oracle.

RMAN Backups Concepts


This chapter describes the basic concepts involved in backing up the database with the Recovery Manager (RMAN) utility.

This chapter contains these topics:

- [About RMAN Channels](#)
- [About RMAN Backups](#)
- [Multiple Copies of RMAN Backups](#)
- [RMAN Backup Options: Naming, Sizing, and Speed](#)
- [RMAN Backup Types](#)
- [Control File and Server Parameter File Autobackups](#)
- [Backup Retention Policies](#)
- [Backup Optimization](#)
- [Restartable Backups](#)
- [Managing Backup Windows and Performance: BACKUP... DURATION](#)
- [How RMAN Responds to Backup Errors](#)
- [Tests and Integrity Checks for Backups](#)

About RMAN Channels

An RMAN **channel** represents one stream of data to a device type, and corresponds to one server session. Most backup and recovery commands in RMAN are executed by server sessions. As illustrated in [Figure 2-1](#), each channel establishes a connection from the RMAN client to a target or auxiliary database instance by starting a server session on the instance. The server session performs the backup, restore, and recovery.

Figure 2–1 Channel Allocation

You can use the `CONFIGURE CHANNEL` command to configure channels for use with disk or tape in all RMAN sessions using **automatic channel allocation**, or allocate channels manually within a `RUN` block. RMAN comes preconfigured with one `DISK` channel that you can use for backups to disk.

When you run a command that requires a channel without allocating a channel explicitly, then RMAN automatically allocates the channels with the options specified in the `CONFIGURE` command. For the `BACKUP` command, RMAN allocates only a single type of channel, such as `DISK`. For the `RESTORE` command and maintenance commands (for example, `DELETE`), RMAN allocates all necessary channels for the device types required to execute the command.

To specify the device type to use for an operation explicitly, use the `ALLOCATE CHANNEL` command, which must be used within a `RUN` block, or `ALLOCATE CHANNEL FOR MAINTENANCE`, which must be executed at the RMAN prompt.

In a Real Application Clusters configuration, there are special considerations regarding channel allocation and backups. See *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide* for more details.

How and when the `ALLOCATE CHANNEL` or `CONFIGURE CHANNEL` commands cause the media manager to allocate resources is vendor-specific. Some media managers allocate resources when you issue the command; others do not allocate resources until you open a file for reading or writing.

See Also: *Oracle Database Backup and Recovery Reference* for `ALLOCATE CHANNEL` syntax and *Oracle Database Backup and Recovery Reference* on `ALLOCATE CHANNEL FOR MAINTENANCE`

Automatic and Manual Channel Allocation

You can use the automatic channel allocation feature to configure a set of persistent, automatic channels for use in all RMAN sessions. You can use the manual channel allocation feature to specify channels for commands used within a `RUN` block.

RMAN allocates automatic channels according to the settings in these commands:

- `CONFIGURE DEVICE TYPE ... PARALLELISM`
- `CONFIGURE DEFAULT DEVICE TYPE`

- CONFIGURE CHANNEL

For example, you can issue the following commands at the RMAN prompt:

```
# since you do not manually allocate channels, RMAN uses preconfigured channels
BACKUP DATAFILE 3;
RESTORE TABLESPACE users;
```

When you run a command that requires channels, and no channels have been allocated using the `ALLOCATE` command, RMAN automatically allocates channels according to values set with the `CONFIGURE` command in the following cases:

- You use commands such as `BACKUP`, `RESTORE`, or `DELETE` outside of a `RUN` block.
- You use commands within a `RUN` block but do not allocate any channels within the `RUN` block.

You can override automatic channel allocation settings by manually allocating channels within a `RUN` block. Manual channels always override automatic channels. For example, you override automatic channel allocation when you issue a command as follows:

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt;
  BACKUP DATABASE PLUS ARCHIVELOG;
}
```

RMAN optimizes automatic channel allocation by leaving automatic channels allocated so long as each new command requires exactly the same channel configuration as the previous command. For example, RMAN can use the same preallocated channels for the following series of commands:

```
BACKUP DATAFILE 1;
BACKUP CURRENT CONTROLFILE;
BACKUP ARCHIVELOG ALL;
```

If you issue a command such as `ALLOCATE` or `CONFIGURE`, then RMAN automatically releases the preallocated channels.

See Also: "[Configuring Channels](#)" on page 5-9 to learn how to configure automatic channels

Automatic Channel Device Configuration and Parallelism

The `CONFIGURE DEVICE TYPE . . . PARALLELISM` command specifies the number of automatic channels to allocate for a specified device type. For example, if you configure parallelism to 3 for a device type, then RMAN allocates three channels for the device type when using automatic channels.

You can change a parallelism setting by issuing another `CONFIGURE DEVICE TYPE . . . PARALLELISM` command. This example configures `PARALLELISM 2` and then changes it to 3:

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 2;
CONFIGURE DEVICE TYPE DISK PARALLELISM 3;
```

The parallelism setting defines the number of channels for a device that RMAN allocates in parallel. It does not have to correspond to the actual number of channels configured for the device. For example, you can manually configure four different `sbt` channels and set `PARALLELISM` for `sbt` to 2, 1, or 10.

You can view the default setting for parallelism by running the `SHOW DEVICE TYPE` command. For example:

```
RMAN> SHOW DEVICE TYPE;
RMAN configuration parameters are:
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; #default
```

As always when the `SHOW` command is used to view the value of a parameter, RMAN includes a "#default" comment at the end of the line if the RMAN default value has not been overridden.

The following example configures the default device to `sbt` and then displays the resulting configuration using the `SHOW DEVICE TYPE` command:

```
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
new RMAN configuration parameters:
CONFIGURE DEFAULT DEVICE TYPE TO 'sbt';
new RMAN configuration parameters are successfully stored

RMAN> SHOW DEVICE TYPE;
RMAN configuration parameters are:
CONFIGURE DEVICE TYPE SBT PARALLELISM 1; # default
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; #
default
```

See Also: ["Configuring Channel Parallelism"](#) on page 5-9

Automatic Channel Default Device Types

Run the `CONFIGURE DEFAULT DEVICE TYPE` command to specify a default device type for automatic channels. For example, you may make backups to tape most of the time and only occasionally make a backup to disk. In this case, configure channels for disk and tape devices, but make `sbt` the default device type:

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 1; # configure device disk
CONFIGURE DEVICE TYPE sbt PARALLELISM 2; # configure device sbt
CONFIGURE DEFAULT DEVICE TYPE TO sbt;
```

Now, RMAN will, by default, use `sbt` channels for backups. For example, if you run the following command:

```
BACKUP TABLESPACE users;
```

RMAN only allocates channels of type `sbt` during the backup because `sbt` is the default device.

You can override the default device for backups by specifying the target device as part of the command. For example:

```
BACKUP DEVICE TYPE sbt DATABASE;
```

If the default device type is `DISK`, then the preceding command overrides this default and uses the `sbt` channel configuration. Note that this command fails unless you have configured the `sbt` device or configured `sbt` channels.

When restoring files, RMAN allocates all automatic channels according to the settings configured for each device type. The default device type configuration is irrelevant. For example, if you configure `PARALLELISM` to 3 for the default `sbt` device and `PARALLELISM` to 2 for `DISK`, then RMAN automatically allocates three `sbt` channels and two `DISK` channels during the restore.

Automatic Channel Naming Conventions

RMAN uses the following convention for channel naming: `ORA_devicetype_n`, where *devicetype* refers to the user's device type (such as `DISK` or `sbt_tape`) and *n* refers to the channel number.

Note: The `sbt` and `sbt_tape` device types are synonymous, but RMAN output always displays `sbt_tape` whether the input is `sbt` or `sbt_tape`.

For example, RMAN names the first `DISK` channel `ORA_DISK_1`, the second `ORA_DISK_2`, and so forth. RMAN names the first `sbt` channel `ORA_SBT_TAPE_1`, the second `ORA_SBT_TAPE_2`, and so forth. When you parallelize channels, RMAN always allocates channels in numerical order, starting with 1 and ending with the parallelism setting (`CONFIGURE DEVICE TYPE . . . PARALLELISM n`), as in this example:

```
ORA_SBT_TAPE_1
ORA_SBT_TAPE_2
ORA_SBT_TAPE_3
```

Automatic channel allocation also applies to maintenance commands. If RMAN allocates an automatic maintenance channel, then it uses the same naming convention as any other automatically allocated channel. If you manually allocate a maintenance channel using `ALLOCATE CHANNEL FOR MAINTENANCE`, then RMAN uses the following convention for channel naming: `ORA_MAINT_devicetype_n`, where *devicetype* refers to the user's device type (for example, `DISK` or `sbt`) and *n* refers to the channel number. For example, RMAN uses these names for two manually allocated disk channels:

```
ORA_MAINT_DISK_1
ORA_MAINT_DISK_2
```

Note that if you run the `CONFIGURE DEVICE TYPE` command to configure default settings for a device type and do not run `CONFIGURE CHANNEL` for this device type, then RMAN allocates all channels without other channel control options. For example, assume that you configure the `sbt` device and run a backup as follows:

```
CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
BACKUP DEVICE TYPE sbt DATABASE;
```

In effect, RMAN does the following:

```
RUN
{
  ALLOCATE CHANNEL ORA_SBT_TAPE_1 DEVICE TYPE sbt;
  BACKUP DATABASE;
}
```

Channel names beginning with the `ORA_` prefix are reserved by RMAN for its own use. You cannot manually allocate a channel with a name that begins with `ORA_`.

Automatic Channel Generic Configurations

The `CONFIGURE CHANNEL DEVICE TYPE` command configures generic settings that are used for all automatic channels of the specified device type. In other words, the command creates a template of settings that RMAN uses for all channels allocated on the device. For example, you can configure disk and tape channels as follows:

```
CONFIGURE CHANNEL DEVICE TYPE sbt PARMS='ENV=(NSR_SERVER=bksvr1)';
CONFIGURE CHANNEL DEVICE TYPE DISK RATE 5M FORMAT='?/oradata/%U';
```

Because you do not specify channel numbers for these channels, the channel settings are generic to all automatic channels of the specified type. The configuration acts as a template. For example, if you set `PARALLELISM` for `DISK` to 10, and the default device type is `DISK`, then RMAN allocates ten disk channels using the settings in the `CONFIGURE CHANNEL DEVICE TYPE DISK` command.

See Also: ["Configuring Channels"](#) on page 5-9

Automatic Channel-Specific Configurations

You can also configure parameters that apply to a specific automatic channel. If you are using a media manager that requires different settings on each channel, then you may find it useful to configure individual channels.

You can mix a `CONFIGURE CHANNEL` command that creates a generic configuration with a `CONFIGURE CHANNEL` command that creates a specific configuration. A generic automatic channel creates a configuration that can be used for any channel that is not explicitly configured.

For example, assume that you run these commands:

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 3;
CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE = 2M;
CONFIGURE CHANNEL 3 DEVICE TYPE DISK MAXPIECESIZE = 900K;
```

In this scenario, RMAN allocates `ORA_DISK_1` and `ORA_DISK_2` with option `MAXPIECESIZE = 2M`, using the settings for the `DISK` channel with no number. RMAN allocates `ORA_DISK_3` with `MAXPIECESIZE = 900K` because this channel was manually assigned a channel number. RMAN always allocates the number of channels specified in the parallelism parameter.

See Also: ["Configuring Channels"](#) on page 5-9

Clearing Automatic Channel Settings

You can specify the `CLEAR` option for any `CONFIGURE` command. The `CLEAR` option returns the specified configuration to its default value. Assume you run these commands:

```
CONFIGURE DEVICE TYPE DISK CLEAR;      # returns DISK to default PARALLELISM 1
                                         # and backup type to BACKUPSET
CONFIGURE DEFAULT DEVICE TYPE CLEAR;    # returns to default device type of DISK
CONFIGURE CHANNEL DEVICE TYPE sbt CLEAR; # removes all options for sbt channel
CONFIGURE CHANNEL 3 DEVICE TYPE DISK CLEAR; # removes configurations for 3rd ch.
```

Each `CONFIGURE . . . CLEAR` command removes the user-entered settings and returns the configuration to its default value.

The only way to find out the default setting for parameters set through `CONFIGURE` is to use `CONFIGURE... CLEAR` to un-set the parameter, so that it takes on the default value, and then run `SHOW ALL` to view all parameters. For any parameter for which the value is currently set to RMAN'S default, RMAN includes a `"#default"` comment at the end of that line of the output from `SHOW ALL`.

See Also: *Oracle Database Backup and Recovery Reference* for the default settings for each `CONFIGURE` command

Determining Channel Parallelism to Match Hardware Devices

RMAN can perform the I/O required for many commands in parallel, to make optimal use of your hardware resources. To perform I/O in parallel, however, the I/O must be associated with a single RMAN command, not a series of commands. For example, it can be more efficient to back up three datafiles using a command such as:

```
BACKUP DATAFILE 5,6,7;
```

rather than issuing the commands

```
BACKUP DATAFILE 5;
BACKUP DATAFILE 6;
BACKUP DATAFILE 7;
```

When all three datafiles are backed up in one command, RMAN recognizes the opportunity for parallelism and can use multiple channels to do the I/O in parallel. When three separate commands are used, RMAN can only perform the backups one at a time, regardless of available channels and I/O devices.

The number of channels available (whether allocated in a RUN block or configured in advance) for use with a device at the moment that you run a command determines whether RMAN will read from or write to that device in parallel while carrying out the command. Failing to allocate the right number of channels adversely affects RMAN performance during I/O operations.

As a rule, the number of channels used in carrying out an individual RMAN command should match the number of physical devices accessed in carrying out that command. If manually allocating channels for a command, allocate one for each device; if configuring automatic channels, configure the PARALLELISM setting appropriately.

When backing up to tape, you should allocate one channel for each tape drive. When backing up to disk, allocate one channel for each physical disk, unless you can optimize the backup for your disk topography by using multiple disk channels. Each manually allocated channel uses a separate connection to the target or auxiliary database.

The following script creates three backups sequentially: three separate BACKUP commands are used to back up one file each. Only one channel is active at any one time because only one file is being backed up in each command.

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt;
  ALLOCATE CHANNEL c2 DEVICE TYPE sbt;
  ALLOCATE CHANNEL c3 DEVICE TYPE sbt;
  BACKUP DATAFILE 5;
  BACKUP DATAFILE 6;
  BACKUP DATAFILE 7;
}
```

The following statement uses **parallelization** on the same example: one RMAN BACKUP command backs up three datafiles, with all three channels in use. The three channels are **concurrently active**—each server session copies one of the datafiles to a separate tape drive.

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt;
  ALLOCATE CHANNEL c2 DEVICE TYPE sbt;
```

```

ALLOCATE CHANNEL c3 DEVICE TYPE sbt;
BACKUP DATAFILE 5,6,7;
}

```

See Also: *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide* for information about parallelization in a Real Application Clusters (RAC) configuration.

Channel Control Options for Manual and Automatic Channels

Whether you allocate channels manually or automatically, you can use channel control commands and options to do the following:

- Control the operating system resources RMAN uses when performing RMAN operations
- Affect the degree of parallelism for a backup or restore command
- Set limits on I/O bandwidth consumption in kilobytes, megabytes, or gigabytes (ALLOCATE CHANNEL . . . RATE, CONFIGURE CHANNEL . . . RATE)
- Set limits on the size of backup pieces (the MAXPIECESIZE parameter specified on the CONFIGURE CHANNEL and ALLOCATE CHANNEL commands)
- Set limits on the size of backup sets (the MAXSETSIZE parameter specified on the BACKUP and CONFIGURE commands)
- Send vendor-specific commands to the media manager (SEND)
- Specify vendor-specific parameters for the media manager (ALLOCATE CHANNEL . . . PARMS, CONFIGURE CHANNEL . . . PARMS)
- Specify which instance performs the operation (ALLOCATE CHANNEL . . . CONNECT, CONFIGURE CHANNEL . . . CONNECT)

In releases 8.1.5 and later of the database, the ALLOCATE CHANNEL command causes RMAN to contact the media manager whenever the type specified is other than DISK. In earlier releases, the ALLOCATE CHANNEL command does not cause RMAN to contact the media manager; RMAN does not call the media manager until a BACKUP, RESTORE, or RECOVER command is issued.

Note: When you specify DEVICE TYPE DISK with any version of RMAN, RMAN does not allocate operating system resources other than for the creation of the server session and does not call the media manager.

Because RMAN has one preconfigured automatic DISK channel, you do not have to manually allocate a maintenance channel when running CHANGE, CROSSCHECK, or DELETE against a disk file (that is, an ARCHIVELOG, DATAFILECOPY, or CONTROLFILECOPY).

A maintenance channel is useful only for a maintenance task; you cannot use it as an input or output channel for a backup or restore.

See Also: *Oracle Database Backup and Recovery Reference* for ALLOCATE CHANNEL syntax, and *Oracle Database Backup and Recovery Reference* for CONFIGURE syntax

Channel Failover

A `BACKUP` command is decomposed into multiple independent backup steps by RMAN. Each independent step can be executed on any channel allocated for the type of device used in the command. If you have multiple channels allocated, and one channel fails or encounters a problem during a backup step, then RMAN attempts to complete the work on another channel. Typically, such retrievable errors can occur when a media manager encounters problems with one of several tape drives, or when an instance fails in a RAC environment.

RMAN reports a message in `V$RMAN_OUTPUT` and in the output to the interactive session or log file when it encounters such problems, as in the following example (refer to bold text):

```
channel ORA_SBT_TAPE_1: backup set failed, re-triable on other channel
ORA-19506: failed to create sequential file, name="/bkup/63d3c3og_1_1", parms=""
ORA-27028: skgfgcre: sbtbackup returned error
ORA-19511: Error received from media manager layer, error text: failed to open
           file /bkup/63d3c3og_1_1 for backup, errno = 2
channel ORA_SBT_TAPE_2: finished piece 1 at 06-SEP-01 piece handle=5ld3blun_1_1
                       comment=API Version 2.0,MMS Version 3.2.0.0
channel ORA_SBT_TAPE_2: backup set complete, elapsed time: 00:00:04
retrying ORA_SBT_TAPE_1 failed backup step on ORA_SBT_TAPE_2
channel ORA_SBT_TAPE_2: starting full datafile backupset
channel ORA_SBT_TAPE_2: specifying datafile(s) in backupset input datafile
                       fno=00004 name=/oracle/dbs/tbs_12.f input datafile
                       fno=00017 name=/oracle/dbs/tbs_14.f
channel ORA_SBT_TAPE_2: starting piece 1 at 06-SEP-01 piece handle=5ld3buds_1_1
                       comment=API Version 2.0,MMS Version 3.2.0.0
channel ORA_SBT_TAPE_2: backup set complete, elapsed time: 00:00:06
```

Note that if RMAN is executing a script, then the next command in the script will not be executed if there were any errors in the previous command.

See Also: ["Interpreting RMAN Message Output"](#) on page 12-1 to learn more about RMAN message and error reporting

About RMAN Backups

When you execute the `BACKUP` command in RMAN, you create one or more backup sets or image copies. By default, RMAN creates backup sets regardless of whether the destination is disk or a media manager.

About Image Copies

An **image copy** is an exact copy of a single datafile, archived redo log file, or control file. Image copies are not stored in an RMAN-specific format. They are identical to the results of copying a file with operating system commands. RMAN can use image copies during RMAN restore and recover operations, and you can also use image copies with non-RMAN restore and recovery techniques.

To create image copies and have them recorded in the RMAN repository, run the `RMAN BACKUP AS COPY` command (or, alternatively, configure the default backup type for disk as image copies using `CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO COPY` before performing a backup). A database server session is used to create the copy, and the server session also performs actions such as validating the blocks in the file and recording the image copy in the RMAN repository.

You can also use an operating system command such as the UNIX `dd` command to create image copies, though these will not be validated, nor are they recorded in the RMAN repository. You can use the `CATALOG` command to add image copies created with native operating system tools in the RMAN repository.

Using RMAN-Created Image Copies

If you run a `RESTORE` command, then by default RMAN restores a datafile or control file to its original location by copying an image copy backup to that location. Image copies are chosen over backup sets because of the extra overhead of reading through an entire backup set in search of files to be restored.

However, if you need to restore and recover a current datafile, and if you have an image copy of the datafile available on disk, then you do not actually need to have RMAN copy the image copy back to its old location. You can instead have the database use the image copy in place, as a replacement for the datafile to be restored. The `SWITCH` command updates the RMAN repository indicate that the image copy should now be treated as the current datafile. Issuing the `SWITCH` command in this case is equivalent to issuing the SQL statement `ALTER DATABASE RENAME FILE`. You can then perform recovery on the copy.

User-Managed Image Copies

RMAN can use image copies created by mechanisms outside of RMAN, such as native operating system file copy commands or third-party utilities that leave image copies of files on disk. These copies are known as **user-managed copies** or **operating system copies**.

The RMAN `CATALOG` command causes RMAN to inspect an existing image copy and enter its metadata into the RMAN repository. Once cataloged, these files can be used like any other backup with the `RESTORE` or `SWITCH` commands.

Some sites store their datafiles on mirrored disk volumes, which permit the creation of image copies by **breaking a mirror**. After you have broken the mirror, you can notify RMAN of the existence of a new user-managed copy, thus making it a candidate for a backup operation. You must notify RMAN when the copy is no longer available, by using the `CHANGE . . . UNCATALOG` command. In this example, before **resilvering the mirror** (not including other copies of the broken mirror), you must use a `CHANGE . . . UNCATALOG` command to update the recovery catalog and indicate that this copy is no longer available.

See Also:

- *Oracle Database Backup and Recovery Basics* to learn how to catalog datafile and archived log image copies
- ["Making Split Mirror Backups with RMAN"](#) on page 6-3
- *Oracle Database Backup and Recovery Reference* for `CHANGE` syntax

About Proxy Copies

During a **proxy copy**, RMAN turns over control of the data transfer to a media manager that supports this feature. Proxy copy can only be used with media managers that support it, not with disk. The `PROXY` option of the `BACKUP` command specifies that a backup should be a proxy copy.

For each file that you attempt to back up with the `BACKUP PROXY` command, RMAN queries the media manager to determine whether it can perform a proxy copy. If the media manager cannot proxy copy the file, then RMAN backs the file up as if the

PROXY option had not been used. (Use the PROXY ONLY option to force RMAN to fail if a proxy copy cannot be performed.)

Proxy copy can be used with datafiles or archived redo logs, as shown in these examples:

```
BACKUP DEVICE TYPE sbt PROXY DATAFILE 3;
BACKUP DEVICE TYPE sbt PROXY ONLY DATABASE;
BACKUP DEVICE TYPE sbt PROXY ONLY ARCHIVELOG ALL;
```

The examples assume that sbt channels have been configured with the appropriate parameters.

Note that control files are never backed up with proxy copy. If the PROXY option is specified on an operation backing up a control file, it is silently ignored for the purposes of backing up the control file.

See Also:

- *Oracle Database Reference* for more information about V\$PROXY_DATAFILE and V\$PROXY_ARCHIVEDLOG
- *Oracle Database Backup and Recovery Reference* for the BACKUP command and the PROXY option

Storage of Backups on Disk and Tape

RMAN can create backups on disk or a third-party media device such as a tape drive. If you specify DEVICE TYPE DISK, then your backups are created on disk, in the file name space of the target instance that is creating the backup. You can make a backup on any device that can store a datafile.

To create backups on non-disk media, such as tape, you must use third-party media management software, and allocate channels with device types, such as SBT, that are supported by that software.

Backups of Archived Logs

There are several features of RMAN backups specific to backups of archived redo logs.

Deletion of Archived Logs After Backups

RMAN can delete one or all copies of archived logs from disk after backing them up to backup sets. If you specify the DELETE INPUT option, then RMAN backs up exactly one copy of each specified log sequence number and thread from an archive destination to tape, and then deletes the specific file it backed up while leaving the other copies on disk. If you specify the DELETE ALL INPUT option, then RMAN backs up exactly one copy of each specified log sequence number and thread, and then deletes that log from all archive destinations. Note that there are special considerations related to deletion of archived redo logs in standby database configurations. See *Oracle Data Guard Concepts and Administration* for details.

Backup Failover for Archived Redo Logs

RMAN's **archived redo log failover** allows RMAN to complete a backup even when some archived log destinations are missing logs or have logs with corrupt blocks. If at least one log corresponding to a given log sequence and thread is available in any of the archiving destinations, then RMAN tries to back it up. If RMAN finds a corrupt block in a log file during backup, it searches other destinations for a copy of that log without corrupt blocks.

By default, RMAN only backs up one copy of each distinct log sequence number. For example, assume that you archive logs 121 through 124 to two archiving destinations: /arch1 and /arch2. The control file contains archived log records as follows:

Sequence	Filename in /arch1	Filename in /arch2
121	/arch1/archive1_121.arc	/arch2/archive1_121.arc
122	/arch1/archive1_122.arc	/arch2/archive1_122.arc
123	/arch1/archive1_123.arc	/arch2/archive1_123.arc
124	/arch1/archive1_124.arc	/arch2/archive1_124.arc

However, unknown to RMAN, a user deletes logs 122 and 124 from the /arch1 directory. Then, you run the following backup:

```
BACKUP ARCHIVELOG FROM SEQUENCE 121 UNTIL SEQUENCE 125;
```


With failover, RMAN completes the backup, using logs 122 and 124 in /arch2.

Multiplexed Backup Sets

When creating backup sets, you can **multiplex** files. In this case, RMAN simultaneously reads multiple files from disk and then writes their blocks into the same backup set. (Image copies, by contrast, are never multiplexed.) For example, RMAN can read from two datafiles simultaneously, and then combine the blocks from these datafiles into a single backup piece.

As [Figure 2-2](#) illustrates, RMAN can back up three datafiles into a backup set that contains only one backup piece. This backup piece contains the intermingled data blocks of the three input files.

Figure 2-2 Datafile Multiplexing



Algorithm for Multiplexed Backups

RMAN multiplexing is determined by the following algorithm:

1. The **number of files in each backup set** is determined by computing the lesser of the following values:
 - The default number of files in a backup set (16 for archived logs, and 4 for datafiles)
 - The number of files read by each channel.
2. The **level of multiplexing** is determined by the lesser of the following values:
 - The number of files in the backup set (as calculated by the preceding step)
 - The default number of files that RMAN reads simultaneously on a single channel (8 files for each channel)

Assume that you are backing up twelve datafiles with one RMAN channel. The number of files in each backup set is 4. To determine the level of multiplexing, compare this value to 8 and take the lesser, which is 4. Because the level of multiplexing is 4, the channel includes blocks from four separate datafiles into each backup set.

See Also:

- ["I/O Buffer Allocation"](#) on page 11-2 to learn how multiplexing affects allocation of disk buffers during backups
- *Oracle Database Backup and Recovery Reference* for BACKUP syntax

Multiplexing by the Media Manager

Media manager multiplexing occurs when the media manager writes the concurrent output from multiple RMAN channels to a single sequential device, such as a tape drive.

Caution: Although media manager multiplexing can sometimes provide a performance benefit during backup, it can have a negative impact on restore performance. Oracle recommends using RMAN multiplexing instead of using multiplexing by the media manager.

Manual Parallelization of Backups

When you configure PARALLELISM to greater than 1 or manually allocate multiple channels, RMAN writes multiple backup sets or image copies in parallel. The channels divide the work of backing up the specified files.

Note: You cannot stripe a single backup set across multiple channels.

By default, RMAN determines which channels should back up which database files. You can use the CHANNEL option of the BACKUP command to manually assign a channel to back up specified files. This example shows a parallel backup to the default disk location that uses the default automatic DISK channels:

```
BACKUP
  (DATAFILE 1,2,3
  CHANNEL ORA_DISK_1)
  (DATAFILECOPY '/tmp/system01.dbf', '/tmp/tools01.dbf')
```

```

CHANNEL ORA_DISK_2)
(ARCHIVELOG FROM SEQUENCE 100 UNTIL SEQUENCE 102 THREAD 1
CHANNEL ORA_DISK_3);

```

You can also manually allocate channels as in the following example:

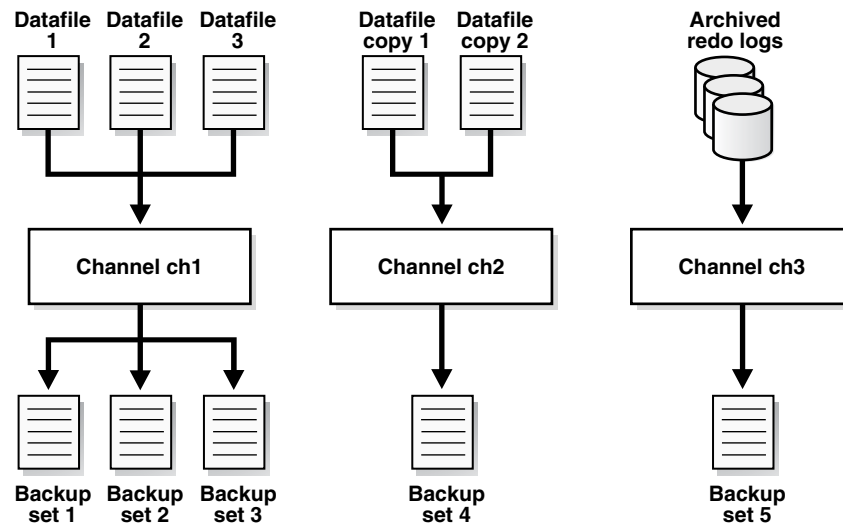
```

RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS="ENV=(BACKUP_SERVER=tape_server1)";
  ALLOCATE CHANNEL c2 DEVICE TYPE sbt PARMS="ENV=(BACKUP_SERVER=tape_server2)";
  ALLOCATE CHANNEL c3 DEVICE TYPE sbt PARMS="ENV=(BACKUP_SERVER=tape_server3)";
  BACKUP
  (DATAFILE 1,2,3
  CHANNEL c1)
  (DATAFILECOPY '/tmp/system01.dbf', '/tmp/tools01.dbf'
  CHANNEL c2)
  (ARCHIVELOG FROM SEQUENCE 100 UNTIL SEQUENCE 102 THREAD 1
  CHANNEL c3);
}

```

Figure 2–3 shows an example of parallelization in which channel ch1 backs up datafiles, channel ch2 backs up datafile copies, and channel ch3 backs up logs.

Figure 2–3 Parallelization of Backups



See Also:

- ["Determining Channel Parallelism to Match Hardware Devices"](#) on page 2-7 for an overview of how allocated channels affect parallelization
- ["Determining How Channels Distribute a Backup Workload: Example"](#) on page 6-17 to learn how to parallelize backups
- *Oracle Database Backup and Recovery Reference* for reference material on the CHANNEL parameter of the BACKUP command

Multiple Copies of RMAN Backups

In RMAN, there are two ways to make multiple, identical copies of backups:

- Duplex your backups within the `BACKUP AS BACKUPSET` command, in which case RMAN creates more than one copy of each backup set
- Back up your files as backup sets or image copies, and then back up the backup sets or image copies using the `RMAN BACKUP BACKUPSET` or `BACKUPCOPY` commands.

Duplexed Backup Sets

When backing up datafiles, archived redo log files, server parameter files and control files into backup pieces, RMAN can **duplex** the backup set, producing up to four identical copies of each backup piece in the backup set on different backup destinations with one `BACKUP` command. (Note that duplexing is not supported for backup operations that produce image copies.)

There are three ways to specify duplexing of backup sets when using the `BACKUP` command:

- Specify a default level of duplexing with `CONFIGURE . . . BACKUP COPIES`
All backup commands that back up data into backup sets will be affected if you use this option, unless you specify different duplexing options for a command using `SET BACKUP COPIES` or provide a `COPIES` option for the `BACKUP` command.
- Use `SET BACKUP COPIES` in a `RUN` block
All commands in the `RUN` block will be affected, overriding any `CONFIGURE... BACKUP COPIES` setting, except those where you provide a `COPIES` option as part of the `BACKUP` command.
- Provide a `COPIES` option to the `BACKUP` command
For this specific `BACKUP` command, files will be duplexed to produce the number of copies you specify.

The `FORMAT` option of the `BACKUP` command specifies the destinations to be used when performing duplexed backups. You can specify up to 4 values for the `FORMAT` option. RMAN uses the second, third, and fourth values only when `BACKUP COPIES`, `SET BACKUP COPIES`, or `CONFIGURE . . . BACKUP COPIES` is specified. The following example creates 3 copies of the backup of datafile 7:

```
BACKUP DEVICE TYPE DISK COPIES 3 DATAFILE 7 FORMAT
'/tmp/%U', '?/oradata/%U', '?/%U';
```

RMAN places the first copy of each backup piece in `/tmp`, the second in `?/oradata`, and the third in the Oracle home. Note that RMAN does not produce 3 backup sets, each with a different unique backup set key. Rather, RMAN produces one backup set with a unique key, and generates 3 identical copies of each backup piece in the set, as shown in this sample `LIST` output:

```
List of Backup Sets
=====
```

```
BS Key   Type LV Size
-----
```

```
1        Full    64K
```

```
List of Datafiles in backup set 1
```

```
File LV Type Ckp SCN    Ckp Time  Name
-----
```

```
7        Full 98410    08-FEB-03 /oracle/oradata/trgt/tools01.dbf
```

```

Backup Set Copy #1 of backup set 1
Device Type Elapsed Time Completion Time Tag
-----
DISK          00:00:01      08-FEB-03      TAG20030208T152314

```

List of Backup Pieces for backup set 1 Copy #1

```

BP Key Pc# Status      Piece Name
-----
1      1  AVAILABLE  /tmp/01dg9tb2_1_1

```

```

Backup Set Copy #2 of backup set 1
Device Type Elapsed Time Completion Time Tag
-----
DISK          00:00:01      08-FEB-03      TAG20030208T152314

```

List of Backup Pieces for backup set 1 Copy #2

```

BP Key Pc# Status      Piece Name
-----
2      1  AVAILABLE  /oracle/oradata/01dg9tb2_1_2

```

```

Backup Set Copy #3 of backup set 1
Device Type Elapsed Time Completion Time Tag
-----
DISK          00:00:01      08-FEB-03      TAG20030208T152314

```

List of Backup Pieces for backup set 1 Copy #3

```

BP Key Pc# Status      Piece Name
-----
3      1  AVAILABLE  /oracle/01dg9tb2_1_3

```

When choosing which `FORMAT` value to use for each backup piece, RMAN uses the first format value for copy number 1, the second format value for copy number 2, and so forth. If the number of format values exceeds the number of copies, then the extra formats are not used. If the number of format values is less than the number of copies, then RMAN reuses the format values, starting with the first one.

See Also:

- ["Duplexing Backup Sets"](#) on page 6-2 to learn how to duplex backups
- *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax
- *Oracle Database Backup and Recovery Reference* for `SET` syntax

Backups of Backup Sets

The RMAN `BACKUP BACKUPSET` command backs up previously created backup sets. Only backup sets that were created on device type `DISK` can be backed up using RMAN. They can be backed up to any available device type.

Note: RMAN issues an error if you attempt to run `BACKUP AS COPY BACKUPSET`.

The `BACKUP BACKUPSET` command uses the default disk channel to copy backup sets from disk to disk. To back up from disk to tape, you must either configure or manually allocate a non-disk channel.

Uses for Backups of Backup Sets

The `BACKUP BACKUPSET` command is a useful way to spread backups among multiple media. For example, you can execute the following `BACKUP` command weekly as part of the production backup schedule:

```
# makes backup sets on disk
BACKUP DEVICE TYPE DISK AS BACKUPSET DATABASE PLUS ARCHIVELOG;
BACKUP DEVICE TYPE sbt BACKUPSET ALL; # copies backup sets on disk to tape
```

Note: Backups to `sbt` that use automatic channels require that you first run the `CONFIGURE DEVICE TYPE sbt` command.

In this way, you ensure that all your backups exist on both disk and tape. You can also duplex backups of backup sets, as in this example:

```
BACKUP COPIES 2 DEVICE TYPE sbt BACKUPSET ALL;
```

(Again, control file autobackups are never duplexed.)

You can also use `BACKUP BACKUPSET` to manage backup space allocation. For example, to keep more recent backups on disk and older backups only on tape, you can regularly run the following command:

```
BACKUP DEVICE TYPE sbt BACKUPSET COMPLETED BEFORE 'SYSDATE-7' DELETE INPUT;
```

This command backs up backup sets that were created more than a week ago from disk to tape, and then deletes them from disk. Note that `DELETE INPUT` here is equivalent to `DELETE ALL INPUT`; RMAN deletes all existing copies of the backup set. If you duplexed a backup to four locations, then RMAN deletes all four copies of the pieces in the backup set.

Backup Optimization When Backing Up Backup Sets

If **backup optimization** is enabled when you issue the command to back up a backup set, and if the identical backup set has already been backed up to the same device type, then RMAN skips the backup of this backup set. For example, when backup optimization is turned on, the following command backs up to tape only those backup sets not already backed up on device type `sbt`:

```
BACKUP DEVICE TYPE sbt BACKUPSET ALL;
```

Backup Failover When Backing Up Backup Sets

When backing up backup sets, if RMAN discovers that one copy of a backup set is corrupted or missing, then it searches for other copies of the same backup set, based on the RMAN repository records about the backup set. This behavior is similar to the behavior of RMAN when backing up archived redo logs that exist in multiple archiving destinations.

For example, assume that backup set with key 872 contains three backup pieces, and that `BACKUP COPIES 3` was issued so that three copies of each backup piece were created, each on a different file system. Also assume that some copies have been deleted or corrupted, so that the following table describes the current status of the backup copies:

Backup Piece Number	Copy Number of the Piece	Status of Copy
1	1	Corrupted

Backup Piece Number	Copy Number of the Piece	Status of Copy
1	2	Intact
1	3	Corrupted
2	1	Missing
2	2	Corrupted
2	3	Intact
3	1	Intact
3	2	Corrupted
3	3	Missing

The following command will cause RMAN to perform automatic failover:

```
BACKUP BACKUPSET 872;
```

RMAN copies only the backup pieces listed as "Intact" in the preceding table in its backup set.

Backups of Image Copies

You can use the following commands to back up existing image copies of database files either as backup sets or as image copies:

- `BACKUP AS COPY COPY OF DATABASE`
- `BACKUP AS BACKUPSET COPY OF TABLESPACE tablespace_name`
- `BACKUP AS BACKUPSET COPY OF DATAFILE datafile`

When using these commands, there must already exist an image copy of every datafile specified in the command. If there are multiple copies of a datafile, the latest one is used. If you specify a tablespace or the whole database, RMAN issues an error if there are any datafiles in the database or tablespace specified for which there are no image copy backups in the RMAN repository.

RMAN Backup Options: Naming, Sizing, and Speed

Recovery Manager provides a number of options to control filenames, sizes of backups and speed during backup.

Generating Filenames for RMAN Backup Pieces

You can either let RMAN determine a unique name for backup pieces or use the `FORMAT` parameter to specify a name. For example, enter:

```
BACKUP TABLESPACE users;
```

RMAN automatically generates unique names for the backup pieces in the default backup location.

The `FORMAT` parameter provides substitution variables that you can use to generate unique filenames. For example, you can run a command as follows:

```
BACKUP TABLESPACE users FORMAT = '/tmp/users_%u%p%c';
```

As described in "[Manual Parallelization of Backups](#)" on page 2-13, you can specify up to four `FORMAT` values. RMAN uses the second, third, and fourth values only when you run `BACKUP COPIES`, `SET BACKUP COPIES`, or `CONFIGURE . . . BACKUP COPIES`.

Note: If you use a media manager, then check your vendor documentation for restrictions on `FORMAT` such as the size of the name, the naming conventions, and so forth.

See Also: *Oracle Database Backup and Recovery Reference* for descriptions of the `FORMAT` parameter and the substitution variables

Generating Filenames for RMAN Image Copies

RMAN can generate filenames for image copies using either `FORMAT` or `DB_FILE_NAME_CONVERT`.

Using `FORMAT` with `RMAN BACKUP AS COPY`

As with backup pieces, `FORMAT` variables are also used to specify the names of image copies. The default format `%U` is defined differently for image copies than for backup pieces. RMAN produces image copies of three types of files: datafiles, control files, and archived logs. The following table describes the meaning of `%U` for each type of file.

Type of File	Meaning of %U
Datafile	data-D-%d_id-%I_TS-%N_FNO-%f_%u
Archived log	arch-D_%d-id-%I_S-%e_T-%h_A-%a_%u
Control file	cf-D_%d-id-%I_%u

Using `DB_FILE_NAME_CONVERT` With `BACKUP AS COPY`

When creating image copies, you can also name the output copies with the `DB_FILE_NAME_CONVERT` option of the `BACKUP` command. This parameter works identically to the initialization parameter `DB_FILE_NAME_CONVERT`. Pairs of filename prefixes are provided to change the names of the output files. If a file is not converted by any of the pairs, then RMAN uses the `FORMAT` specification; if no `FORMAT` is specified, then RMAN uses the default format `%U`.

For example, you can run the following command to copy the datafiles whose filename is prefixed with `/maindisk/oradata/users` so that they are prefixed with `/backups/users_ts`:

```
BACKUP AS COPY TABLESPACE users
  DB_FILE_NAME_CONVERT=('/maindisk/oradata/users', '/backups/users_ts');
```

Note: There are restrictions on using `DB_FILE_NAME_CONVERT` with `BACKUP AS COPY` to convert Oracle Managed Files (OMF) filenames. See *Oracle Database Backup and Recovery Reference* for details on these restrictions.

Using Tags to Identify RMAN Backups

You can assign a user-specified character string called a **tag** to backup sets and image copies (either copies created by RMAN or copies created by an operating system

utility). A tag is a symbolic name for a backup set or file copy, such as `weekly_backup`. You can specify the tag rather than the filename when executing the `RESTORE` or `CHANGE` command. The maximum length of a tag is 30 bytes.

Default RMAN Backup Tag Format

If you do not specify a tag, then RMAN creates a default tag for backups (except for control file autobackups) in the format `TAGYYYYMMDDTHHMMSS`, where `YYYY` is the year, `MM` is the month, `DD` is the day, `HH` is the hour (in 24-hour format), `MM` is the minutes, and `SS` is the seconds. For example, a backup of datafile 1 may receive the tag `TAG20030208T133437`. The date and time refer to when RMAN started the backup, in the time zone of the instance performing the backup. If multiple backup sets are created by one `BACKUP` command, then each backup piece is assigned the same default tag.

How Tags Are Applied

When applied to a backup set, a tag applies to a specific copy of the backup set. If you do not duplex a backup set, that is, make multiple identical copies of it, then the backup set has just one tag. For example,

```
BACKUP COPIES 1 DATAFILE 7 TAG foo;
```

creates one backup set with tag `FOO`. Tags are stored in uppercase, regardless of the case used when entering them. However, you can back up this backup set and give this new copy of the backup set the tag `BAR`. So, the backup set has two identical copies: one tagged `FOO` and the other tagged `BAR`.

When applied to image copies, a tag applies to each individual image copy. For example, you run the following command:

```
# Back up as image copies the datafiles of tablespaces users and tools
# all copies get the TAG 'users_tools'
BACKUP AS COPY TAG users_tools TABLESPACE users, tools;
```

You can also copy an image copy with a specific tag, and give the output copy a different tag, as in the following example:

```
# Create new copies of all image copies of the database that have the tag
# 'full_cold_copy', and give the new copies the tag 'new_full_cold_copy'
BACKUP AS COPY
  COPY OF DATABASE
  FROM TAG=full_cold_copy
  TAG=new_full_cold_copy;

# Create backup sets of the image copy of tablespace users that has the tag
# 'monday_users', and of tablespace SYSTEM which has the tag 'monday_system'.
# All these new backup # sets receive the tag 'for_audit'.
BACKUP AS BACKUPSET TAG for_audit
  COPY OF TABLESPACE users FROM TAG monday_users
  TABLESPACE SYSTEM FROM TAG monday_system;
```

Uniqueness of Backup Tags

Tags do not need to be unique, so multiple backup sets or image copies can have the same tag, for example, `weekly_backup`. When you specify that a datafile should be restored from backups that have a specific tag, and more than one backup of the requested file has the desired tag, RMAN restores the most recent backup that has the desired tag (within any other constraints of the restore command, of course, such as a point in time).

Tags can indicate the intended purpose or usage of different classes of backups or copies. For example, you can tag datafile copies that you intend to use in a SWITCH differently (*for_switch_only*) from file copies that should be used only for RESTORE (*for_restore_only*).

Note: If you specify the FROM tag option to the RESTORE or SWITCH command, then RMAN considers *only* backup sets and image copies with a matching tag when choosing which backup to use

See Also: *Oracle Database Backup and Recovery Reference* for SWITCH syntax, and *Oracle Database Backup and Recovery Reference* for RESTORE syntax

Size of Backup Pieces

RMAN will, by default, put the entire contents of a backup set into one backup piece, regardless of the size of the backup set. If you are backing up to a file system or media manager that has a limit on the maximum file size that can be created, then you may need to restrict the size of backup pieces that RMAN will create. To restrict the size of each backup piece, specify the MAXPIECESIZE option of the CONFIGURE CHANNEL or ALLOCATE CHANNEL commands. This option limits backup piece size to the specified number of bytes. If the total size of the backup set is greater than the specified backup piece size, then multiple physical pieces will be created to hold the backup set contents.

For example, if datafile 1 is 6GB, you can still restrict the backup piece size for disk backups to 2 GB by configuring an automatic disk channel, and then run a backup as follows:

```
CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE = 2G;
BACKUP AS BACKUPSET DATAFILE 1;
```

A LIST BACKUP command reveals that RMAN created five backup pieces rather than one backup piece to conform to the MAXPIECESIZE size restriction:

```
BS Key  Type LV Size      Device Type Elapsed Time Completion Time
-----
29      Full  9728M    DISK          00:00:35    NOV 02 2002 18:29:26
List of Datafiles in backup set 29
File LV Type Ckp SCN    Ckp Time          Name
-----
1      Full 177590    NOV 02 2002 18:28:51 /oracle/oradata/trgt/system01.dbf

Backup Set Copy #1 of backup set 29
Device Type Elapsed Time Completion Time      Tag
-----
DISK          00:00:35    NOV 02 2002 18:29:26 TAG20021102T152701
```

```
List of Backup Pieces for backup set 29 Copy #1
BP Key  Pc# Status      Piece Name
-----
53      1  AVAILABLE  /oracle/dbs/10d85733_1_1
54      2  AVAILABLE  /oracle/dbs/10d85733_2_1
55      3  AVAILABLE  /oracle/dbs/10d85733_3_1
56      4  AVAILABLE  /oracle/dbs/10d85733_4_1
57      5  AVAILABLE  /oracle/dbs/10d85733_5_1
```

This option can be used for media managers that cannot manage a backup piece that spans more than one tape. For example, if a tape can hold 10GB, but the backup set being created must hold 80GB of data, then RMAN must be instructed to create backup pieces of 10GB, small enough to fit on the tapes used with the media manager. The backup set media will in this case consist of eight tapes. Media managers supporting SBT2.0 can return a value to RMAN indicating the largest supported backup piece size, which RMAN will use in planning backup activities.

See Also:

- *Oracle Database Backup and Recovery Reference* for `ALLOCATE CHANNEL` syntax
- *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax

Number and Size of Backup Sets

Use the *backupSpec* clause of the `BACKUP` command to specify the objects that you want to back up as well as specify other options. Each *backupSpec* clause produces at least one backup set. The total number and size of backup sets depends for the most part on an internal RMAN algorithm, although you can tune RMAN behavior to a certain extent with the `MAXSETSIZE` parameter.

Factors Affecting the Number and Size of Backup Sets

In determining the characteristics of the RMAN backup sets, the internal algorithm is influenced by the following factors:

- The number of input files specified in each *backupSpec* clause
- The number of channels that you allocate
- The default number of files in each backup set (4 for datafiles and 16 for archived logs)
- The default number of files read simultaneously by a single channel (8)
- The `MAXSETSIZE` parameter (specified on the `CONFIGURE` and `BACKUP` commands), which specifies a maximum backup set size

The most important rules in the algorithm for backup set creation are:

- Each allocated channel that performs work in the backup job—that is, each channel that is not idle—generates at least one backup set.

Note: RMAN writes backup pieces sequentially; striping a backup piece across multiple output devices is not supported. For example, RMAN does not simultaneously write half of a backup piece to one device and the other half to another device, nor can it write the first piece of a backup set to one device and the second piece to another device.

- RMAN always tries to divide the backup load so that all allocated channels have roughly the same amount of work to do.
- The default number of datafiles in each backup set is determined by an internal RMAN limit (16 for archived logs, 4 for datafiles).

- The number of datafiles multiplexed in a backup set is limited by the lesser of the number of files in each backup set and the default number of files read by a single channel simultaneously (8).
- The maximum size of a backup set is determined by the `MAXSETSIZE` parameter of the `CONFIGURE` or `BACKUP` command.

See Also: *Oracle Database Backup and Recovery Reference* to learn the syntax for the `backupSpec` clause, and [Chapter 11, "Tuning Backup and Recovery"](#) to learn about RMAN buffer management

Overview of the `MAXSETSIZE` Parameter

To specify the maximum size of each backup set, use the `MAXSETSIZE` parameter in the `CONFIGURE` or `BACKUP` command. By limiting the overall size of the backup set, the parameter indirectly limits the number of files in the set and can possibly force RMAN to create additional backup sets.

Specifying `MAXSETSIZE`: Example Assume that you want to back up 50 datafiles, each containing 1000 blocks. To set the maximum size of each backup set to 10 MB, use the following command:

```
BACKUP DATABASE MAXSETSIZE = 10M;
```

Caution: If a datafile being backed up is bigger than `MAXSETSIZE` then your backup will fail. Always ensure that `MAXSETSIZE` is as large as your largest datafile.

See Also: *Oracle Database Backup and Recovery Reference* for information on the `MAXSETSIZE` parameter

I/O Read Rate of Backups

By default, RMAN uses all available I/O bandwidth to read/write to disk. You can limit the I/O resources consumed by a backup job with the `RATE` option of the `ALLOCATE CHANNEL` or `CONFIGURE CHANNEL` commands. The `RATE` option specifies the maximum number of bytes for each second that RMAN reads on the channel.

For example, you can configure automatic channels to limit each channel to read 1 MB a second:

```
CONFIGURE DEVICE TYPE sbt PARALLELISM 2;
CONFIGURE DEFAULT DEVICE TYPE TO sbt;
CONFIGURE CHANNEL DEVICE TYPE sbt RATE 1M;
```

In effect, the `RATE` option throttles RMAN so that a backup job does not consume excessive I/O bandwidth on the computer.

See Also: ["Tuning RMAN Backup Performance: Procedure"](#) on page 11-6 for tips about how to optimize RMAN performance

RMAN Backup Types

As explained in [Table 2-1](#), RMAN backups can be classified in these ways:

- Full or incremental
- Open or closed

- Consistent or inconsistent

Note that these backup classifications apply only to datafile backups. Backups of other files, such as archivelogs and control files, always include the complete file and are never inconsistent.

Table 2–1 Backup Types

Backup Type	Definition
Full	<p>A backup of a datafile that includes every allocated block in the file being backed up. A full backup of a datafile can be an image copy, in which case every data block is backed up. It can also be stored in a backup set, in which case datafile blocks not in use may be skipped, according to rules in <i>Oracle Database Backup and Recovery Reference</i>.</p> <p>A full backup cannot be part of an incremental backup strategy; that is, it cannot be the parent for a subsequent incremental backup.</p>
Incremental	<p>An incremental backup is either a level 0 backup, which includes every block in the file except blocks compressed out because they have never been used, or a level 1 backup, which includes only those blocks that have been changed since the parent backup was taken.</p> <p>A level 0 incremental backup is physically identical to a full backup. The only difference is that the level 0 backup is recorded as an incremental backup in the RMAN repository, so it can be used as the parent for a level 1 backup.</p>
Open	A backup of online, read/write datafiles when the database is open.
Closed	A backup of any part of the target database when it is mounted but not open. Closed backups can be consistent or inconsistent.
Consistent	<p>A backup taken when the database is mounted (but not open) after a normal shutdown. The checkpoint SCNs in the datafile headers match the header information in the control file. None of the datafiles has changes beyond its checkpoint. Consistent backups can be restored without recovery.</p> <p>Note: If you restore a consistent backup and open the database in read/write mode without recovery, transactions after the backup are lost. You still need to perform an <code>OPEN RESETLOGS</code>.</p>
Inconsistent	<p>A backup of any part of the target database when it is open or when a crash occurred or <code>SHUTDOWN ABORT</code> was run prior to mounting.</p> <p>An inconsistent backup requires recovery to become consistent.</p>

Incremental Backups

The goal of an incremental backup is to back up only those data blocks that have changed since a previous backup. You can use RMAN to create incremental backups of datafiles, tablespaces, or the whole database.

During media recovery, RMAN examines the restored files to determine whether it can recover them with an incremental backup. If it has a choice, then RMAN always chooses incremental backups over archived logs, as applying changes at a block level is faster than reapplying individual changes.

RMAN does not need to restore a base incremental backup of a datafile in order to apply incremental backups to the datafile during recovery. For example, you can restore non-incremental image copies of the datafiles in the database, and RMAN can recover them with incremental backups.

Incremental backups allow faster daily backups, use less network bandwidth when backing up over a network, and provide better performance when tape I/O

bandwidth limits backup performance. They also allow recovery of database changes not reflected in the redo logs, such as direct load inserts. Finally, incremental backups can be used to back up NOARCHIVELOG databases, and are smaller than complete copies of the database (though they still require a clean database shutdown).

One effective strategy is to make incremental backups to disk (as image copies), and then back up these image copies to a media manager with `BACKUP AS BACKUPSET`. Then, you do not have the problem of keeping the tape streaming that sometimes occurs when making incremental backups directly to tape. Because incremental backups are not as big as full backups, you can create them on disk more easily.

Incremental Backup Algorithm

Each data block in a datafile contains a system change number (SCN), which is the SCN at which the most recent change was made to the block. During an incremental backup, RMAN reads the SCN of each data block in the input file and compares it to the checkpoint SCN of the parent incremental backup. (If block change tracking is enabled, RMAN does not read the portions of the file known to have not changed since the parent incremental backup.) If the SCN in the input data block is greater than or equal to the checkpoint SCN of the parent, then RMAN copies the block.

One consequence of this mechanism is that RMAN applies all blocks containing changed data during recovery—even if the change is to an object created with the `NOLOGGING` option. Hence, making incremental backups is a safeguard against the loss of changes made by `NOLOGGING` operations.

See Also: *Oracle Database Concepts* for more information about `NOLOGGING` mode

Multilevel Incremental Backups

RMAN can create **multilevel incremental backups**. Each incremental level is denoted by a value of 0 or 1. A level 0 incremental backup, which is the base for subsequent incremental backups, copies all blocks containing data. The only difference between a level 0 incremental backup and a full backup is that a full backup is never included in an incremental strategy.

A level 1 incremental backup can be either of the following types:

- A **differential backup**, which backs up all blocks changed after the most recent incremental backup at level 1 or 0
- A **cumulative backup**, which backs up all blocks changed after the most recent incremental backup at level 0

Incremental backups are differential by default.

Note: Cumulative backups are preferable to differential backups when recovery time is more important than disk space, because fewer incremental backups need to be applied during recovery.

The size of the backup file depends solely upon the number of blocks modified and the incremental backup level.


Differential Incremental Backups

In a differential level 1 backup, RMAN backs up all blocks that have changed since the most recent incremental backup at level 1 (cumulative or differential) or level 0. For example, in a differential level 1 backup, RMAN determines which level 1 backup

occurred most recently and backs up all blocks modified after that backup. If no level 1 is available, RMAN copies all blocks changed since the base level 0 backup.

If no level 0 backup is available, then the behavior varies with the compatibility mode setting. If compatibility is $\geq 10.0.0$, RMAN copies all blocks that have been changed since the file was created. Otherwise, RMAN behaves as it did in previous releases, by generating a level 0 backup.

Figure 2-4 Differential Incremental Backups




In the example shown in [Figure 2-4](#), the following occurs each week:

- Sunday
 - An incremental level 0 backup backs up *all* blocks that have ever been in use in this database.
- Monday - Saturday
 - On each day from Monday through Saturday, a differential incremental level 1 backup backs up all blocks that have changed since the most recent incremental backup at level 1 or 0. The Monday backup copies blocks changed since Sunday level 0 backup, the Tuesday backup copies blocks changed since the Monday level 1 backup, and so forth.

Cumulative Incremental Backups

In a cumulative level 1 backup, RMAN backs up all the blocks used since the most recent level 0 incremental backup. Cumulative incremental backups reduce the work needed for a restore by ensuring that you only need one incremental backup from any particular level. Cumulative backups require more space and time than differential backups, however, because they duplicate the work done by previous backups at the same level.

Figure 2-5 Cumulative Incremental Backups



In the example shown in [Figure 2-5](#), the following occurs each week:

- Sunday
 - An incremental level 0 backup backs up *all* blocks that have ever been in use in this database.
- Monday - Saturday
 - A cumulative incremental level 1 backup copies all blocks changed since the most recent level 0 backup. Because the most recent level 0 backup was created on Sunday, the level 1 backup on each day Monday through Saturday backs up all blocks changed since the Sunday backup.

Planning an Incremental Backup Strategy

Choose a backup scheme according to an acceptable MTTR (mean time to recover). For example, you can implement a three-level backup scheme so that a full or level 0 backup is taken monthly, a cumulative level 1 is taken weekly, and a differential level 1 is taken daily. In this scheme, you never have to apply more than a day's worth of redo for complete recovery.

When deciding how often to take full or level 0 backups, a good rule of thumb is to take a new level 0 whenever 20% or more of the data has changed. If the rate of change to your database is predictable, then you can observe the size of your incremental backups to determine when a new level 0 is appropriate. The following query displays the number of blocks written to a backup set for each datafile with at least 50% of its blocks backed up:

```
SELECT FILE#, INCREMENTAL_LEVEL, COMPLETION_TIME, BLOCKS, DATAFILE_BLOCKS
FROM V$BACKUP_DATAFILE
WHERE INCREMENTAL_LEVEL > 0
AND BLOCKS / DATAFILE_BLOCKS > .5
ORDER BY COMPLETION_TIME;
```

Compare the number of blocks in differential or cumulative backups to a base level 0 backup. For example, if you only create level 1 cumulative backups, then take a new level 0 backup when the most recent level 1 backup is about half of the size of the base level 0 backup.

See Also: *Oracle Database Backup and Recovery Basics* to learn how make incremental backups

Control File and Server Parameter File Autobackups

Having recent backups of your control file and server parameter file is extremely valuable in many recovery situations. To increase the likelihood that you will have such backups, the database supports control file and server parameter file autobackups. RMAN can automatically back up the control file and server parameter file (SPFILE) in situations in which the RMAN repository data for your database has been updated in a way that affects RMAN's ability to restore your database.

With a control file autobackup, RMAN can recover the database even if the current control file, recovery catalog, and server parameter file are inaccessible. Because the path used to store the autobackup follows a well-known format, RMAN can search for and restore the server parameter file from that autobackup.

After you have started the instance with the restored server parameter file, RMAN can restore the control file from the autobackup. After you mount the control file, use the RMAN repository in the mounted control file to restore the datafiles.

A control file autobackup lets you restore the RMAN repository contained in the control file when the control file is lost and you have no recovery catalog. You do not need a recovery catalog or target database control file to restore the control file autobackup. For example, you can issue:

```
RESTORE CONTROLFILE FROM AUTOBACKUP;
```

After you restore and mount the control file, you have the backup information necessary to restore and recover the database. You can connect to the target instance in NOCATALOG mode and recover the database. You can also create a new recovery catalog and register the target database. The RMAN repository records in the control file will be copied to the new recovery catalog.

The automatic backup of the control file occurs independently of any backup of the current control file explicitly requested as part of your backup command. You can turn the autobackup feature on or off by running the following commands:

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;  
CONFIGURE CONTROLFILE AUTOBACKUP OFF;
```

Oracle recommends that `CONFIGURE CONTROLFILE AUTOBACKUP` be set to `ON`.

How RMAN Performs Control File Autobackups

The first channel allocated during the backup job creates the autobackup and places it into its own backup set; for autobackups after database structural changes, the default disk channel makes the backup. If a server parameter file is used, it is backed up in the same backup set as the control file during a control file autobackup.

After the control file autobackup completes, the database writes a message containing the complete path of the backup piece and the device type to the alert log.

The RMAN behavior when the `BACKUP` command includes datafile 1 depends on the `CONFIGURE CONTROLFILE AUTOBACKUP` setting. If control file autobackups are `ON`

and the backup includes datafile 1, RMAN writes the control file and SPFILE to a separate autobackup backup set. If control file autobackups are OFF and the backup includes datafile 1, then RMAN includes the current control file and SPFILE in the same backup set as the datafiles.

The control file autobackup filename has a default format of %F for all device types, so that RMAN can guess the file location and restore it without a repository. The substitution variable %F is defined in the description of the CONFIGURE command in *Oracle Database Backup and Recovery Basics*. You can specify a different format with the CONFIGURE CONTROLFILE AUTOBACKUP FORMAT command. All autobackup formats must include the %F variable.

The SET CONTROLFILE AUTOBACKUP FORMAT command, which you can specify either within a RUN block or at the RMAN prompt, overrides the configured autobackup format in the session only. The order of precedence is:

1. SET within a RUN block
2. SET at RMAN prompt
3. CONFIGURE CONTROLFILE AUTOBACKUP FORMAT

You can configure the autobackup format even when CONFIGURE CONTROLFILE AUTOBACKUP is set to OFF, but RMAN does not generate autobackups in this case. For RMAN to make autobackups, you must set CONFIGURE CONTROLFILE AUTOBACKUP to ON.

See Also:

- *Oracle Database Backup and Recovery Reference* for BACKUP syntax
- *Oracle Database Backup and Recovery Reference* for RESTORE syntax

When RMAN Performs Control File Autobackups

By default, control file autobackups are turned off, and no control file autobackups are performed. If CONFIGURE CONTROLFILE AUTOBACKUP is ON, then RMAN automatically backs up the control file and the current server parameter file (if used to start up the database) in one of two circumstances: when a successful backup must be recorded in the RMAN repository, and when a structural change to the database affects the contents of the control file which therefore must be backed up.

Control File Autobackups After Backup Activities

This means that the control file is backed up in the following situations:

- After every BACKUP command issued at the RMAN prompt.
- At the end of a RUN block, if the last command in the block was BACKUP.
- Whenever a BACKUP command within a RUN block is followed by a command that is not BACKUP.

The first channel allocated during the backup job creates the autobackup and places it into its own backup set. The control file autobackup may be written to disk or tape.

Control File Autobackups After Database Structural Changes

The control file is also automatically backed up after database structural changes such as adding a new tablespace, altering the state of a tablespace or datafile (for example, bringing it online), adding a new online redo log, renaming a file, adding a new redo

thread, and so on. Losing this information would compromise your ability to recover the database.

This backup is performed by the server process itself, rather than one of the RMAN channels. This type of autobackup, unlike autobackups that occur after a successful backup, is always created on disk. You can use `CONFIGURE CONTROLFILE AUTOBACKUP FOR DEVICE TYPE DISK` to set the location for this disk based control file autobackup. Note that a failure of the automatic control file autobackup after a structural change never causes the associated structural change to fail. For example, if you add a datafile, and if the resulting control file autobackup fails, then the datafile addition is still successful.

Backup Retention Policies

You can use the `CONFIGURE RETENTION POLICY` command to create a persistent and automatic **backup retention policy**. When a backup retention policy is in effect, RMAN considers backups of datafiles and control files as **obsolete**, that is, no longer needed for recovery, according to criteria that you specify in the `CONFIGURE` command. You can then use the `REPORT OBSOLETE` command to view obsolete files and `DELETE OBSOLETE` to delete them.

As you produce backups over time, older backups become obsolete as they are no longer needed to satisfy the retention policy. RMAN can identify the obsolete files for you, but it does not automatically delete them. You must use the `DELETE OBSOLETE` command to delete files that are no longer needed to satisfy the retention policy.

If you have a flash recovery area configured, however, then the database automatically deletes unnecessary files from the flash recovery area based on its internal disk quota rules. The disk quota rules are distinct from the backup retention policy rules, but the database will never delete files in violation of the retention policy to satisfy the disk quota.

The term **obsolete** does not mean the same as **expired**. A backup is obsolete when `REPORT OBSOLETE` or `DELETE OBSOLETE` determines, based on the user-defined retention policy, that it is not needed for recovery. A backup is considered expired only when RMAN performs a crosscheck and cannot find the file. In short, *obsolete* means "not needed," whereas *expired* means "not found."

From the perspective of a retention policy, a datafile backup is a full or level 0 backup of an individual datafile or control file. It does not matter whether the backup is a datafile image copy, a proxy copy, or part of a backup set. For datafile copies and proxy copies, if RMAN determines that the copy or proxy copy is not needed, then the copy or proxy copy can be deleted. For datafile backups in backup sets, RMAN cannot delete the backup set until all of the individual datafile backups within the backup set are obsolete.

Besides affecting full or level 0 datafile and control file backups, the retention policy affects archived redo logs and level 1 incremental backups. First, RMAN decides which datafile and control file backups are obsolete. Then, RMAN considers as obsolete all archived logs and incremental level 1 backups that are not needed to recover the oldest datafile or control file backup that must be retained.

Note: RMAN cannot implement an automatic retention policy if backups are deleted by non-RMAN methods, for example, through the media manager's tape retention policy. The media manager should never expire a tape until all RMAN backups on that tape have been removed from the media manager's catalog.

There are two mutually exclusive options for implementing a retention policy: **redundancy** and **recovery window**. If no retention policy is configured by the user, then the `REPORT OBSOLETE` and `DELETE OBSOLETE` commands use a default retention policy of `REDUNDANCY 1`.

To configure a retention policy based on a recovery window, use the following command:

- `CONFIGURE RETENTION POLICY TO RECOVERY WINDOW`

To configure a retention policy based on redundancy, use the following command:

- `CONFIGURE RETENTION POLICY TO REDUNDANCY`

You can also disable the retention policy completely, meaning that RMAN does not consider any backup to be obsolete. To do so, use the following command:

```
CONFIGURE RETENTION POLICY TO NONE;
```

Recovery Window

A recovery window is a period of time that begins with the current time and extends backward in time to the **point of recoverability**. The point of recoverability is the earliest time for a hypothetical point-in-time recovery, that is, the earliest point to which you can recover following a media failure. For example, if you implement a recovery window of one week, then this window of time must extend back exactly seven days from the present so that you can restore a backup and recover it to this point. You implement this retention policy as follows:

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

This command ensures that for each datafile one backup that is older than the point of recoverability must be retained. For example, if the recovery window is 7, then there must always exist one backup of each datafile that satisfies the following condition:

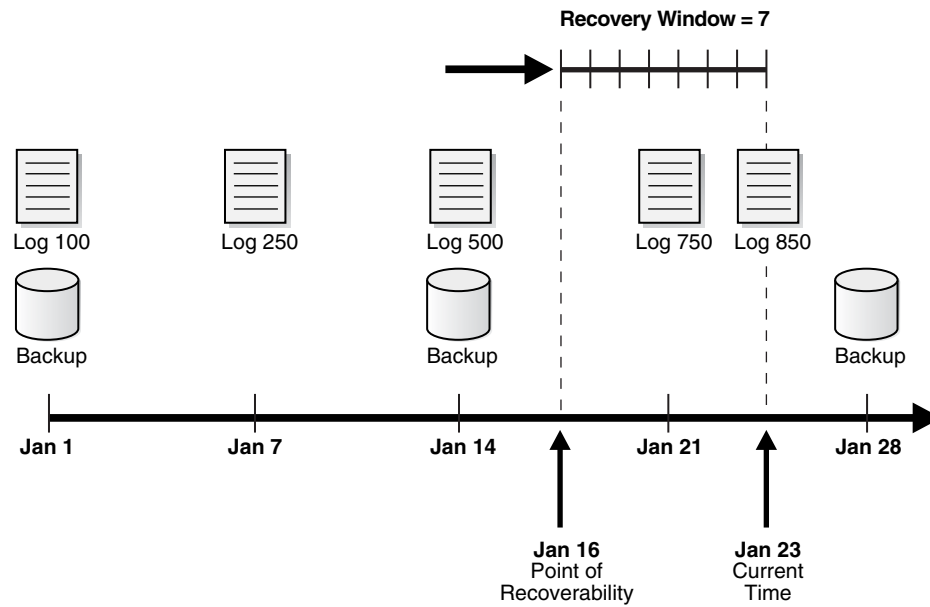
```
SYSDATE - BACKUP CHECKPOINT TIME >= 7
```

All backups older than the most recent backup that satisfied this condition are obsolete.

Assume the following retention policy illustrated in [Figure 2-6](#). The retention policy has the following aspects:

- The recovery window is 7 days.
- Database backups are scheduled every two weeks on these days:
 - January 1
 - January 15
 - January 29
 - February 12
- The database runs in `ARCHIVELOG` mode, and archived logs are saved on disk only as long as needed for the retention policy.

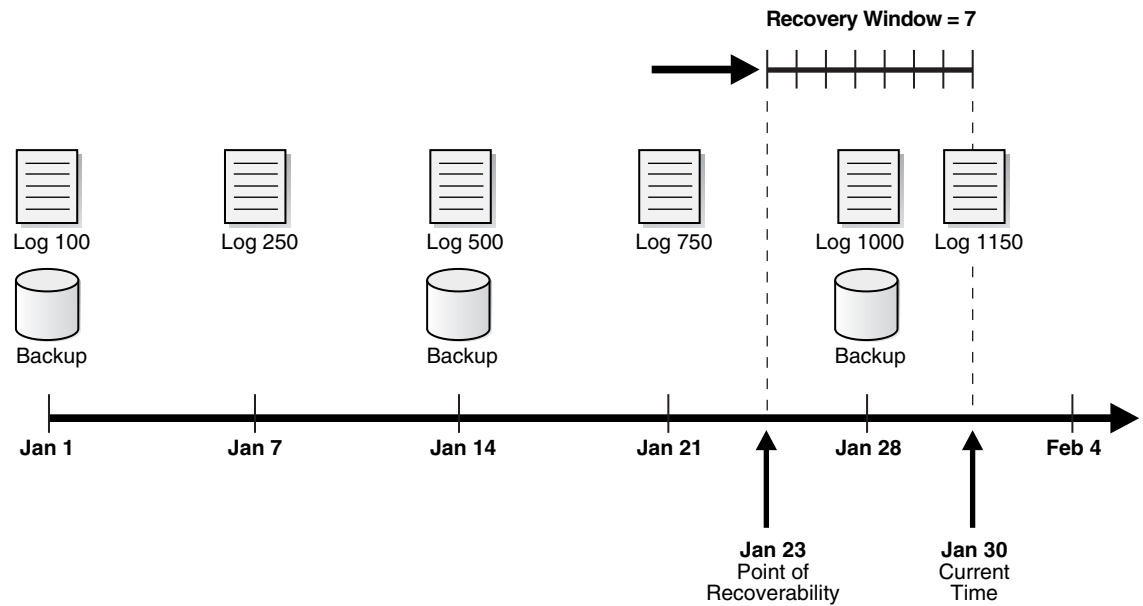
Figure 2–6 Recovery Window, Part 1



As illustrated in [Figure 2–6](#), the current time is January 23 and the point of recoverability is January 16. Hence, the January 14 backup is needed for recovery, and so are the archived logs from log sequence 500 through 850. The logs before 500 and the January 1 backup are obsolete because they are not needed for recovery to a point within the window.

Assume the same scenario a week later, as depicted in [Figure 2–7](#).

Figure 2–7 Recovery Window, Part 2



In this scenario, the current time is January 30 and the point of recoverability is January 23. Note how the January 14 backup is *not* obsolete even though a more recent backup (January 28) exists in the recovery window. This situation occurs because

restoring the January 28 backup does not enable you to recover to the earliest time in the window, January 23. To ensure recoverability to any point within the window, you must save the January 14 backup as well as all archived redo logs from log sequence 500 to 1150.

Backup Redundancy

A redundancy-based retention policy specifies how many backups of each datafile must be retained. For example, you can specify:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 2;
```

Although the recovery window is the best practice for specifying a retention policy, it can complicate disk space usage planning because the number of backups that must be kept by the recovery window is not constant and depends on the backup schedule. Use the `CONFIGURE RETENTION POLICY TO REDUNDANCY n` command to implement a retention policy that maintains a constant number of backups of each datafile. `RECOVERY WINDOW` and `REDUNDANCY`-based retention policies are mutually exclusive.

The default retention policy is `REDUNDANCY = 1`, to maintain compatibility with the behavior of `REPORT OBSOLETE` in earlier RMAN releases. You can also run the following command to disable the retention policy altogether:

```
CONFIGURE RETENTION POLICY TO NONE;
```

If the retention policy is configured to `NONE`, then `REPORT OBSOLETE` and `DELETE OBSOLETE` do not consider any backups to be obsolete.

Batch Deletes of Obsolete Backups

Run the `REPORT OBSOLETE` command to determine which backups are currently obsolete according to the retention policy.

If you configure the retention policy to `NONE`, then RMAN does not consider any backups as obsolete. Consequently, RMAN issues an error when you run `REPORT OBSOLETE` without any other options and the retention policy is set to `NONE`.

A companion command, `DELETE OBSOLETE`, deletes all files which are obsolete according to the retention policy. You should run `DELETE OBSOLETE` periodically to minimize space wasted by storing obsolete backups. For example, you can run `DELETE OBSOLETE` in a weekly script.

You can also specify the `REDUNDANCY` or `RECOVERY WINDOW` options on the `REPORT` or `DELETE` commands, to override the configured retention policy.

The `REPORT OBSOLETE` and `DELETE OBSOLETE` commands work in two steps:

1. For each datafile for which there are full backup, datafile copy, or level 0 incremental backups, RMAN identifies the oldest full or level 0 backup or copy that is not obsolete under the retention policy being tested. Any full backup, level 0 incremental backup, or datafile copy of a datafile older than the one identified in this step is considered obsolete.
2. Any archived logs and level 1 incremental backups that are older than the oldest non-obsolete full backup are then obsolete because there is no full or level 0 backup to which they can be applied.

See Also:

- *Oracle Database Backup and Recovery Basics* to generate reports and delete backups
- *Oracle Database Backup and Recovery Reference* for DELETE syntax
- *Oracle Database Backup and Recovery Reference* for REPORT syntax

Exempting Backups from the Retention Policy

You may want to store a **long-term backup**, potentially offsite, for much longer than the time dictated by the retention policy. For example, you may make a database backup on the first day of every year to satisfy some regulatory requirement, independent of backups taken for your ongoing backup and recovery strategy.

Such long-term backups should be recorded in the RMAN repository, but they must be exempted from the retention policy because RMAN would quickly consider them obsolete, and they would be removed the next time the DELETE OBSOLETE command is used.

You can exempt a backup from the retention policy by using the KEEP option with the BACKUP command when you create the backup, or the KEEP option of the CHANGE command to exempt an existing backup. Note that backups exempted from the retention policy are still fully valid backups, which can be used in restore and recovery operations like any other if RMAN judges them to be the best choice available.

You can change the exempt status of a backup using the CHANGE . . . KEEP and CHANGE . . . NOKEEP commands. The NOKEEP option (default) indicates that the backup is not immune from the configured retention policy.

You can specify the LOGS option to save archived logs for a possible incomplete recovery of the long-term backup. When LOGS is specified, all logs more recent than the backup are kept as long as the backup is kept. In other words, KEEP UNTIL TIME . . . LOGS means that RMAN will keep all logs required to recover the backup as long as the backup is kept. If you specify NOLOGS, then RMAN does not keep the logs required to recover the backup. Note that if you use KEEP UNTIL TIME . . . with an inconsistent backup, you must use the LOGS option, or that backup will become unusable when the logs required to recover it are deleted as obsolete.

You can specify an end date using the UNTIL clause, or either specify that the backup should be kept FOREVER. If you specify UNTIL, then RMAN will not mark the backup as obsolete until after the UNTIL date has passed. Note that it is an error to specify KEEP FOREVER with the LOGS option, as this would require keeping all redo logs forever.

The following commands are examples of long-term backups:

```
# Creates a backup and exempts it from retention policy until last day of 2007
BACKUP DATABASE KEEP UNTIL TIME "TO_DATE('31-DEC-2007' 'dd-mon-yyyy')" NOLOGS;
```

```
# Specifies that backupset 2 is no longer exempt from the retention policy
CHANGE BACKUPSET 2 NOKEEP;
```

```
# Creates a backup that is indefinitely exempt from the retention policy
BACKUP TABLESPACE users KEEP FOREVER NOLOGS;
```

See Also: *Oracle Database Backup and Recovery Reference* for CHANGE syntax

Relationship Between Retention Policy and Flash Recovery Area Rules

The RMAN status `OBSOLETE` is always determined in reference to a retention policy. For example, if a database backup is `OBSOLETE` in the RMAN repository, it is because it is either not needed for recovery to a point within the recovery window, or it is redundant.

If you configure a flash recovery area, then the database uses an internal algorithm to select files in the flash recovery area that are no longer needed to meet the configured retention policy. These backups have status `OBSOLETE`, and are eligible for deletion to satisfy the disk quota rules.

There is one important difference between the flash recovery area criteria for `OBSOLETE` status and the disk quota rules for deletion eligibility. Assume that archived logs 1000 through 2000, which are on disk, are needed for the currently enabled recovery window and so are not obsolete. If you back up these logs to tape, then the retention policy still considers the disk logs as required, that is, not obsolete. Nevertheless, the flash recovery area disk quota algorithm considers the logs on disk as eligible for deletion because they have already been backed up to tape. The logs on disk do not have `OBSOLETE` status in the repository, but they are eligible for deletion by the flash recovery area. Note, though, that the retention policy is never violated when determining which files to delete from the flash recovery area to satisfy the disk quota rules.

Backup Optimization

If you enable backup optimization, then the `BACKUP` command skips backing up files when the identical file has already been backed up to the specified device type.

Rules for Identifying Identical Files for Backup Optimization

Table 2–2 describes criteria that RMAN uses to determine whether a file is identical to a file that it already backed up.

Table 2–2 Criteria to Determine an Identical File

Type of File	Criteria to Determine an Identical File
Datafile	The datafile must have the same DBID, checkpoint SCN, creation SCN, and <code>RESETLOGS</code> SCN and time as a datafile already in a backup. The datafile must be offline-normal, read-only, or closed normally.
Archived redo log	Same thread, sequence number, and <code>RESETLOGS</code> SCN and time.
Backup set	Same backup set recid and stamp.

If RMAN determines that a file is identical and it has already been backed up, then it is a candidate to be skipped. However, RMAN must do further checking to determine whether to skip the file, because both the retention policy and the backup duplexing feature are factors in the algorithm that determines whether RMAN has sufficient backups on the specified device type.

Backup Optimization Algorithm

Table 2–3 describes the algorithm that backup optimization uses when determining whether to skip the backup of an identical file.

Table 2–3 Backup Optimization Algorithm

For an Identical ...	Backup Optimization Algorithm
Datafile	<p>With a recovery window-based retention policy:</p> <p>For backups to tape, RMAN takes another backup of a file, even if a backup of an identical file exists, if the most recent backup is older than the configured recovery window. This is done to allow media to be recycled after the media expires.</p> <p>For backups to disk, RMAN skips taking the backup if an identical file is available from a backup on disk, even if that backup is older than the beginning of the recovery window. The retention policy causes RMAN to retain the old backup for as long as it is needed.</p> <p>With a redundancy-based retention policy:</p> <p>RMAN sets $r=1$ by default and searches for values of n in this order of precedence (that is, values higher on the list override values lower on the list):</p> <ol style="list-style-type: none"> 1. If CONFIGURE RETENTION POLICY TO REDUNDANCY r is enabled, then RMAN only skips datafiles when $n=r+1$ backups exist. 2. BACKUP . . . COPIES n 3. SET BACKUP COPIES n 4. CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE . . . TO n <p>RMAN skips backup only if at least n backups of an identical file exist on the specified device. If RMAN does not skip the backup, then it makes the backup exactly as specified.</p>
Archived log	<p>By default, $n=1$. RMAN searches for values of n in this order of precedence (that is, values higher on the list override values lower on the list):</p> <ol style="list-style-type: none"> 1. BACKUP . . . COPIES n 2. SET BACKUP COPIES n 3. CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE . . . TO n <p>RMAN skips backup only if at least n backups of an identical file exist on the specified device. If RMAN does not skip the backup, then it makes the backup exactly as specified.</p>
Backup set	<p>By default, $n=1$. RMAN searches for other values of n in this order of precedence (that is, values higher on the list override values lower on the list):</p> <ol style="list-style-type: none"> 1. BACKUP . . . COPIES n 2. SET BACKUP COPIES n <p>RMAN skips backup only if at least n backups of an identical file exist on the specified device. If RMAN does not skip the backup, then it makes the backup exactly as specified.</p>

For example, assume that at 9 a.m. you back up three copies of all existing archived logs to tape:

```
BACKUP DEVICE TYPE sbt COPIES 3 ARCHIVELOG ALL;
```

Later, you enable the following configuration setting in preparation for another backup:

```
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE sbt TO 4;
CONFIGURE BACKUP OPTIMIZATION ON;
```

Then, you run the following archived log backup at noon:

```
BACKUP DEVICE TYPE sbt COPIES 2 ARCHIVELOG ALL;
```

In this case, the BACKUP . . . COPIES setting overrides the CONFIGURE . . . COPIES setting, so RMAN sets $n=2$. RMAN skips the backup of a log only if at least two copies of the log exist on the *sbt* device. Because three copies of each log exist on *sbt* of all the logs generated before 9 a.m., RMAN skips the backups of these logs. However,

RMAN backs up two copies of all logs generated after 9 a.m. because these logs have not yet been backed up to tape.

At this point, three copies of the logs created before 9 a.m. exist on tape, and two copies of the logs created after 9 a.m. exist on tape. Assume that you run the following backup at 3 p.m.:

```
RUN
{
  SET BACKUP COPIES 3;
  BACKUP DEVICE TYPE sbt ARCHIVELOG ALL;
}
```

In this case, RMAN sets $n=3$ and so will not back up the logs created before 9 a.m. because three copies already exist on tape. However, only two copies of the logs created after 9 a.m. exist on tape, so RMAN does not optimize backups of these logs. Hence, RMAN backs up three copies of the logs created after 9 a.m.

Requirements for Backup Optimization

Backup optimization is used when the following conditions are true:

- The `CONFIGURE BACKUP OPTIMIZATION ON` command has been run to enable backup optimization.
- You run `BACKUP DATABASE`, `BACKUP ARCHIVELOG` with `ALL` or `LIKE` options, or `BACKUP BACKUPSET ALL`.
- Only one type of channel is allocated, that is, you do not mix channels of type `DISK` and `sbt` in the same backup command.

For example, assume that you run these commands:

```
BACKUP DEVICE TYPE sbt DATABASE PLUS ARCHIVELOG;
BACKUP DEVICE TYPE sbt BACKUPSET ALL;
```

If none of these files has changed since the last backup, then RMAN does not back up the files again, nor signal an error if it skips all files specified in the command.

Overriding and Disabling Backup Optimization

To override backup optimization and back up all files whether or not they have changed, specify the `FORCE` option on the `BACKUP` command. For example:

```
BACKUP DEVICE TYPE sbt BACKUPSET ALL FORCE;
```

To disable backup optimization on a persistent basis, use the following command:

```
RMAN> CONFIGURE BACKUP OPTIMIZATION OFF;
```

Effect of Retention Policies on Backup Optimization for SBT Backups

Backup optimization is not always applied when backing up to SBT devices. The exceptions to normal backup optimization behavior for recovery window-based and redundancy-based retention policies are described in the following sections.

Note: Use caution when enabling backup optimization if you use a media manager with its own internal expiration policy. Run `CROSSCHECK` periodically to synchronize the RMAN repository with the media manager. Otherwise, RMAN may skip backups due to optimization without recognizing that the media manager has discarded backups stored on tape.

Backup Optimization for SBT Backups with Recovery Window Retention Policy

If backup optimization is enabled, and if a recovery window retention policy is in effect, then when performing SBT backups RMAN always backs up datafiles whose most recent backup is older than the recovery window. For example, assume that:

- Today is February 21.
- The recovery window is 7 days.
- The most recent backup of tablespace `tools` to tape is January 3.
- Tablespace `tools` is read-only.

On February 21, when you issue a command to back up tablespace `tools` to tape, RMAN backs it up even though it did not change after the January 3 backup (because it is read-only). RMAN makes the backup because no backup of the tablespace exists within the 7-day recovery window.

This behavior allows the media manager to expire old tapes. Otherwise, the media manager would be forced to keep the January 3 backup of tablespace `tools` indefinitely. By making a more recent backup of tablespace `tools` on February 21, RMAN allows the media manager to expire the tape containing the obsolete January 3 backup.

Backup Optimization for SBT Backups With Redundancy Retention Policy

Assume that you configure a retention policy for redundancy. In this case, RMAN only skips backups of offline or read-only datafiles to SBT when there are $r + 1$ backups of the files, where r is set in `CONFIGURE RETENTION POLICY TO REDUNDANCY r` .

For example, assume that you enable backup optimization and set the following retention policy:

```
CONFIGURE DEFAULT DEVICE TYPE TO sbt;
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE RETENTION POLICY TO REDUNDANCY 2;
```

With these settings, RMAN only skips backups when three identical files are already backed up. Also assume that you have never backed up the `users` tablespace, which is read/write, and that you perform the actions described in [Table 2-4](#) over the course of the week.

Table 2-4 Effect of Redundancy Setting on Backup Optimization

Day	Action	Result	Redundant Backup
Monday	Take tablespace <code>users</code> offline normal.		
Tuesday	<code>BACKUP DATABASE</code>	The <code>users</code> tablespace is backed up.	
Wednesday	<code>BACKUP DATABASE</code>	The <code>users</code> tablespace is backed up.	

Table 2–4 (Cont.) Effect of Redundancy Setting on Backup Optimization

Day	Action	Result	Redundant Backup
Thursday	BACKUP DATABASE	The users tablespace is backed up.	Tuesday backup
Friday	BACKUP DATABASE	The users tablespace is <i>not</i> backed up.	Tuesday backup
Saturday	BACKUP DATABASE	The users tablespace is <i>not</i> backed up.	Tuesday backup
Sunday	DELETE OBSOLETE	The Tuesday backup is deleted.	
Monday	BACKUP DATABASE	The users tablespace is backed up.	Wednesday backup

The backups on Tuesday, Wednesday, and Thursday back up the offline `users` tablespace to satisfy the condition that three backups must exist (one more than redundancy setting). The Friday and Saturday backups do not back up the `users` tablespace because of backup optimization. Note that the Tuesday backup of `users` is obsolete beginning on Thursday.

On Sunday, you delete all obsolete backups, which removes the Tuesday backup of `users`. The Tuesday backup is obsolete because of the retention policy setting. The whole database backup on Monday then backs up the `users` tablespace to satisfy the condition that three backups must exist (one more than redundancy setting). In this way, you can recycle your tapes over time.

See Also: ["Backing Up Files Using Backup Optimization"](#) on page 6-11, and ["Configuring Backup Optimization"](#) on page 5-16

Restartable Backups

Using the **restartable backups** feature, RMAN can back up only those files that have not been backed up since a specified date. Use this feature after a backup fails to back up the parts of the database missed by the failed backup.

The unit of restartability is a backup set. If the backup generates multiple backup sets, then the backups that completed successfully do not have to be rerun. If the entire database is written into one backup set, and if the backup fails halfway through, then the entire backup has to be restarted.

To take better advantage of restartable backups, you can use set the `MAXSETSIZE` parameter of the `BACKUP` command. If, for instance, you set `MAXSETSIZE` to 10MB for a given backup command, a new backup set is produced for each 10MB of backup output. If the backup fails after some backup sets have been produced and must be restarted, the data backed up in those backup sets will not have to be backed up again. (Note that `MAXSETSIZE` must be large enough that any file can be accommodated in a single backup piece, because files cannot span backup pieces.)

For example, if the largest datafile is less than 10 MB, then you can back up the database daily as follows:

```
BACKUP DATABASE MAXSETSIZE = 10M;
```

Then, after a failure you can back up all files in the database that were not backed up in the last 24 hours by issuing:

```
BACKUP DATABASE NOT BACKED UP SINCE TIME 'SYSDATE-1';
```

If the `SINCE TIME` is later than the completion time, then RMAN backs up the file. If you use "BACKUP DATABASE NOT BACKED UP" without the `SINCE TIME` parameter, then RMAN only backs up files that have never been backed up.

When determining whether a file has been backed up, RMAN compares the `SINCE TIME` date with the completion time of the most recent backup of the file. The completion time for a backup piece is the completion time of the entire backup set, not an individual backup piece; in other words, all files in the same backup set have the same completion time.

See Also: ["Restarting a Backup After It Partially Completes"](#) on page 6-11 and *Oracle Database Backup and Recovery Reference* for BACKUP syntax

Managing Backup Windows and Performance: BACKUP... DURATION

A **backup window** is a period of time during which a backup activity must complete. For example, you may want to restrict your database backup activities to a window of time when user activity on your system is low, such as between 2:00 AM and 6:00 AM.

The BACKUP command supports a `DURATION` argument which lets you specify how long a given backup job is allowed to run. You could, for example, run the following command at 2:00AM:

```
BACKUP DURATION 4:00 TABLESPACE users;
```

RMAN backs up the specified data at the maximum possible speed. If the backup is not complete in four hours, the backup is interrupted. Any completed backupsets are retained and can be used in restore operations, even if the entire backup is not complete. Any incomplete backupsets are discarded.

Controlling RMAN Behavior when Backup Window Ends with PARTIAL

By default, when a BACKUP... DURATION command runs out of time before the backup completes, RMAN reports an error. (The effect of this is that if the command is running in a RUN block, the RUN block terminates.) You can control this behavior by adding the `PARTIAL` option to the BACKUP... DURATION command, as in this example:

```
BACKUP DURATION 4:00 PARTIAL TABLESPACE users FILESPERSET 1;
```

When `PARTIAL` is used, no error is reported when a backup command is interrupted due to the end of the backup window. Instead, a message showing which files could not be backed will be displayed. If the BACKUP command is part of a RUN block, then the remaining commands in the RUN block will continue to execute.

When using `DURATION` the least recently backed up files are backed up first. Thus, if you retry a job that was interrupted when the available duration expired, each successive attempt covers more of the files needing backup.

Note also the use of `FILESPERSET 1` in this example. With this option, RMAN puts each file into its own backupset. This way, when a backup is interrupted at the end of the backup window, only the backup of the file currently being backed up is lost. All backup sets completed during the window are saved, minimizing the lost work due to the end of the backup window.

Managing Backup Performance with MINIMIZE TIME and MINIMIZE LOAD

When using `DURATION` you can run the backup with the maximum possible performance, or run as slowly as possible while still finishing within the allotted time, to minimize the performance impact of backup tasks. To maximize performance, use the `MINIMIZE TIME` option with `DURATION`, as shown in this example:

```
BACKUP DURATION 4:00 PARTIAL MINIMIZE TIME DATABASE FILESPERSET 1;
```

To extend the backup to use the full time available, use the `MINIMIZE LOAD` option, as in this example:

```
BACKUP DURATION 4:00 PARTIAL MINIMIZE LOAD DATABASE FILESPERSET 1;
```

RMAN monitors the progress of the running backup, and periodically estimates how long the backup will take to complete at its present rate. If RMAN estimates that the backup will finish before the end of the backup window, it slows down the rate of backup so that the full available duration will be used. This reduces the overhead on the database associated with the backup.

Note: Note these issues when using `DURATION` and `MINIMIZE LOAD` with a tape backup:

- Efficient backup to tape requires tape streaming. If you use `MINIMIZE LOAD`, RMAN may reduce the rate of backup to the point where tape streaming is not optimal.
- RMAN will hold the tape resource for the entire duration of the backup window. This prevents the use of the tape resource for any other purpose during the backup window.

Because of these concerns, it is not recommended that `MINIMIZE LOAD` be used with tape. See ["Factors Affecting Backup Speed to Tape"](#) on page 11-4 for more details on efficient tape handling.

How RMAN Responds to Backup Errors

RMAN detects and responds to two primary types of backup errors: I/O errors and corrupt blocks. This section contains the following topics:

- [Handling I/O Errors in RMAN Backup: NOT BACKED UP SINCE](#)
- [Handling Corrupt Datafile Blocks in RMAN Backup: MAXCORRUPT](#)

Handling I/O Errors in RMAN Backup: NOT BACKED UP SINCE

Any I/O errors that RMAN encounters when reading files or writing to the backup pieces or image copies cause RMAN to terminate the backup job in progress. For example, if RMAN tries to back up a datafile but the datafile is not on disk, then RMAN terminates the backup.

If multiple channels are being used, or redundant copies of backups are being created, RMAN may be able to continue the backup without user intervention.

If an error occurs during `BACKUP AS BACKUPSET`, then RMAN needs to rewrite the backup sets that it was writing at the time of the error. However, for any backup job that creates multiple backup sets, if some backup sets were complete during the error then RMAN retains those backups.

The `NOT BACKED UP SINCE` option of the `BACKUP` command backs up only files that are not already backed up since the time specified. You can use this clause after an interrupted backup to resume the backup.

Handling Corrupt Datafile Blocks in RMAN Backup: `MAXCORRUPT`

When RMAN encounters a corrupt datafile block during a backup, the behavior depends upon whether RMAN has encountered this corrupt block during a previous backup. If the block is already identified as corrupt, then it is included in the backup. If the block is not previously identified as corrupt, then RMAN's default behavior is to stop the backup.

You can override this behavior using the `SET MAXCORRUPT` command with `BACKUP` in a `RUN` block. Setting `MAXCORRUPT` allows a specified number of previously undetected block corruptions in datafiles during the execution of an `RMAN BACKUP` command. If RMAN detects more than this number of new corrupt blocks while taking the backup, then the backup job aborts, and no backup is created.

As RMAN finds corrupt blocks during the backup process, it writes the corrupt blocks to the backup with a special header indicating that the block has media corruption. If the backup completes without exceeding the specified `MAXCORRUPT` limit, then the database records the address of the corrupt blocks and the type of corruption found (logical or physical) in the control file. You can access these records through the `V$DATABASE_BLOCK_CORRUPTION` view.

Note: If the backup job aborts because more than `MAXCORRUPT` corrupt blocks are found, the `V$DATABASE_BLOCK_CORRUPTION` view is not populated, because the information used to populate the view is only available if a backup is successfully created. In such a situation, you can run `BACKUP VALIDATE` on the datafiles to be backed up, to populate `V$DATABASE_BLOCK_CORRUPTION` and use block media recovery to repair the corrupt blocks. See *Oracle Database Backup and Recovery Basics* for details on using `BACKUP... VALIDATE`, and "[Performing Block Media Recovery with RMAN](#)" on page 7-12 for more details on block media recovery.

See Also:

- "[Tests and Integrity Checks for Backups](#)" on page 2-42 for more information about fractured and corrupt blocks
- "[Restartable Backups](#)" on page 2-39 for more information about the `NOT BACKED UP SINCE` clause
- *Oracle Database Reference* for a description of `V$DATABASE_BLOCK_CORRUPTION`
- *Oracle Database Backup and Recovery Reference* for `SET MAXCORRUPT` syntax

Tests and Integrity Checks for Backups

The database prevents operations that result in unusable backup files or corrupt restored datafiles. The database server automatically does the following:

- Blocks access to datafiles while they are being restored or recovered
- Allows only one restore operation for each datafile at a time

- Ensures that incremental backups are applied in the correct order
- Stores information in backup files to allow detection of corruption

You can use the `BACKUP VALIDATE` and `RESTORE VALIDATE` commands to test backup and restore operations without creating output files. In this way, you can check your datafiles for possible problems. If you run RMAN with the following configuration, then it detects all types of corruption that are possible to detect:

- In the initialization parameter file, set `DB_BLOCK_CHECKSUM=TRUE`
- In the RMAN `BACKUP` and `RESTORE` commands, do *not* specify the `MAXCORRUPT` option, do *not* specify the `NOCHECKSUM` option, but *do* specify the `CHECK LOGICAL` option

See *Oracle Database Backup and Recovery Basics* for more details on `BACKUP VALIDATE` and `RESTORE VALIDATE`.

Detecting Physical Block Corruption With RMAN BACKUP... VALIDATE

RMAN depends upon database server sessions to perform backups, and the database server can detect many types of physically corrupt blocks during the backup process. Each new corrupt block not previously encountered in a backup is recorded in the control file and in the `alert.log`. By default, error checking for physical corruption is enabled.

At the end of a backup, RMAN stores the corruption information in the recovery catalog and control file. Access this data using the `V$DATABASE_BLOCK_CORRUPTION` view.

If the server session encounters a datafile block during a backup that has already been identified as corrupt by the database, then the server session copies the corrupt block into the backup and the corrupt block is recorded the control file as either a logical or media corruption. RMAN copies the block in case the user wants to try to salvage the contents of the block.

If RMAN encounters a datafile block that has media corruption but that has not already been identified as corrupt by the database, then it writes the block to the backup with a reformatted header indicating that the block has media corruption.

Detection of Logical Block Corruption

Besides testing for media corruption, the database can also test data and index blocks for logical corruption, such as corruption of a row piece or index entry. If RMAN finds logical corruption, then it logs the block in the `alert.log`. If `CHECK LOGICAL` was used, the block is also logged in the server session trace file. By default, error checking for logical corruption is disabled.

For `BACKUP` commands the `MAXCORRUPT` parameter sets the total number of physical and logical corruptions permitted in a file. If the sum of physical and logical corruptions for a file is less than its `MAXCORRUPT` setting, the RMAN command completes successfully. If `MAXCORRUPT` is exceeded, the command terminates and RMAN does not read the rest of the file. `V$DATABASE_BLOCK_CORRUPTION` is populated with corrupt block ranges if the command succeeds. Otherwise, you must set `MAXCORRUPT` higher and re-run the backup to find out the corrupt block ranges.

See Also: *Oracle Database Backup and Recovery Reference* for `BACKUP . . . MAXCORRUPT` syntax

Detection of Fractured Blocks During Open Backups

One danger in making online backups is the possibility of inconsistent data within a block. For example, assume that you are backing up block 100 in datafile `users.dbf`. Also, assume that the copy utility reads the entire block while DBWR is in the middle of updating the block. In this case, the copy utility may read the old data in the top half of the block and the new data in the bottom top half of the block. The result is called a **fractured block**, meaning that the data contained in this block is not consistent. at a given SCN.

When performing backups of an open tablespace *without* using RMAN, you must put tablespaces in **backup mode** to prevent the creation of fractured blocks in your backup. When not in backup mode, the database records only changed bytes in the redo stream. When a tablespace is in backup mode, each time a block is changed the database writes the before-image of the entire block to the redo stream before modifying the block. Then, the database also records the changes to the block in the redo log. During user-managed recovery using SQL*Plus, the database applies both the captured block images and the recorded block changes from the redo logs. Applying the block images repairs any possible fractured blocks in the backup being restored and recovered.

RMAN does not require that you put datafiles into backup mode. During an RMAN backup, a database server session reads each block of the datafile and checks whether each block is fractured by comparing the block header and footer. If a block is fractured, the session re-reads the block. If the same fracture is found, then the block is considered permanently corrupt. If `MAXCORRUPT` is exceeded, the backup stops.

Backup Validation with RMAN

You can run the `BACKUP . . . VALIDATE` command to check datafiles for physical and logical corruption, or to confirm that all database files exist in the correct locations. No backup is taken, but all specified files are scanned to verify that they can be backed up. All corruptions are recorded in the `V$DATABASE_BLOCK_CORRUPTION` view.

The following example shows how to check your entire database and archived redo log files for physical and logical corruption:

```
BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
```

You cannot use the `MAXCORRUPT` or `PROXY` parameters with the `VALIDATE` option.

See Also: *Oracle Database Backup and Recovery Reference* for `BACKUP` syntax and "[Validating Backups with RMAN](#)" on page 6-11 for more details on using `BACKUP VALIDATE`.

RMAN Recovery Concepts

This chapter describes the basic concepts involved in using RMAN to restore, recover, and duplicate databases.

This chapter contains these topics:

- [Restoring Files with RMAN](#)
- [Datafile Media Recovery with RMAN](#)
- [Block Media Recovery with RMAN](#)
- [Database Duplication with RMAN](#)
- [Physical Standby Database Creation with RMAN](#)

Restoring Files with RMAN

Use the RMAN `RESTORE` command to restore the following types of files from disk or other media:

- Database (all datafiles)
- Tablespaces
- Control files
- Archived redo logs
- Server parameter files

Because a backup set is in a proprietary format, you cannot simply copy it as you would a backup database file created with an operating system utility; you must use the RMAN `RESTORE` command to extract its contents. In contrast, the database can use image copies created by the RMAN `BACKUP AS COPY` command without additional processing.

See Also: *Oracle Database Backup and Recovery Reference* for `RESTORE` syntax and prerequisites

Mechanics of Datafile Restore Operations

When you issue a `RESTORE` command, RMAN directs a server session to restore a suitable backup to either:

- The default location, overwriting the files with the same name currently there
- A new location, which you can specify with the `SET NEWNAME` command

To restore a datafile, either mount the database, or keep it open and take the datafile to be restored offline. When RMAN performs a restore, it creates the restored files as datafile image copies and records them in the repository.

The following table describes the behavior of the `RESTORE`, `SET NEWNAME`, and `SWITCH` commands.

Run SET NEWNAME?	RESTORE Behavior	Run SWITCH?
No	RMAN restores the files to their current path names.	N/A
Yes	RMAN restores the files to the path names specified by <code>SET NEWNAME</code> and creates repository records for each datafile copy after the restore.	If yes, then RMAN replaces the current datafile names in the control file to the names of the restored files and current datafile names (if exists) are stored as datafile copies. If no, then RMAN doesn't update current datafile names and the restored file is retained as datafile copies.

For example, if you restore datafile `*/oradata/trgt/tools01.dbf` to its default location, then RMAN restores the file `*/oradata/trgt/tools01.dbf` and overwrites any file that it finds with the same filename. If you run a `SET NEWNAME` command before you restore a file, then RMAN creates a datafile copy with the name that you specify. For example, assume that you run the following commands:

```
RUN {
  SET NEWNAME FOR DATAFILE '*/oradata/trgt/tools01.dbf' TO '/tmp/tools01.dbf';
  RESTORE DATAFILE '*/oradata/trgt/tools01.dbf';
}
```

In this case, RMAN creates a datafile copy of `*/oradata/trgt/tools01.dbf` named `/tmp/tools01.dbf` and records it in the repository. To update the control file to use the datafile copy at `*/oradata/trgt/tools01.dbf` to `/tmp/tools01.dbf` as the datafile, use the `SWITCH` command as shown in the following example:

```
SWITCH DATAFILE '/tmp/tools01.dbf' TO DATAFILECOPY '*/oradata/trgt/tools01.dbf';
```

The `SWITCH` command is the RMAN equivalent of the SQL statement `ALTER DATABASE RENAME FILE`.

See Also: *Oracle Database Backup and Recovery Reference* for `SET NEWNAME` syntax, and *Oracle Database Backup and Recovery Reference* for `SWITCH` syntax

How RMAN Selects Backups for Use in Restore Operations

RMAN uses the record of available backups in the RMAN repository to select the best available backups for use in the restore operation. The most recent backup available, or the most recent backup satisfying any `UNTIL` clause specified in the `RESTORE` command, is always the preferred choice. If two backups are from the same point in time, RMAN prefers to use image copies over backup sets because RMAN can restore more quickly from image copies than from backup sets (especially those stored on tape).

All specifications of the `RESTORE` command must be satisfied before RMAN restores a backup. Unless limited by the `DEVICE TYPE` clause, the `RESTORE` command searches for backups on all device types of configured channels.

If no available backup in the repository satisfies all the specified criteria, then RMAN returns an error during the compilation phase of the restore job.

If you use only manually allocated channels, a backup job may fail if there is no usable backup on the media for which you allocated channels. Configuring automatic channels makes it more likely that RMAN can find and restore a backup that satisfies the specified criteria.

See Also: ["Configuring Channels"](#) on page 5-9

Restore Failover

During a `RESTORE` operation, if a backup piece, image copy or proxy copy is inaccessible (for instance, deleted from the device) or a block in the backup is corrupted, then RMAN automatically looks for another usable copy of this backup piece or image copy, on the same device or another device, based on the information in the RMAN repository. If no usable copies are available, then RMAN searches for prior backups. RMAN searches all prior backups for the most recent available backup usable in the current operation until it has exhausted all possibilities. If all of the backups are unusable or no backups exist, then RMAN will try to re-create the datafile.

Restore failover is also used when there are errors restoring archivelogs during `RECOVER`, `BLOCKRECOVER`, and `FLASHBACK DATABASE` commands.

When RMAN performs restore failover to another backup of the same file, you will see a message similar to this one in the output of RMAN:

```
failover to piece handle=/u01/backup/db_1 tag=BACKUP_031009
```

Also, details about block corruptions will be printed in the alert log and trace files.

When restore failover cannot locate another copy of the same backup and searches for a prior backup, the message generated is similar to this example:

```
ORA-19624: operation failed, retry possible
ORA-19505: failed to identify file "/u01/backup/db_1"
ORA-27037: unable to obtain file status
SVR4 Error: 2: No such file or directory
Additional information: 3
failover to previous backup
```

Restore Optimization

RMAN uses restore optimization to avoid restoring datafiles from backup when possible. If a datafile is already present in the correct location and its header contains the expected information, then RMAN does not restore the datafile from backup.

Note: Restore optimization only checks the datafile header. It does not scan the datafile body for corrupted blocks.

You can use the `FORCE` option of the `RESTORE` command to override this behavior and restore the requested files unconditionally.

Restore optimization is particularly useful in cases where an operation that restores several datafiles is interrupted. For example, assume that a full database restore encounters a power failure after all except one of the datafiles has been restored. If you run the same RESTORE operation again, then RMAN only restores the single datafile that was not restored during the previous attempt.

Restore optimization is also used when duplicating a database. If a datafile at the duplicate is in the correct place with the correct header contents, then the datafile is not duplicated. Unlike RESTORE, however, DUPLICATE does not support a FORCE option. To force RMAN to duplicate a datafile that is skipped due to restore optimization, delete the datafile from the duplicate before running the DUPLICATE command.

See Also: *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide* for description of RESTORE behavior in a RAC configuration

Datafile Media Recovery with RMAN

The concept of **datafile media recovery** is the application of online or archived redo logs or incremental backups to a restored datafile in order to update it to the current time or some other specified time. Use the RMAN RECOVER command to perform media recovery and apply logs or incremental backups automatically.

RMAN Media Recovery: Basic Steps

If possible, make the recovery catalog available to perform the media recovery. If it is not available, or if you do not maintain a recovery catalog, then RMAN uses metadata from the target database control file. If both the control file and recovery catalog are lost, then you can still recover the database—assuming that you have backups of the datafiles and at least one autobackup of the control file.

The generic steps for media recovery using RMAN are as follows:

1. Place the database in the appropriate state: mounted or open. For example, mount the database when performing whole database recovery, or open the database when performing online tablespace recovery.
2. To perform incomplete recovery, use the SET UNTIL command to specify the time, SCN, restore point, or log sequence number at which recovery terminates. Alternatively, specify the UNTIL clause on the RESTORE and RECOVER commands.
3. Restore the necessary files with the RESTORE command.
4. Recover the datafiles with the RECOVER command.
5. Place the database in its normal state. For example, open the database, or bring the recovered tablespaces online.

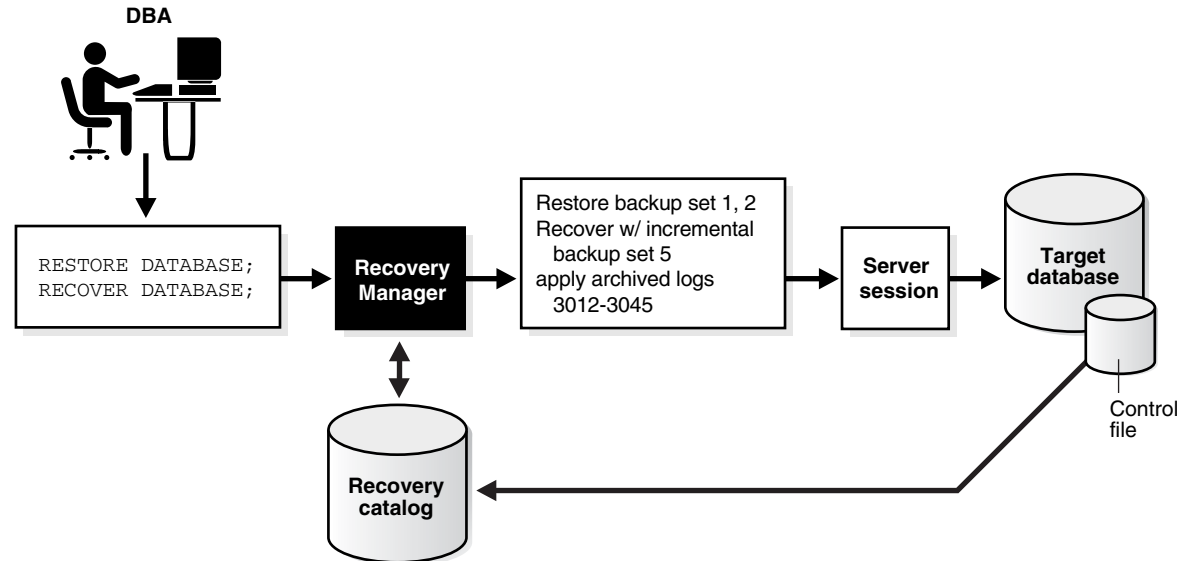
Figure 3–1 illustrates an example of RMAN media recovery using data from the RMAN repository. The repository in this example is stored in a recovery catalog. The contents of the recovery catalog have been synchronized with the record of backups in the control file.

The DBA enters the following commands in the RMAN client:

```
RESTORE DATABASE;  
RECOVER DATABASE;
```


RMAN then queries the repository, which in this example is a recovery catalog. RMAN then decides which backup sets to restore, and which incremental backups and archived logs to use for recovery. A server session on the target database instance performs the actual work of restore and recovery.

Figure 3–1 Performing RMAN Media Recovery



See Also:

- [Chapter 7, "Advanced RMAN Recovery Techniques"](#) for detailed restore and recovery procedures
- ["Control File and Server Parameter File Autobackups"](#) on page 2-28
- *Oracle Database Backup and Recovery Reference* for RESTORE syntax
- *Oracle Database Backup and Recovery Reference* for RECOVER syntax

Mechanics of Recovery: Incremental Backups and Redo Logs

If RMAN has a choice between applying an incremental backup or applying redo to the restored datafiles to meet a recovery objective, then it always chooses an incremental backup. If overlapping levels of incremental backup are available, then RMAN automatically chooses the one covering the longest period of time.

RMAN does not need to apply incremental backups to a restored level 0 incremental backup: it can also apply archived logs. RMAN restores the datafiles that it needs from available backups, applies incremental backups to the datafiles if they are available, and then applies archived logs.

How RMAN Searches for Archived Redo Logs During Recovery

If RMAN cannot find an incremental backup, then it looks in the repository for the names of archived redo logs to use for recovery. The database records an archived log in the control file whenever one of the following occurs:

- The archiver process archives a redo log

- RMAN restores an archived log
- The RMAN `BACKUP AS COPY` command copies a log
- The RMAN `CATALOG` command catalogs a user-managed backup of an archived log

If you use a recovery catalog, then RMAN propagates archived log data into the recovery catalog during resynchronization, classifying archived logs as image copies.

You can view the log information using the `LIST ARCHIVELOG` command or the `V$ARCHIVED_LOG` control file view.

During recovery, RMAN looks for the needed logs using the filenames specified in the `V$ARCHIVED_LOG` view. If the logs were created in multiple destinations or were generated by the `BACKUP AS COPY`, `CATALOG`, or `RESTORE` commands, then multiple, identical copies of each log sequence number exist on disk. All copies of a log sequence number listed as `AVAILABLE` are candidates for use in recovery, regardless of how or where they are stored.

Logs that have been deleted or uncataloged through RMAN are not considered available for recovery. For example, assume that the database archives log 100 to directories `/dest1` and `/dest2`. The RMAN repository indicates that `/dest1/log100.arc` and `/dest2/log100.arc` exist. If you delete `/dest1/log100.arc` with the `DELETE` command, then the repository indicates that only `/dest2/log100.arc` is available for recovery.

If the RMAN repository indicates that no copies of a needed log sequence number exist on disk, then RMAN looks in backups and restores archived redo logs as needed to perform the media recovery. By default, RMAN restores the archived redo logs to the flash recovery area, if one of the archive log destinations is set to `USE_DB_RECOVERY_FILE_DEST`. Otherwise, it restores the logs to the first local archiving destination specified in the initialization parameter file. You can run the `SET ARCHIVELOG DESTINATION` command to specify a different restore location. If you specify the `DELETE ARCHIVELOG` option on `RECOVER`, then RMAN deletes the archived logs after restoring and applying them. If you also specify `MAXSIZE integer` on the `RECOVER` command, then RMAN restores archived logs until the disk space allowed by `MAXSIZE` is consumed, then applies redo from the logs and deletes the restored logs to free space, until there is room enough to restore another archived log. RMAN continues restoring, applying and deleting logs, within the `MAXSIZE` limit, until recovery is complete.

RMAN Behavior When the Repository Is Not Synchronized

If an archived log is deleted from disk but is still listed in the RMAN repository, then RMAN does not perform automatic failover during recovery. For example, if the repository indicates that `/dest1/log100.arc` is on disk when in fact this log was deleted using an operating system command, and if RMAN attempts to apply this log file during recovery, then recovery terminates with an error. RMAN does not automatically attempt to apply other copies of log 100 that are listed as available in the repository.

This situation can sometimes occur when you delete an archived log with an operating system utility and then fail to run `CROSSCHECK` to synchronize the repository with the actual disk contents. If you run a `CROSSCHECK`, then the RMAN repository is synchronized, and recovery can proceed by applying available copies of the log or restoring a backup of the log if no disk copies are available.

See Also: *Oracle Database Backup and Recovery Reference* for `CROSSCHECK` syntax

Incomplete or Point-In-Time Recovery

RMAN can perform either complete or point-in-time recovery (also called incomplete recovery). You can specify a time, SCN, restore point, or log sequence number as a limit for incomplete recovery with the `SET UNTIL` command or with an `UNTIL` clause specified directory on the `RESTORE` and `RECOVER` commands. The easiest method is run the `SET UNTIL` command before issuing the `RESTORE` and `RECOVER` commands. After performing incomplete recovery, you must open the database with the `RESETLOGS` option.

See Also: *Oracle Database Backup and Recovery Reference* for the `UNTIL` clause syntax

Tablespace Point-in-Time Recovery

Recovery Manager automated Tablespace Point-in-Time Recovery (TSPITR) enables you to recover one or more tablespaces to a point in time that is different from that of the rest of the database. RMAN TSPITR is most useful in these cases:

- To recover from an erroneous drop or truncate table operation
- To recover a table that has become logically corrupted
- To recover from an incorrect batch job or other DML statement that has affected only a subset of the database
- In cases where there are multiple logical schemas in separate tablespaces of one physical database, and where one schema must be recovered to a point different from that of the rest of the physical database
- For VLDBs (very large databases), even if a full database point-in-time recovery would suffice, you might choose to do tablespace point-in-time recovery rather than restore the whole database from a backup and perform a complete database roll forward

Similar to a table export, RMAN TSPITR enables you to recover a consistent data set; however, the data set is the entire tablespace rather than a single object.

See Also: [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#) to learn how to perform TSPITR using RMAN

Block Media Recovery with RMAN

Although datafile media recovery is the principal form of recovery, you can also use the RMAN `BLOCKRECOVER` command to perform **block media recovery**. Block media recovery recovers an individual corrupt datablock or set of datablocks within a datafile. In cases when a small number of blocks require media recovery, you can selectively restore and recover damaged blocks rather than whole datafiles.

Block media recovery provides several advantages over datafile media recovery. For example, block media recovery

- Lowers the **Mean Time to Recovery (MTTR)** because only blocks needing recovery are restored and only necessary corrupt blocks undergo recovery. Block media recovery minimizes redo application time and avoids I/O overhead during recovery.
- Allows affected datafiles to remain online during recovery of the blocks. Without block-level recovery, if even a single block is corrupt, then you must restore a backup of the entire datafile and apply all redo generated for that file after the backup was created.

Note these restrictions of block media recovery:

- You can only perform block media recovery with RMAN. No SQL*Plus recovery interface is available.
- You can only perform complete recovery of individual blocks. In other words, you cannot stop recovery before all redo has been applied to the block.
- You can only recover blocks marked media corrupt. The `V$DATABASE_BLOCK_CORRUPTION` view indicates which blocks in a file were marked corrupt since the most recent `BACKUP` or `BACKUP . . . VALIDATE` command was run against the file.
- You must have a full RMAN backup. Incremental backups are not used by block media recovery. Proxy backups are also not used by block media recovery. Only full backups and archived log files are used.
- Block media recovery is able to restore blocks from parent incarnation backups and recover the corrupted blocks through a `RESETLOGS`.
- Blocks that are marked media corrupt are not accessible to users until recovery is complete. Any attempt to use a block undergoing media recovery results in an error message indicating that the block is media corrupt.

See Also: ["Performing Disaster Recovery"](#) on page 7-10 and *Oracle Database Backup and Recovery Reference* for `BLOCKRECOVER` syntax

When to Use Block Media Recovery

Block media recovery is not intended for cases where the extent of data loss or corruption is unknown and the entire datafile requires recovery. In such cases, datafile media recovery is the best solution. Block media recovery is not a replacement for traditional datafile media recovery, but a supplement to it.

In most cases, the database marks a block as media corrupt, invalidates the block in the instances (or all enabled instances in a Real Application Clusters configuration), and then writes it to disk when the corruption is first encountered. No subsequent read of the block will be successful until the block is recovered. You can only perform block recovery on blocks that are marked corrupt. This corrupt status effectively takes the block offline in all database instances and prevents user access during recovery.

Block media recovery is most useful for data losses that affect specific blocks. Block-level data loss usually results from intermittent, random I/O errors that do not cause widespread data loss, as well as memory corruptions that get written to disk. Typically, these types of block corruption are reported in the following locations:

- Error messages in standard output
- The alert log
- User trace files
- Results of the SQL commands `ANALYZE TABLE` and `ANALYZE INDEX`
- Results of the `DBVERIFY` utility
- Third-party media management output

For example, you may discover the following messages in a user trace file:

```
ORA-01578: ORACLE data block corrupted (file # 7, block # 3)
ORA-01110: data file 7: '/oracle/oradata/trgt/tools01.dbf'
ORA-01578: ORACLE data block corrupted (file # 2, block # 235)
ORA-01110: data file 2: '/oracle/oradata/trgt/undotbs01.dbf'
```

You can then specify the corrupt blocks in the BLOCKRECOVER command as follows:

```
BLOCKRECOVER
  DATAFILE 7 BLOCK 3
  DATAFILE 2 BLOCK 235;
```

Block Media Recovery When Redo Is Missing


Like datafile media recovery, block media recovery cannot generally survive a missing or inaccessible archived log (although it will attempt restore failover when looking for usable copies of archived redo log files, as described in "Restore Failover" on page 3-3). Nevertheless, block media recovery can survive gaps in the redo stream if the missing or corrupt redo records do not affect the blocks being recovered. Whereas datafile recovery requires an unbroken series of redo changes from the beginning of recovery to the end, block media recovery only requires an unbroken set of redo changes for the blocks being recovered.

Note: Each block is recovered independently during block media recovery, so recovery may be successful for a subset of blocks.

When RMAN first detects missing or corrupt redo records during block media recovery, it does not immediately signal an error because the block undergoing recovery may become a **new block** later in the redo stream. When a block is new all previous redo for that block becomes irrelevant because the redo applies to an old incarnation of the block. For example, the database can new a block when users delete all the rows recorded in the block or drop a table.

Assume that media recovery is performed on block 13 as depicted in Figure 3-2.

Figure 3-2 Performing RMAN Media Recovery



After block recovery begins, RMAN discovers that change 120 is missing. RMAN does not terminate recovery in the hope that block 13 will be new later in the redo stream. Assume that in change 140 a user drops the table `EMPLOYEE` stored in block 13. At this point, the database formats block 13 as a new block. Because the redo for block 13 in change 120 related to the `EMPLOYEE` table, and the `EMPLOYEE` table was dropped in change 140, RMAN can skip this missing change and apply the redo between changes 140 and 160.

Database Duplication with RMAN

Use the RMAN `DUPLICATE` command to create a copy of the target database in another location. The command restores backups of the primary database files and creates a new database.

As part of the duplication, RMAN manages the following:

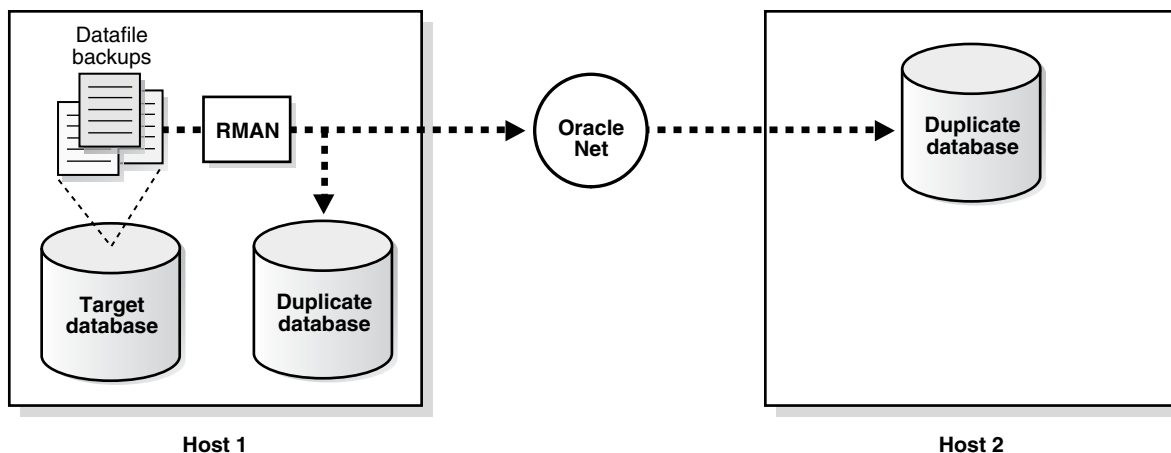
- Restores the target datafiles into the duplicate database and performs incomplete recovery using all available archived log and incremental backups
- Opens the duplicate database with the `RESETLOGS` option after incomplete recovery to create the online redo logs
- Generates a new, unique database identifier for the duplicate database

Note the following features of RMAN duplication. You can:

- Skip read-only tablespaces with the `SKIP READONLY` clause (read-only tablespaces are included by default). You can also exclude any tablespace with the `SKIP TABLESPACE` clause so long as it is not the `SYSTEM` or `SYSAUX` tablespace and does not contain rollback or undo data. If you omit tablespaces, then you can add them later.
- Create the duplicate database in a new host. If the same directory structure is available, then you can use the `NOFILENAMECHECK` option and reuse the target datafile filenames for the duplicate datafiles.
- Create the duplicate database by using the `SET UNTIL` command or `UNTIL` clause of the `DUPLICATE` command to recover it to a past time. By default, the `DUPLICATE` command creates the database using the most recent backups of the target database and then performs recovery to the most recent consistent point contained in the incremental and archived redo log backups.
- Use the duplicate database without a recovery catalog.
- Register the duplicate database in the same recovery catalog as the target database. This option is possible because the duplicate database receives a new database identifier during duplication. If you copy the target database with operating system utilities, then the database identifier of the copied database remains the same so you cannot register it in the same recovery catalog (unless you change its `DBID` with the `DBNEWID` utility, described in *Oracle Database Utilities*).

Figure 3-3 illustrates a case of database duplication. In this example, RMAN creates two duplicate database by using one set of datafile backups: one database on the local host and one database on a remote host.

Figure 3-3 Creating a Duplicate Database from Backups



The method you use to duplicate your database depends on whether you are creating your duplicate database on the same or a different host and whether the duplicate

directory structure is the same as your target database directory structure. For example, in some cases you can keep the same directory structure and filenames in your duplicate database, while other times you must rename the files.

See Also:

- [Chapter 13, "Creating and Updating Duplicate Databases with RMAN"](#) to learn how to make a duplicate database
- *Oracle Database Backup and Recovery Reference* for DUPLICATE command syntax
- *Oracle Database Utilities* to learn how to use the DBNEWID utility

Physical Standby Database Creation with RMAN

You can use the RMAN DUPLICATE command to create a physical standby database. (Note that RMAN cannot be used to create a logical standby database, because it is not a block-for-block duplicate of the primary database.)

RMAN automates the following steps of the creation procedure:

1. Restores the standby control file.
2. Restores the primary datafile backups.
3. Optionally, RMAN recovers the standby database (after the control file has been mounted) up to the specified time or to the latest archived redo log generated.
4. RMAN leaves the database mounted so that the user can activate it, place it in manual or managed recovery mode, or open it in read-only mode.

RMAN cannot fully automate creation of the standby database because you must manually create an initialization parameter file for the standby database, start the standby instance without mounting the control file, and perform any Oracle Net setup required before performing the creation of the standby. Also, you must have RMAN backups of all datafiles available and a control file backup that is usable as a standby control file, and those backups must be accessible by the standby instance under the same name.

Once the standby database is created, RMAN can be used to back up the datafiles, control file and archived redo logs of the standby. Backups of datafiles and archived redo logs taken from a physical standby database are fully interchangeable with primary backups. In other words, you can restore a backup of a physical standby datafile to the primary database, and you can restore a backup of a primary datafile to the physical standby database. The standby control file backups can be used to restore the standby control file without needing to re-instantiate the standby in cases where the standby control file is lost.

See Also: *Oracle Data Guard Concepts and Administration* to learn how to create and back up a physical standby database with RMAN

Connecting to Databases with RMAN

This chapter gives detailed instructions for starting the Recovery Manager (RMAN) command-line interface and making database connections. This chapter contains these topics:

- [Starting RMAN Without Connecting to a Database](#)
- [Connecting to a Target Database and a Recovery Catalog](#)
- [Connecting to an Auxiliary Database](#)
- [Hiding Passwords When Connecting to Databases](#)
- [Sending RMAN Output Simultaneously to the Terminal and a Log File](#)
- [Using the RMAN Pipe Interface](#)

Starting RMAN Without Connecting to a Database

You can start RMAN at the operating system command line without connecting to a database by issuing the `RMAN` command without any arguments. For example, enter:

```
% rman
```

If you did not specify the `CMDFILE`, `SCRIPT` or `@` option at the command line, then RMAN displays the RMAN prompt:

```
RMAN>
```

After the RMAN prompt is displayed, you can issue further commands to connect to the target database, recovery catalog database, or auxiliary database.

If you start RMAN without specifying either `CATALOG` or `NOCATALOG` on the command line, then RMAN makes no recovery catalog connection. The first time a command is issued that requires the RMAN repository, RMAN performs the operation in `NOCATALOG` mode if you have not connected to a recovery catalog yet. After that point, the `CONNECT CATALOG` command can no longer be used without exiting and restarting the RMAN client.

Connecting to a Target Database and a Recovery Catalog

The following examples use the following sample connection data, and assume that a recovery catalog is being used.

Value	Description
<code>SYS</code>	User with <code>SYSDBA</code> privileges

Value	Description
oracle	The password for connecting as SYSDBA specified in the target database's password file
trgt	The net service name for the target database
rman	User that owns the recovery catalog schema. This is a user defined in the recovery catalog database that has been granted the RECOVERY_CATALOG_OWNER role.
cat	The password for connecting to the recovery catalog as user RMAN
catdb	The net service name for the recovery catalog database

Connecting to the Target Database and Recovery Catalog from the Command Line

You can specify either operating system or Oracle Net authentication information on the command line when you start the RMAN client.

This example shows how to use operating system authentication to connect to the target database and Oracle Net authentication for the recovery catalog:

```
% rman TARGET / CATALOG rman/cat@catdb
```

This example shows how to use Oracle Net authentication to connect to both the target database and the recovery catalog:

```
% rman TARGET SYS/oracle@trgt CATALOG rman/cat@catdb
```

Connecting to the Target Database and Recovery Catalog from the RMAN Prompt

You can also start RMAN and connect to the target database from the RMAN prompt. The following example uses operating system authentication for the target database and Oracle Net authentication for the recovery catalog:

```
% rman
RMAN> CONNECT TARGET /
RMAN> CONNECT CATALOG rman/cat@catdb
```

The following example uses Oracle Net password file authentication for the target database, which requires that the target database be using a password file, that defines the password for user SYS to be 'oracle'. Oracle Net authentication is also used for the recovery catalog.

```
% rman
RMAN> CONNECT TARGET SYS/oracle@trgt
RMAN> CONNECT CATALOG rman/cat@catdb
```

See Also: *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide* for details on connecting RMAN to a RAC cluster.

Connecting to an Auxiliary Database

To use the DUPLICATE command, you need to connect to an auxiliary instance. To perform RMAN TSPITR, you may also need to connect to an auxiliary instance if, for example, you do not let RMAN manage the auxiliary instance for you.

See Also:

- [Chapter 13, "Creating and Updating Duplicate Databases with RMAN"](#) for more details on using the DUPLICATE command
- [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#) for more details on performing TSPITR

If the auxiliary instance uses a password file for authentication, then you can connect using a password for either local or remote access. If you are connecting remotely through a net service name, then authentication through a password file is mandatory.

The following dummy values have been substituted into the following examples:

Value	Description
aux	The password for connecting as SYSDBA specified in the auxiliary database's orapwd file
auxdb	The net service name for the auxiliary database

Connecting to an Auxiliary Database from the Command Line

To launch RMAN connected to an auxiliary instance from the operating system command line, enter the following:

```
% rman AUXILIARY SYS/aux@auxdb
```

To connect to target, auxiliary, and recovery catalog databases, launch the RMAN client with these command line arguments:

```
% rman TARGET SYS/oracle@trgt AUXILIARY SYS/aux@auxdb CATALOG rman/cat@catdb
```

Connecting to an Auxiliary Database from the RMAN Prompt

To launch RMAN without connecting to an auxiliary, and connect to the auxiliary database from the RMAN prompt, enter the following commands:

```
% rman
RMAN> CONNECT AUXILIARY SYS/aux@auxdb
```

To connect to the target, auxiliary, and recovery catalog databases from within RMAN, enter the following commands:

```
% rman
RMAN> CONNECT TARGET SYS/oracle@trgt
RMAN> CONNECT CATALOG rman/cat@catdb
RMAN> CONNECT AUXILIARY
SYS/aux@auxdb
```

Diagnosing Connection Problems

When diagnosing errors RMAN encounters in connecting to the target, catalog and auxiliary databases, using SQL*Plus to connect to the databases directly can reveal underlying problems with the connection information or the databases.

Diagnosing Target and Auxiliary Database Connection Problems

RMAN always connects to target and auxiliary databases using the SYSDBA role. Thus, when using SQL*Plus to diagnose connection problems to the target or auxiliary databases, request a SYSDBA connection to reproduce RMAN's behavior.

For example, if the following RMAN command encountered connection errors:

```
RMAN> CONNECT target sys/oracle@target
```

you would reproduce the connection attempt with the SQL*Plus command:

```
SQL> CONNECT sys/oracle@target AS SYSDBA
```

Diagnosing Recovery Catalog Connection Problems

When RMAN connects to the recovery catalog database, it **does not** use the SYSDBA role. So, when you are using SQL*Plus to diagnose connection problems to the recovery catalog database, you must enter the catalog connect string **exactly** as it was entered into RMAN. Do not also specify `AS SYSDBA`.

Hiding Passwords When Connecting to Databases

If you create an RMAN command file which uses a `CONNECT` command with database level credentials (user name and password), then anyone with read access to this file can learn the password. There is no secure way to incorporate a `CONNECT` string with a password into a command file.

It is also possible, using the `ps` command under Unix or some similar command under other operating systems, to view command lines and arguments entered into the shell or other host operating system command line interpreter. Therefore, it is risky to invoke RMAN with a command line like this example:

```
% rman TARGET sys/oracle@target
```

To connect to RMAN from the operating system command line and hide authentication information, you can start RMAN without connecting to databases, and then enter `CONNECT` commands at the RMAN prompt. You can also start RMAN without a password in the connect string, as in this example:

```
% rman TARGET sys@target
```

RMAN will prompt for a password in such a case.

If you create an RMAN command file which uses a `CONNECT` command that includes authentication information, RMAN does not echo the connect string when you run the command file with the "@" command. This prevents connect strings from appearing in any log files that contain RMAN output.

For example, create a command file `listbkup.rman` which reads:

```
CONNECT target sys/oracle@target
LIST BACKUP;
```

Then execute this script by running RMAN with the @ command line option:

```
% rman @listbkup.rman
```

When the command file executes, RMAN replaces the connection string with an asterisk, as shown in the following output:

```

Recovery Manager: Release 10.2.0.1.0 - Production

Copyright (c) 1995, 2005, Oracle. All rights reserved.

RMAN> connect target *
2> list backup;
3>
connected to target database: RDBMS (DBID=771530996)

using target database control file instead of recovery catalog

List of Backup Sets
=====
...rest of output omitted

```

Sending RMAN Output Simultaneously to the Terminal and a Log File

If you specify the `LOG` option at the command line, then RMAN displays command input but does not display the RMAN output. The easiest way to send RMAN output both to a log file and to standard output is to use the UNIX `tee` command or its equivalent on another operating system. For example:

```
% rman | tee rman.log
RMAN>
```

In this way, both input and output are visible within the RMAN command-line interface.

Using the RMAN Pipe Interface

The RMAN pipe interface is an alternative method for issuing commands to RMAN and receiving the output from those commands. With this interface, RMAN obtains commands and sends output by using the `DBMS_PIPE` PL/SQL package instead of the operating system shell. Using this interface, it is possible to write a portable programmatic interface to RMAN.

The pipe interface is invoked by using the `PIPE` command-line parameter for the RMAN client. RMAN uses two private pipes: one for receiving commands and the other for sending output. The names of the pipes are derived from the value of the `PIPE` parameter. For example, you can invoke RMAN with the following command:

```
% rman PIPE abc TARGET SYS/oracle@trgt
```

RMAN opens the two pipes in the target database: `ORA$RMAN_ABC_IN`, which RMAN uses to receive user commands, and `ORA$RMAN_ABC_OUT`, which RMAN uses to send all output back to RMAN. All messages on both the input and output pipes are of type `VARCHAR2`.

Note that RMAN does not permit the pipe interface to be used with public pipes, because they are a potential security problem. With a public pipe, any user who knows the name of the pipe can send commands to RMAN and intercept its output.

If the pipes are not already initialized, then RMAN creates them as private pipes. If you want to put commands on the input pipe before starting RMAN, you must first create the pipe by calling `DBMS_PIPE.CREATE_PIPE`. Whenever a pipe is not explicitly created as a private pipe, the first access to the pipe automatically creates it as a public pipe, and RMAN returns an error if it is told to use a public pipe.

Note: If multiple RMAN sessions can run against the target database, then you must use unique pipe names for each session of RMAN. The `DBMS_PIPE.UNIQUE_SESSION_NAME` function is one method that can be used to generate unique pipe names.

Executing Multiple RMAN Commands In Succession Through a Pipe: Example

This scenario assumes that the application controlling RMAN wants to run multiple commands in succession. After each command is sent down the pipe and executed and the output returned, RMAN will pause and wait for the next command.

1. Start RMAN by connecting to a target database (required) and specifying the `PIPE` option. For example, issue:

```
% rman PIPE abc TARGET SYS/oracle@trgt
```

You can also specify the `TIMEOUT` option, which forces RMAN to exit automatically if it does not receive any input from the input pipe in the specified number of seconds. For example, enter:

```
% rman PIPE abc TARGET SYS/oracle@trgt TIMEOUT = 60
```

2. Connect to the target database and put the desired commands on the input pipe by using `DBMS_PIPE.PACK_MESSAGE` and `DBMS_PIPE.SEND_MESSAGE`. In pipe mode, RMAN issues message `RMAN-00572` when it is ready to accept input instead of displaying the standard RMAN prompt.
3. Read the RMAN output from the output pipe by using `DBMS_PIPE.RECEIVE_MESSAGE` and `DBMS_PIPE.UNPACK_MESSAGE`.
4. Repeat steps 2 and 3 to execute further commands with the same RMAN instance that was started in step 1.
5. If you used the `TIMEOUT` option when starting RMAN, RMAN terminates automatically after not receiving any input for the specified length of time. To force RMAN to terminate immediately, send the `EXIT` command.

Executing RMAN Commands In a Single Job Through a Pipe: Example

This scenario assumes that the application controlling RMAN wants to run one or more commands as a single job. After running the commands that are on the pipe, RMAN will exit.

1. After connecting to the target database, create a pipe (if it does not already exist under the name `ORA$RMAN_pipe_IN`).
2. Put the desired commands on the input pipe. In pipe mode, RMAN issues message `RMAN-00572` when it is ready to accept input instead of displaying the standard RMAN prompt.
3. Start RMAN with the `PIPE` option, and specify `TIMEOUT = 0`. For example, enter:

```
% rman PIPE abc TARGET SYS/oracle@trgt TIMEOUT = 0
```

4. RMAN reads the commands that were put on the pipe and executes them by using `DBMS_PIPE.PACK_MESSAGE` and `DBMS_PIPE.SEND_MESSAGE`. When it has exhausted the input pipe, RMAN exits immediately.
5. Read RMAN output from the output pipe by using `DBMS_PIPE.RECEIVE_MESSAGE` and `DBMS_PIPE.UNPACK_MESSAGE`.

See Also: *PL/SQL Packages and Types Reference* for documentation on the DBMS_PIPE package

Part II

Advanced RMAN Backup and Recovery Topics

Part II describes how to use the RMAN utility to perform advanced backup and recovery operations, and explains RMAN performance tuning and troubleshooting.

This part contains these chapters:

- [Chapter 4, "Connecting to Databases with RMAN"](#)
- [Chapter 5, "Configuring the RMAN Environment: Advanced Topics"](#)
- [Chapter 6, "Making Backups with RMAN: Advanced Topics"](#)
- [Chapter 7, "Advanced RMAN Recovery Techniques"](#)
- [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#)
- [Chapter 9, "RMAN Backup and Repository Maintenance"](#)
- [Chapter 10, "Managing the Recovery Catalog"](#)
- [Chapter 11, "Tuning Backup and Recovery"](#)
- [Chapter 12, "Recovery Manager Troubleshooting"](#)

Configuring the RMAN Environment: Advanced Topics

This chapter describes how to perform setup and configuration tasks. This chapter contains these topics:

- [Configuring the Flash Recovery Area: Advanced Topics](#)
- [Configuring RMAN to Make Backups to a Media Manager](#)
- [Configuring Channels](#)
- [Configuring the Maximum Size of Backup Sets and Pieces](#)
- [Configuring Backup Optimization](#)
- [Configuring Backup Duplexing: CONFIGURE... BACKUP COPIES](#)
- [Configuring Tablespaces for Exclusion from Whole Database Backups](#)
- [Setting the Snapshot Control File Location](#)
- [Setting Up RMAN for Use with a Shared Server](#)

See Also: *Oracle Database Backup and Recovery Basics* for basic RMAN configuration information

Configuring the Flash Recovery Area: Advanced Topics

To take maximum advantage of the flash recovery area, it should be used to store and manage as many different types of file as possible: online redo logs, archived redo logs, control files and RMAN's own working files for backup and restore operations.

This section contains the following topics:

- [Configuring Online Redo Log Creation in the Flash Recovery Area](#)
- [Configuring Control File Creation in the Flash Recovery Area](#)
- [Archived Redo Log Creation in the Flash Recovery Area](#)
- [RMAN File Creation in the Flash Recovery Area](#)

Configuring Online Redo Log Creation in the Flash Recovery Area

The following statements can create online redo logs in the flash recovery area:

- `CREATE DATABASE`
- `ALTER DATABASE ADD LOGFILE`

- ALTER DATABASE ADD STANDBY LOGFILE
- ALTER DATABASE OPEN RESETLOGS

The default size of an online log created in the flash recovery area is 100 MB. The log member filenames are automatically generated by the database.

The initialization parameters that determine where online redo log files are created are `DB_CREATE_ONLINE_LOG_DEST_n`, `DB_RECOVERY_FILE_DEST` and `DB_CREATE_FILE_DEST`. Details of the effect of various combinations of these parameters on online redo log creation can be found in *Oracle Database SQL Reference* in the description of the `LOGFILE` clause of the `CREATE DATABASE` statement.

Configuring Control File Creation in the Flash Recovery Area

The initialization parameters `CONTROL_FILES`, `DB_CREATE_ONLINE_LOG_DEST_n`, `DB_RECOVERY_FILE_DEST`, and `DB_CREATE_FILE_DEST` all interact to determine the location where the database control files are created.

For a full description of how these parameters interact, see the "Semantics" section of the description of `CREATE CONTROLFILE` in *Oracle Database SQL Reference*.

If the database creates an Oracle managed control file, and if the database uses a server parameter file, then the database sets the `CONTROL_FILES` initialization parameter in the server parameter file. If the database uses a client-side initialization parameter file, then you must set the `CONTROL_FILES` initialization parameter manually in the initialization parameter file.

Archived Redo Log Creation in the Flash Recovery Area

It is recommended that you use the flash recovery area as an archived log location because the archived logs are automatically managed by the database. Whatever archiving scheme you choose, it is always advisable to create multiple copies of archived logs.

You have the following basic options, listed from most to least recommended:

1. Enable archiving to the flash recovery area *only* and use disk mirroring to create the redundancy needed to protect the archived redo logs.
2. Enable archiving to the flash recovery area and set other `LOG_ARCHIVE_DEST_n` initialization parameter to locations outside the flash recovery area.
3. Set `LOG_ARCHIVE_DEST_n` initialization parameters to archive *only* to non-flash recovery area locations.

If you want to use the flash recovery area, you cannot use the `LOG_ARCHIVE_DEST` and `LOG_ARCHIVE_DUPLEX_DEST` initialization parameters. You must use instead the `LOG_ARCHIVE_DEST_n` parameters, which have somewhat different semantics. Once your database is using `LOG_ARCHIVE_DEST_n`, you can configure a flash recovery area.

Rules for Initialization Parameters Affecting Redo Log File Destinations

The interactions among different initialization parameters affecting redo log archiving distinctions are as follows:

- If `LOG_ARCHIVE_DEST` (and, optionally, `LOG_ARCHIVE_DUPLEX_DEST`) is set, these parameters will specify the only redo log archiving destinations.
- If `DB_RECOVERY_FILE_DEST` is specified (that is, if a flash recovery area is configured) and no `LOG_ARCHIVE_DEST_n` is specified, then `LOG_ARCHIVE_`

`DEST_10` is implicitly set to the flash recovery area. (You can override this behavior by explicitly setting `LOG_ARCHIVE_DEST_10` to an empty string.)

- If you set any local destinations for `LOG_ARCHIVE_DEST_n`, then archived redo logs are stored only in the destinations you specify using those parameters. In this case, redo log files are not archived in the flash recovery area by default. If you have a flash recovery area configured, you can explicitly add the flash recovery area to the set of archiving destinations by setting one of the `LOG_ARCHIVE_DEST_n` parameters to `LOCATION=USE_DB_RECOVERY_FILE_DEST` (note that this does not have to be `LOG_ARCHIVE_DEST_10`).
- If you do not set any value for `LOG_ARCHIVE_DEST`, `LOG_ARCHIVE_DEST_n`, or `DB_RECOVERY_FILE_DEST`, then the redo logs are archived to a default location that is platform-specific. On Solaris, for example, the default is: `~/dfs`.

Filenames for Archived Redo Log Files in the Flash Recovery Area

The generated filenames for the archived redo logs in the flash recovery area are Oracle Managed Filenames and are not determined by `LOG_ARCHIVE_FORMAT`.

RMAN File Creation in the Flash Recovery Area

This section describes RMAN commands or implicit actions (such as control file autobackup) that can create files in the flash recovery area, and how to control whether a specific command creates files there or in some other destination. The assumption in all cases is that a flash recovery area has already been configured for your database. The commands are:

- **BACKUP**

Do not specify a `FORMAT` option to the `BACKUP` command, and do not configure a `FORMAT` option for disk backups. In such a case, RMAN creates backup pieces and image copies in the flash recovery area, with names in Oracle Managed Files name format.

- **Control File Autobackup**

RMAN can create control file autobackups in the flash recovery area. Use the RMAN command `CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK CLEAR` to clear any configured format option for the control file autobackup location on disk. Control file autobackups will be placed in the flash recovery area when no other destination is configured.

- **RESTORE ARCHIVELOG**

Explicitly or implicitly (as in the case of), set one of the `LOG_ARCHIVE_DEST_n` parameters to `'LOCATION=USE_DB_RECOVERY_FILE_DEST'`. If you do not specify `SET ARCHIVELOG DESTINATION` to override this behavior, then restored archived redo log files will be stored in the flash recovery area.

- **RECOVER DATABASE or TABLESPACE, BLOCKRECOVER, and FLASHBACK DATABASE**

These commands restore archived redo logs from backup for use during media recovery, as required by the command. RMAN restores any redo log files needed during these operations to the flash recovery area, and delete them once they are applied during media recovery.

To direct the restored archived redo logs to the flash recovery area, set one of the `LOG_ARCHIVE_DEST_n` parameters to `'LOCATION=USE_DB_RECOVERY_FILE_`

DEST", and make sure you are not using SET ARCHIVELOG DESTINATION to direct restored archived logs to some other destination.

Configuring RMAN to Make Backups to a Media Manager

On most platforms, to back up to and restore from sequential media such as tape you must integrate a media manager with your Oracle database. A media manager is not an Oracle product and must be obtained from a third-party vendor. If you choose to use RMAN with a media manager, then you must obtain all product-specific information from the vendor.

This section describes the generic steps for configuring RMAN for use with a media manager. The actual steps depend on the media management product that you install and the platform on which you are running the database.

Read the following sections in order when configuring the media manager:

1. [Prerequisites for Using a Media Manager with RMAN](#)
2. [Locating the Media Management Library: The SBT_LIBRARY Parameter](#)
3. [Testing Whether the Media Manager Library Is Integrated Correctly](#)
4. [Configuring SBT Channels for Use with a Media Manager](#)

See Also: ["Media Management"](#) on page 1-7 for an overview of media management software and its implications for RMAN

Prerequisites for Using a Media Manager with RMAN

Before you can begin using RMAN with a media manager, you must install it and make sure that RMAN can communicate with it. Instructions for this procedure should be available in the media manager vendor's software documentation.

In general, you should begin by installing and configuring the media management software on the target host or production network. Ensure that you can make non-RMAN backups of operating system files on the target database host. This step makes later troubleshooting much easier, by confirming that the basic integration of the media manager with the target host has been successful. Refer to your media management documentation to learn how to back up files to the media manager outside of RMAN.

Then, obtain and install the third-party media management module for integration with the database server. This module contains the media management library that the Oracle database loads and uses when accessing the media manager. It is generally a third-party product which must be purchased separately. Contact your media management vendor for details.

Locating the Media Management Library: The SBT_LIBRARY Parameter

When allocating or configuring channels for RMAN to use to communicate with a media manager, specify the SBT_LIBRARY parameter to provide the path to the media management software library. When RMAN actually allocates channels to communicate with a media manager, it attempts to load the library indicated by the SBT_LIBRARY parameter.

If you do not provide a value for this parameter, RMAN looks in a platform-specific default location. On UNIX, the default library filename is \$ORACLE_HOME/lib/libobk.so, with the extension name varying according to platform: .so,

.sl, .a, and so forth. On Windows the default library location is %ORACLE_HOME%\bin\orasbt.dll.

Note: The default media management library file is *not* part of the standard database installation. It is only present if you install third-party media management software.

If the database is unable to locate a media management library in the location specified by the `SBT_LIBRARY` parameter or the default location, then RMAN issues an `ORA-27211` error and exits.

Whenever channel allocation fails, the database writes a trace file to the `USER_DUMP_DEST` directory. The following shows sample output:

```
SKGFQ OSD: Error in function sbtinit on line 2278
SKGFQ OSD: Look for SBT Trace messages in file /oracle/rdbms/log/sbtio.log
SBT Initialize failed for /oracle/lib/libobk.so
```

See Also:

- Your operating system specific database documentation and the documentation supplied by your media vendor for instructions on how to achieve media manager integration on your platform
- ["After Installation of Media Manager, RMAN Channel Allocation Fails: Scenario"](#) on page 12-12 for troubleshooting scenarios involving media manager problems

Testing Whether the Media Manager Library Is Integrated Correctly

After you have confirmed that the database server can load the media management library, test to make sure that RMAN can back up to the media manager. The process for testing the media management library is described in the following sections:

- [Configuring Media Management Software for RMAN Backups](#)
- [Testing `ALLOCATE CHANNEL` on the Media Manager](#)
- [Testing a Backup to the Media Manager](#)

Configuring Media Management Software for RMAN Backups

After installing the media management software, perform whatever configuration that your vendor requires so that the software can accept RMAN backups. Depending on the type of media management software that you installed, you may have to define media pools, configure users and classes, and so forth.

Then, determine which `PARMS` settings are needed for the `ALLOCATE CHANNEL` or `CONFIGURE CHANNEL` commands as well as the recommended `FORMAT` string for the `BACKUP` command (if needed). The `PARMS` parameter sends instructions to the media manager. For example, the following vendor-specific `PARMS` setting instructs the media manager to back up to a volume pool called `oracle_tapes`:

```
PARMS= 'ENV= (NSR_DATA_VOLUME_POOL=oracle_tapes) '
```

Refer to your third-party vendor documentation for the appropriate settings.

See Also:

- *Oracle Database Backup and Recovery Reference* for `ALLOCATE CHANNEL` syntax
- *Oracle Database Backup and Recovery Reference* for channel control options

Configuring Backup Piece Names and Sizes for a Media Manager

To work with restrictions on file names and sizes imposed by your media manager, you may need to configure RMAN settings that control the naming and size of backup pieces.

Configuring Backup Piece Names for RMAN Backups to a Media Manager You may need to manage the naming of backup pieces to be written to the media manager, so that backup pieces have unique names. A backup piece name is determined by the `FORMAT` string specified in the `BACKUP` command, the `CONFIGURE CHANNEL` command, or the `ALLOCATE CHANNEL` command. The media manager considers the backup piece name as the filename of the backup file, so this name must be unique in the media manager catalog.

You can use the substitution variables provided by RMAN to generate unique backup piece names. If you do not specify the `FORMAT` parameter, then RMAN automatically generates a unique filename with the `%U` substitution variable.

Note: Refer to your media management documentation to determine the string character limit for the media manager. For example, some media managers only support a 14-character backup piece name, and some require special `FORMAT` strings. The unique backup piece names generated by `%U` are less than 14 characters.

See Also: *Oracle Database Backup and Recovery Reference* for the complete list of variables allowable in format strings with the `BACKUP` command

Configuring Backup Piece Sizes for RMAN Backups to a Media Manager Some media managers have limits on the maximum size of files that they can back up or restore. You must ensure that RMAN does not produce backup sets larger than limits imposed by your media manager.

To limit backup piece sizes, use the parameter `MAXPIECESIZE`, which you can set in the `CONFIGURE CHANNEL` and `ALLOCATE CHANNEL` commands. Refer to the `*.rcv` scripts in the `demo` subdirectory on your system, which is located in an operating system specific location (`$ORACLE_HOME/rdbms` on UNIX) for an example.

See Also: *Oracle Database Backup and Recovery Reference* and "[Size of Backup Pieces](#)" on page 2-21 for details on how to set `MAXPIECESIZE`

Testing `ALLOCATE CHANNEL` on the Media Manager

Use the following steps to confirm that RMAN is able to load the media management library when allocating a channel for your media manager.

1. Start RMAN and connect to the target database. For example, enter:


```
% rman TARGET /
```

2. Run the `ALLOCATE CHANNEL` command with the `PARMS` required by your media management software. For example, run this command:

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt
  PARMS=' SBT_LIBRARY=/mediavendor/lib/libobk.so ENV=(NSR_SERVER=tape_srv,NSR_
GROUP=oracle_tapes)';
}
```

If you do not receive an error message, then the database successfully loaded the media management library. If you receive the `ORA-27211` error, the media management library could not be loaded:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03009: failure of allocate command on c1 channel at 11/30/2001 13:57:18
ORA-19554: error allocating device, device type: SBT_TAPE, device name:
ORA-27211: Failed to load Media Management Library
Additional information: 25
```

In this case, you must check your media management installation to make sure that the library is correctly installed, and re-check the value for the `SBT_LIBRARY` parameter as described in ["Locating the Media Management Library: The SBT_LIBRARY Parameter"](#) on page 5-4.

For any other errors, check the trace file in `USER_DUMP_DEST` directory for more information.

See Also: ["After Installation of Media Manager, RMAN Channel Allocation Fails: Scenario"](#) on page 12-12 for a troubleshooting scenario

Testing a Backup to the Media Manager

After testing a channel allocation on the media manager, make a test backup. For example, to test whether your backup goes successfully to tape, you might run the following command:

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt
  PARMS=' SBT_LIBRARY=/mediavendor/lib/libobk.so
ENV=(NSR_SERVER=tape_srv,NSR_GROUP=oracle_tapes)';
  BACKUP CURRENT CONTROLFILE;
}
```

The specifics of your `PARMS` and `FORMAT` settings depend on the media management software that you are using.

If the backup succeeds, then you are ready to make backups to your media manager.

Possible failures include the following cases:

Case	Response
The backup hangs.	A hanging backup usually indicates that the media manager is waiting to mount a tape. Check if there are any media manager jobs in "tape mount request" mode and fix the problem. Ensure that the steps in "Configuring RMAN to Make Backups to a Media Manager" on page 5-4 are correctly done. Refer to "Backup Job Is Hanging: Scenario" on page 12-13 if the problem persists.
The backup fails with an ORA-19511	This error indicates that the media management software is not correctly configured. Ensure that the steps in "Configuring RMAN to Make Backups to a Media Manager" on page 5-4 are correctly done. Also, ensure that you have the correct PARMS and FORMAT strings required by your media management software.

See Also: ["Testing the Media Management API"](#) on page 12-7 and ["RMAN Troubleshooting Scenarios"](#) on page 12-12 for more information about troubleshooting RMAN with a media manager

Configuring SBT Channels for Use with a Media Manager

This section describes how to configure channels specifically for use with a media manager. For an overview of configured channels and how they are used, refer to the section ["Configuring Channels"](#) on page 5-9. The following setup procedure references the sections in ["Configuring Channels"](#) where it is appropriate.

To configure channels for use with a media manager:

1. Configure a generic channel of DEVICE TYPE `sbt` as described in ["Configuring Channel Settings for a Device Type"](#) on page 5-9. In the configuration enter all parameters that you tested in the section ["Testing a Backup to the Media Manager"](#) on page 5-7. For example, assume that your media vendor requires PARMS settings as follows:

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE sbt
      PARMS='SBT_LIBRARY=/mediavendor/lib/libobk.so
      ENV=(NSR_SERVER=tape_svr,NSR_CLIENT=oracleclnt,NSR_GROUP=ora_tapes)'
      FORMAT "BACKUP_%U";
```

2. After configuring the channel, test by backing up something small, such as the control file:

```
RMAN> BACKUP DEVICE TYPE sbt CURRENT CONTROLFILE;
```

3. Check your configuration by running the following command:

```
RMAN> SHOW CHANNEL FOR DEVICE TYPE sbt;
```

4. Configure the default device to `sbt` so that RMAN sends all backups to the media manager. For example:

```
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
```

1. After configuring the default device, make a test backup to determine whether it is really going to the media manager:

```
RMAN> BACKUP CURRENT CONTROLFILE;
```

1. Check your configuration by running the following command:

```
RMAN> SHOW DEFAULT DEVICE TYPE;
```

- If you use more than one tape device, then you must specify the channel parallelism as described in ["Configuring Channel Parallelism"](#) on page 5-9. Assume that you want to back up to your media manager using two tape drives in parallel. In this case, you can run the following commands:

```

RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 2;
RMAN> BACKUP DATABASE;

```

Configuring Channels

You can configure persistent settings for your channels, such as channel parameters, parallelism, and the default device type for backups. The configured settings are stored in the RMAN repository. If you configure channel settings, then you do not have to use `ALLOCATE CHANNEL` commands with every RMAN backup, restore, recovery or maintenance command. Configuring persistent channel settings greatly simplifies the use of RMAN.

You can always override configured channels with `ALLOCATE CHANNEL` for a particular backup job surrounded by a `RUN` block.

By default, RMAN has preconfigured a disk channel so that you can back up to disk without doing any manual configuration. You may, however, want to parallelize the channels for disk or tape devices to improve performance.

See Also: ["About RMAN Channels"](#) on page 2-1 for a conceptual overview of configured and allocated channels, and *Oracle Database Backup and Recovery Reference* for syntax

Configuring Channel Parallelism

Configuring parallelism for a device type specifies the number of server sessions to be used for I/O to that device type. By default, channel parallelism for each configured device is set to 1. As a rule, allocating one channel for each physical device is best. If you are backing up to only one disk location or only one tape drive, then you need only one channel.

The `CONFIGURE DEVICE TYPE . . . PARALLELISM integer` command specifies how many channels (up to 254) RMAN should allocate for jobs on the specified device type. This command allocates three channels for jobs on device type `DISK` :

```
RMAN> CONFIGURE DEVICE TYPE DISK PARALLELISM 3;
```

These commands back up to a media manager using two tape drives in parallel:

```

RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt; # default backup device is tape
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 2; # configure two tape channels
RMAN> BACKUP DATABASE; # backup goes to two tapes, in two parallel streams

```

Each configured `sbt` channel will back up roughly half the total data.

See Also: ["Determining Channel Parallelism to Match Hardware Devices"](#) on page 2-7

Configuring Channel Settings for a Device Type

By default, RMAN allocates a one `DISK` channel with default options, and uses it for backup commands.

Note: This disk channel allocated by default is not the same channel as the **default channel**, a disk channel which RMAN creates when it first connects to the target instance, and generally does not use for activities such as backups and restores that require large amounts of I/O.

However, you may want to change the default DISK channel settings, for example, to specify a degree of parallelism or output locations for disk backups. Also, if you use a media manager, you must configure any required options for it, such as PARMs, FORMAT, MAXPIECESIZE, and so forth. By configuring channel settings, you define which parameters are used for channels RMAN allocates when you use configured channels for a backup job.

Use the CONFIGURE CHANNEL command to configure options for DISK and sbt channels. CONFIGURE CHANNEL takes the same options used to specify one-time options with ALLOCATE CHANNEL.

For example, you can configure default parameters for disk and tape channels as in this example:

```

RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT = '?/bkup_%U';
RMAN> CONFIGURE CHANNEL DEVICE TYPE sbt
      PARMs='SBT_LIBRARY=/mediavendor/lib/libobk.so ENV=(NSR_SERVER=tape_svr,NSR_
CLIENT=oracleclnt,NSR_GROUP=ora_tapes)';

```

You can configure generic channel settings for a device type, that is, a template that is used for any channels created based on configured settings for that device. If you set the PARALLELISM for a device, and then make the device default, then RMAN uses the generic configured channel settings for each parallelized channel.

Note that if you use CONFIGURE CHANNEL to specify generic channel settings for a device, any previous settings are discarded, even if the settings are not in conflict. For example, after the second CONFIGURE CHANNEL command, which specifies only a FORMAT for configured disk channels, the MAXPIECESIZE for the disk channel is returned to its default value:

```

RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2G;
RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT = /tmp/%U;

```

You can also configure default settings for individual channels from a group of parallelized channels by specifying a channel number.

Showing the Configured Channel Settings

The SHOW CHANNEL, SHOW DEVICE TYPE and SHOW DEFAULT DEVICE TYPE commands are used to display the current configured channel settings.

Showing the Currently Configured Channel Settings

After connecting to the target database and recovery catalog (if you use one), issue the SHOW CHANNEL command to display the currently configured channel settings. For example, connect the RMAN client to the target and, if applicable, the recovery catalog. Then enter:

```
RMAN> SHOW CHANNEL;
```

Sample output for SHOW CHANNEL follows:

```
RMAN configuration parameters are:
```

```
CONFIGURE CHANNEL DEVICE TYPE SBT RATE 1500K;
```

Showing the Configured Device Types

Issue the `SHOW DEVICE TYPE` command to display the configured devices and their `PARALLELISM` and backup type settings.

To show the default device type and currently configured settings for disk and sbt devices:

After connecting to the target database and recovery catalog (if you use one), run the `SHOW DEVICE TYPE` command. For example, enter:

```
SHOW DEVICE TYPE;      # shows the CONFIGURE DEVICE TYPE ... PARALLELISM settings
```

Sample output for `SHOW DEVICE TYPE` follows:

```

RMAN configuration parameters are:
CONFIGURE DEVICE TYPE DISK PARALLELISM 2 BACKUP TYPE TO COPY;
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 2 BACKUP TYPE TO BACKUPSET;
```

Note: As with all `SHOW` commands, the output of `SHOW DEVICE TYPE` is in the form of a valid `RMAN CONFIGURE` command. You can in fact enter one command, like those shown in the preceding sample output, to configure the backup type and parallelism simultaneously. Refer to the syntax diagrams for `CONFIGURE` in *Oracle Database Backup and Recovery Reference* for details on all of the possible ways of combining arguments to the `CONFIGURE` command.

Showing the Default Device Type

Issue the `SHOW DEFAULT DEVICE TYPE` command to display the settings for the default device type for backups. When you issue the `BACKUP` command, RMAN allocates only default channels of the type set by the `CONFIGURE DEFAULT DEVICE TYPE` command. This default device type setting is not in effect when you use commands other than `BACKUP`. Note that you cannot disable the default device type: it is always either `DISK` (default setting) or `sbt`.

To show the default device type for backups:

After connecting to the target database and recovery catalog (if you use one), run the `SHOW DEFAULT DEVICE TYPE` command. For example, enter:

```
SHOW DEFAULT DEVICE TYPE;      # shows the CONFIGURE DEFAULT DEVICE TYPE setting
```

Sample output for `SHOW DEFAULT DEVICE TYPE` follows:

```

RMAN configuration parameters are:
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT';
```

Manually Overriding Configured Channels

If you manually allocate a channel during a job, then RMAN disregards any configured channel settings. For example, assume that the default device type is configured to `sbt`, and you execute this command:

```

RMAN> RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE DISK;
```

```
    BACKUP TABLESPACE users;
}
```

In this case, RMAN uses only the disk channel that you manually allocated within the RUN block, overriding any defaults set by using `CONFIGURE DEVICE TYPE`, `CONFIGURE DEFAULT DEVICE`, or `CONFIGURE CHANNEL` settings.

See Also:

- ["About RMAN Channels"](#) on page 2-1 to learn about configured and allocated channels
- *Oracle Database Backup and Recovery Reference* for `ALLOCATE` syntax
- *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax

Configuring a Specific Channel for a Device Type

Besides configuring a generic channel for a device, you can also configure one or more specific channels for each device type by manually assigning your own channel numbers to the channels. Run the `CONFIGURE CHANNEL n` command (where *n* is a positive integer less than 255) to configure a specific channel. When manually numbering channels, you must specify one or more channel options (for example, `MAXPIECESIZE` or `FORMAT`) for each channel. When you use that specific numbered channel in a backup, the configured settings for that channel will be used instead of the configured generic channel settings.

Configure specific channels by number when it is necessary to control the parameters set for each channel separately. This could arise in the following situations:

- When running a Real Application Clusters (RAC) configuration in which individual nodes do not have access to the full set of backups, so different nodes must be configured with different connect strings so that all backups are accessible from some node
- When running a Real Application Cluster and using a media manager with multiple tape drives requiring different `PARMS` settings

Configuring Specific Channels: Examples

In this example, you want to send disk backups to two different disks. Configure disk channels as follows:

```
CONFIGURE DEFAULT DEVICE TYPE TO disk;          # backup goes to disk
CONFIGURE DEVICE TYPE sbt PARALLELISM 2;      # two channels used in parallel
CONFIGURE CHANNEL 1 DEVICE TYPE DISK FORMAT '/disk1/%U' # 1st channel to disk1
CONFIGURE CHANNEL 2 DEVICE TYPE DISK FORMAT '/disk2/%U' # 2nd channel to disk2
BACKUP DATABASE; # backup - first channel goes to disk1 and second to disk2
```

In this example, assume that you have two tape drives and want each tape drive to use tapes from a different tape pool. Configure your default output device and default sbt channels as follows:

```
CONFIGURE DEFAULT DEVICE TYPE TO sbt;          # backup goes to sbt
CONFIGURE DEVICE TYPE sbt PARALLELISM 2;      # two sbt channels will be allocated by default
# Assume media manager takes NSR_DATA_VOLUME_POOL to
# specify a pool
# Configure channel 1 to pool named first_pool
```

```

CONFIGURE CHANNEL 1 DEVICE TYPE sbt
  PARMS '
SBT_LIBRARY=/mediavendor/lib/libobk.so
  ENV=(NSR_DATA_VOLUME_POOL=first_pool)';
# configure channel 2 to pool named second_pool
CONFIGURE CHANNEL 2 DEVICE TYPE sbt
  PARMS '
SBT_LIBRARY=/mediavendor/lib/libobk.so
  ENV=(NSR_DATA_VOLUME_POOL=second_pool)';
BACKUP DATABASE; # first stream goes to 'first_pool' and second to 'second_pool'

```

Mixing Generic and Specific Channels

When parallelizing, RMAN always allocates channels beginning with CHANNEL 1 and ending with channel number equal to the PARALLELISM setting.

If you configure settings for a specific channel using CONFIGURE CHANNEL with a channel number, RMAN uses those specified configured settings. Otherwise, it uses the generic configuration for channels for that device type, as specified by the CONFIGURE CHANNEL command without a channel number.

Assume you enter the following channel configuration:

```

# disk channel configuration
CONFIGURE DEVICE TYPE DISK PARALLELISM 4;
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT = '/tmp/backup_%U';
CONFIGURE CHANNEL 2 DEVICE TYPE DISK MAXPIECESIZE = 20M;
CONFIGURE CHANNEL 4 DEVICE TYPE DISK MAXPIECESIZE = 40M;

# sbt channel configuration
CONFIGURE DEVICE TYPE sbt PARALLELISM 3;
CONFIGURE CHANNEL DEVICE TYPE sbt
  PARMS='SBT_LIBRARY=oracle.disksbt, ENV=(BACKUP_DIR=?/oradata)';
CONFIGURE CHANNEL 3 DEVICE TYPE sbt
  PARMS='SBT_LIBRARY=oracle.disksbt, ENV=(BACKUP_DIR=/tmp)';

```

The following table illustrates the channel names and channel settings that RMAN allocates when the default device is DISK and PARALLELISM for DISK is set to 4.

Channel Name	Setting
ORA_DISK_1	FORMAT = '/tmp/backup_%U'
ORA_DISK_2	MAXPIECESIZE = 20M
ORA_DISK_3	FORMAT = '/tmp/backup_%U'
ORA_DISK_4	MAXPIECESIZE = 40M

The following table illustrates the channel names and channel settings that RMAN allocates when the default device is sbt and PARALLELISM for sbt is set to 3.

Channel Name	Setting
ORA_SBT_TAPE_1	PARMS=' ENV=(BACKUP_DIR=?/oradata) '
ORA_SBT_TAPE_2	PARMS=' ENV=(BACKUP_DIR=?/oradata) '
ORA_SBT_TAPE_3	PARMS=' ENV=(BACKUP_DIR=/tmp) '

Relationship Between CONFIGURE CHANNEL and Parallelism Setting

The PARALLELISM setting is not constrained by the number of specifically configured channels. For example, if you back up to 20 different tape devices, then you can configure 20 different `sbt` channels, each with a manually assigned number (from 1 to 20) and each with a different set of channel options. In such a situation, you can set PARALLELISM to any value up to the number of devices, in this instance 20.

RMAN always numbers parallel channels starting with 1 and ending with the PARALLELISM setting. For example, if the default device is `sbt` and PARALLELISM for `sbt` is set to 3, then RMAN names the channels as follows:

```
ORA_SBT_TAPE_1
ORA_SBT_TAPE_2
ORA_SBT_TAPE_3
```

RMAN always uses the name `ORA_SBT_TAPE_n` even if you configure DEVICE TYPE `sbt` (not the synonymous `sbt_tape`). RMAN always allocates the number of channels specified in PARALLELISM, using specifically configured channels if you have configured them and generic channels if you have not.

See Also: ["Automatic Channel-Specific Configurations"](#) on page 2-6 for concepts about manually numbered channels, and ["Configuring Specific Channels: Examples"](#) on page 5-12

Clearing Channel and Device Settings

To clear a configuration is to return it to its default settings. You can clear channel and device settings by using these commands:

- `CONFIGURE DEVICE TYPE . . . CLEAR`
- `CONFIGURE DEFAULT DEVICE TYPE CLEAR`
- `CONFIGURE CHANNEL DEVICE TYPE . . . CLEAR`
- `CONFIGURE CHANNEL n DEVICE TYPE . . . CLEAR` (where *n* is an integer)

Each `CONFIGURE . . . CLEAR` command clears only itself. For example, `CONFIGURE DEVICE TYPE . . . CLEAR` does not clear `CONFIGURE DEFAULT DEVICE TYPE`. The `CONFIGURE DEVICE TYPE . . . CLEAR` command removes the configuration for the specified device type and returns it to the default (PARALLELISM 1).

Note: You cannot specify any other options when clearing a device type.

The `CONFIGURE DEFAULT DEVICE TYPE . . . CLEAR` command clears the configured default device and returns it to `DISK` (the default setting).

The `CONFIGURE CHANNEL DEVICE TYPE . . . CLEAR` command erases the channel configuration for the specified device type. RMAN does not change the PARALLELISM setting for the device type because PARALLELISM is specified through a separate `CONFIGURE` command.

If you have manually assigned options to configured channels, then clear the options for these channels individually by specifying the channel number in `CONFIGURE CHANNEL n DEVICE TYPE . . . CLEAR`. For example, assume that you run the following:

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE = 1800K;
RMAN> CONFIGURE CHANNEL 3 DEVICE TYPE DISK FORMAT = /tmp/%U;
```



```

RMAN> CONFIGURE CHANNEL 3 DEVICE TYPE DISK CLEAR;

```

In this case, RMAN clears the settings for CHANNEL 3, but leaves the settings for the generic DISK channel (the channel with no number manually assigned) intact.

See Also: ["Clearing Automatic Channel Settings"](#) on page 2-6

Configuring the Maximum Size of Backup Sets and Pieces

The CONFIGURE MAXSETSIZE command limits the size of backup sets created on a channel. This CONFIGURE setting applies to any channel, whether manually allocated or configured, when the BACKUP command is used to create backup sets.

You can set MAXSETSIZE in bytes (default), kilobytes (K), megabytes (M), and gigabytes (G). The default value is given in bytes and is rounded down to the lowest kilobyte value. For example, if you set the maximum set size to 2000, then RMAN rounds down this value to 1 kilobyte (1024 bytes). If you set the maximum set size to 2049, then RMAN rounds down this value to 2 kilobytes (2048 bytes).

The value set by the CONFIGURE MAXSETSIZE command is a default for the given channel. You can override the configured MAXSETSIZE value by specifying a MAXSETSIZE option for an individual BACKUP command.

Assume that you issue the following commands at the RMAN prompt:

```

CONFIGURE DEFAULT DEVICE TYPE TO sbt;
CONFIGURE CHANNEL DEVICE TYPE sbt PARMS 'ENV=(NSR_DATA_VOLUME_POOL=first_pool)';
CONFIGURE MAXSETSIZE TO 7500K;
BACKUP TABLESPACE users;
BACKUP TABLESPACE tools MAXSETSIZE 5G;

```

The results will be as follows:

- The backup of the `users` tablespace uses the configured `sbt` channel and the configured default MAXSETSIZE setting of 7500K.
- The backup of the `tools` tablespace uses the MAXSETSIZE setting of 5G used in the BACKUP command.

Note: There is no equivalent to MAXSETSIZE for controlling the size of image copies. Since an image copy is an exact duplicate of the file being backed up, its size must be identical to the source file.

This fact can present a problem with some older operating systems which limit the size of individual files. If you are using a raw partition to store a 10GB datafile, and your operating system only supports 4GB files on the file system, you cannot take image copy backups of that file.

See Also: *Oracle Database Backup and Recovery Reference* for BACKUP syntax

Showing the Default Maximum Size of Backup Sets: SHOW MAXSETSIZE

You can use SHOW MAXSETSIZE to view the maximum backup set size set using CONFIGURE MAXSETSIZE. The size of a backup set is measured in the total bytes of the included backup pieces. After connecting to the target database and recovery catalog (if you use one), issue the SHOW MAXSETSIZE command. For example, enter:

```
SHOW MAXSETSIZE;          # shows the CONFIGURE MAXSETSIZE settings
```

Sample output for SHOW MAXSETSIZE follows:

```
RMAN configuration parameters are:  
CONFIGURE MAXSETSIZE TO 3072K;
```

Configuring Backup Optimization

Run the CONFIGURE command to enable and disable backup optimization. Backup optimization skips the backup of files in certain circumstances if the identical file or an identical version of the file has already been backed up. Full details on the backup optimization algorithm are provided in ["Backup Optimization"](#) on page 2-12.

Note that backup optimization applies only to the following commands:

- BACKUP DATABASE
- BACKUP ARCHIVELOG with ALL or LIKE options
- BACKUP BACKUPSET ALL

You can override optimization at any time by specifying the FORCE option on the BACKUP command. For example, you can run:

```
BACKUP DATABASE FORCE;  
BACKUP ARCHIVELOG ALL FORCE;
```

By default, backup optimization is configured to OFF. To enable backup optimization, run the following command:

```
CONFIGURE BACKUP OPTIMIZATION ON;
```

To disable backup optimization, run the following command:

```
CONFIGURE BACKUP OPTIMIZATION OFF;
```

To clear the current backup optimization setting, that is, return backup optimization to its default setting of OFF, run this command:

```
CONFIGURE BACKUP OPTIMIZATION CLEAR;
```

See Also:

- ["Backup Optimization Algorithm"](#) on page 2-35 for the complete criteria that determine whether a file is identical and the conditions under which backup optimization is operative
- ["Backing Up Files Using Backup Optimization"](#) on page 6-11 for examples of how to optimize RMAN backups

Displaying Backup Optimization Setting: SHOW BACKUP OPTIMIZATION

You can use SHOW BACKUP OPTIMIZATION to view the current settings of backup optimization as configured with the CONFIGURE BACKUP OPTIMIZATION command. After connecting to the target database and recovery catalog (if you use one), issue the SHOW BACKUP OPTIMIZATION command. For example, enter:

```
SHOW BACKUP OPTIMIZATION;
```

Sample output for SHOW BACKUP OPTIMIZATION follows:

```
RMAN configuration parameters are:
```

```
CONFIGURE BACKUP OPTIMIZATION ON;
```

Configuring Backup Duplexing: CONFIGURE... BACKUP COPIES

Use the `CONFIGURE . . . BACKUP COPIES` command to specify how many copies of each backup piece should be created on the specified device type for the specified type of file. This feature is known as **duplexing**. The `CONFIGURE` settings for duplexing only affect backups of datafiles, control files and archived logs into backup sets, and do not affect image copies.

Note: Control file autobackups on disk are a special case and are *never* duplexed: RMAN always creates one and only one copy.

To configure the number of backup set copies, specify an integer. The following examples show possible configurations:

```
# Makes 2 disk copies of each datafile and control file backup set
# (autobackups excluded)
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
# Makes 3 copies of every archived redo log backup to tape
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE sbt TO 3;
```

If you use the duplexing feature in conjunction with multiple `FORMAT` strings, then you can name each individual backup set copy. For example, assume that you configure `BACKUP COPIES` to 3. Then, you can issue:

```
BACKUP DATABASE FORMAT '/tmp/%U', '?/dbs/%U', '?/oradata/%U';
```

RMAN generates 3 identical copies of each backup piece in the backup set, and names each piece according to the specified `FORMAT` string: the first copy is placed in the `/tmp` directory, the second in the `?/dbs` directory, and the third in the `?/oradata` directory. Note that you can specify the `FORMAT` string on the `BACKUP`, `CONFIGURE CHANNEL`, and `ALLOCATE CHANNEL` commands.

To return a `BACKUP COPIES` configuration to its default value, run the same `CONFIGURE` command with the `CLEAR` option, as in this example:

```
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE sbt CLEAR;
```

By default, `CONFIGURE . . . BACKUP COPIES` is set to 1 for each device type.

Note: If you do not want to create a persistent copies configuration, then you can specify copies with the `BACKUP COPIES` and `SET BACKUP COPIES` commands.

See Also:

- ["Manual Parallelization of Backups"](#) on page 2-13 for concepts
- *Oracle Database Backup and Recovery Reference* for `BACKUP` syntax
- *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax
- *Oracle Database Backup and Recovery Reference* for `SET` syntax

Showing the Configured Degree of Duplexing: SHOW... BACKUP COPIES

SHOW... BACKUP COPIES lets you view how you have used CONFIGURE... BACKUP COPIES command to set the number of identical copies that RMAN makes of each of several types of backup.

After connecting to the target database and recovery catalog (if you use one), run the SHOW ARCHIVELOG BACKUP COPIES or SHOW DATAFILE BACKUP COPIES commands. For example, enter:

```
SHOW DATAFILE BACKUP COPIES;      # shows CONFIGURE DATAFILE BACKUP COPIES setting
```

Sample output for SHOW DATAFILE BACKUP COPIES follows:

```
RMAN configuration parameters are:  
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE SBT_TAPE TO 1; # default  
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
```

Configuring Tablespaces for Exclusion from Whole Database Backups

You can run CONFIGURE EXCLUDE FOR TABLESPACE to exempt the specified tablespace from the BACKUP DATABASE command. The exclusion condition applies to any datafiles that you add to this tablespace in the future.

This tablespace exclusion feature is useful when you do not want to make a specified tablespace part of the regular backup schedule, as in these cases:

- A tablespace is easy to rebuild, so it is more cost-effective to rebuild it than back it up every day.
- A tablespace contains temporary or test data that you do not need to back up.
- A tablespace does not change often and therefore should be backed up on a different schedule from other backups.

For example, you can exclude testing tablespaces `cwmlite` and `example` from whole database backups as follows:

```
CONFIGURE EXCLUDE FOR TABLESPACE cwmlite;  
CONFIGURE EXCLUDE FOR TABLESPACE example;
```

If you run the following command, then RMAN backs up all tablespaces in the database except `cwmlite` and `example`:

```
BACKUP DATABASE;
```

You can still back up the configured tablespaces by explicitly specifying them in a BACKUP command or by specifying the NOEXCLUDE option on a BACKUP DATABASE command. For example, you can enter one of the following commands:

```
# backs up the whole database, including cwmlite and example  
BACKUP DATABASE NOEXCLUDE;  
BACKUP TABLESPACE cwmlite, example; # backs up only cwmlite and example
```

You can disable the exclusion feature for `cwmlite` and `example` as follows:

```
CONFIGURE EXCLUDE FOR TABLESPACE cwmlite CLEAR;  
CONFIGURE EXCLUDE FOR TABLESPACE example CLEAR;
```

RMAN includes these tablespaces in future whole database backups.

Showing the Tablespaces Excluded from Backups

`SHOW EXCLUDE` shows how you have used the `CONFIGURE EXCLUDE` command to exclude tablespaces from whole database backups.

After connecting to the target database and recovery catalog (if you use one), run the `SHOW EXCLUDE` command. For example, enter:

```
RMAN> SHOW EXCLUDE;      # shows the CONFIGURE EXCLUDE setting
```

Sample output for `SHOW EXCLUDE` follows:

```
RMAN configuration parameters are:
CONFIGURE EXCLUDE FOR TABLESPACE 'OLD_ACCOUNTS';
```

See Also:

- *Oracle Database Backup and Recovery Reference* for `BACKUP` syntax
- *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax

Configuring Auxiliary Instance Datafile Names: CONFIGURE AUXNAME

When performing tablespace point-in-time recovery (TSPITR) or duplicating a database using RMAN, you may want to set the names of datafiles in the auxiliary instance before starting the actual TSPITR or database duplication.

The command for doing so is:

```
CONFIGURE AUXNAME FOR datafileSpec TO 'filename';
```

where *datafileSpec* identifies some datafile by its original name or datafile number, and *filename* is the new path for the specified file.

For example, you might configure a new auxiliary name for datafile 2 as follows:

```
CONFIGURE AUXNAME FOR DATAFILE 2 TO '/newdisk/datafiles/df2.df';
```

As with other settings, this `CONFIGURE` setting is persistent across RMAN sessions until cleared using `CONFIGURE... CLEAR`, as shown here:

```
CONFIGURE AUXNAME FOR DATAFILE 2 CLEAR;
```

If you are performing TSPITR or running the `DUPLICATE` command, then by using `CONFIGURE AUXNAME` you can preconfigure the filenames for use on the auxiliary database without manually specifying the auxiliary filenames during the procedure.

When renaming files with the `DUPLICATE` command, `CONFIGURE AUXNAME` is an alternative to `SET NEWNAME`. The difference is that after you set the `AUXNAME` the first time, you do not need to reset the filename when you issue another `DUPLICATE` command: the `AUXNAME` setting remains in effect until you issue `CONFIGURE AUXNAME ... CLEAR`. In contrast, you must reissue the `SET NEWNAME` command every time you rename files.

See [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#) for more details on using `CONFIGURE AUXNAME` in connection with TSPITR, and [Chapter 13, "Creating and Updating Duplicate Databases with RMAN"](#) for more on using `CONFIGURE AUXNAME` in performing database duplication.

Showing the Default Filenames Configured for Auxiliary Channels

To view auxiliary datafile names currently configured for your database, you can use the `SHOW AUXNAME` command:

```
RMAN> SHOW AUXNAME;
```

Sample output follows:

```
RMAN configuration parameters are:
CONFIGURE AUXNAME FOR DATAFILE '/oracle/oradata/trgt/tools01.dbf' TO
'/tmp/tools01.dbf';
```

Setting the Snapshot Control File Location

When RMAN needs to resynchronize from a read-consistent version of the control file, it creates a temporary **snapshot control file**. RMAN needs a snapshot control file only when resynchronizing with the recovery catalog or when making a backup of the current control file.

The default value for the snapshot control file is platform-specific and depends on the Oracle home. For example, the default filename on some UNIX platforms in Oracle Database 10g is `$ORACLE_HOME/dbs/snapcf_@.f`. Note that if you have a flash recovery area configured, the default location for the snapshot control file is *not* the flash recovery area.

In general, you should only need to set the snapshot control file location when you are upgrading to the current release from a release earlier than 8.1.7. In these earlier releases, the default location for the snapshot control file was not dependent on the Oracle home, whereas in the current release the default location is dependent on the Oracle home.

Default Location of the Snapshot Control File

By default, the location of the snapshot control file is determined by the rules in the following table:

If you ...	Then ...
Create a new database in the current release	The snapshot control file location uses the default value. In this case, the default snapshot control file location changes if you change the Oracle home.
Upgrade to the current release from a release prior to 8.1.7	The snapshot control file location is not set to the default value. Instead, RMAN uses the snapshot location that is already stored in the control file. In this case, the snapshot control file location does not change if you change the Oracle home.

Viewing the Configured Location of the Snapshot Control File

You can see the current snapshot location by running the `SHOW` command. This example shows a snapshot location that is determined by the default rule:

```
RMAN> SHOW SNAPSHOT CONTROLFILE NAME;
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/dbs/snapcf_trgt.f'; # default
```

This example shows a snapshot control file that has a nondefault filename:

```
RMAN> SHOW SNAPSHOT CONTROLFILE NAME;
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/oradata/trgt/snap_trgt.ctl';
```

Setting the Location of the Snapshot Control File

Use the `CONFIGURE SNAPSHOT CONTROLFILE NAME TO 'filename'` command to change the name of the snapshot control file. Subsequent snapshot control files that RMAN creates use the specified filename.

For example, start RMAN and then enter:

```
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/oradata/trgt/snap_trgt.ctl';
```

You can also set the snapshot control file name to a raw device:

```
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/dev/vgd_1_0/rlvt5';
```

If one RMAN job is already backing up the control file while another needs to create a new snapshot control file, you may see the following message:

```
waiting for snapshot control file enqueue
```

Under normal circumstances, a job that must wait for the control file enqueue waits for a brief interval and then successfully retrieves the enqueue. Recovery Manager makes up to five attempts to get the enqueue and then fails the job. The conflict is usually caused when two jobs are both backing up the control file, and the job that first starts backing up the control file waits for service from the media manager.

See Also: ["Backup Fails Because of Control File Enqueue: Scenario"](#) on page 12-18, *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide* for handling of snapshot control files in RAC configurations, and *Oracle Database Backup and Recovery Reference* for `CONFIGURE` syntax

To reset the snapshot control file location to the default, run the `CONFIGURE SNAPSHOT CONTROLFILE LOCATION CLEAR` command.

Showing the Current Snapshot Control File Name

Issue the `SHOW SNAPSHOT CONTROLFILE` command to display the value set by `CONFIGURE SNAPSHOT CONTROLFILE NAME`.

Note: In releases prior to Oracle9i, the `CONFIGURE SNAPSHOT CONTROLFILE` command was called `SET SNAPSHOT CONTROLFILE`.

To show the snapshot control file filename:

After connecting to the target database and recovery catalog (if you use one), run the `SHOW SNAPSHOT CONTROLFILE` command. For example, enter:

```
SHOW SNAPSHOT CONTROLFILE NAME; # shows CONFIGURE SNAPSHOT CONTROLFILE setting
```

Sample output for `SHOW SNAPSHOT CONTROLFILE` follows:

```
RMAN configuration parameters are:
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/dbs/cf_snap.f';
```

See Also: ["Setting the Snapshot Control File Location"](#) on page 5-20 to learn about the snapshot control file and its function

Setting Up RMAN for Use with a Shared Server

RMAN cannot connect to the target database through a shared server dispatcher. RMAN requires a dedicated server process. Nevertheless, you can connect specified sessions to dedicated servers, even when the target is configured for a shared server.

To ensure that RMAN does not connect to a dispatcher when the target database is configured for a shared server, the net service name used by RMAN must include `(SERVER=DEDICATED)` in the `CONNECT_DATA` attribute of the connect string.

Oracle Net configuration varies greatly from system to system. The following procedure illustrates only one method. This scenario assumes that the following service name in the `tnsnames.ora` connects to the target database using the shared server architecture, where `inst1` is a value of the `SERVICE_NAMES` initialization parameter:

```
inst1_shs =
  (DESCRIPTION=
    (ADDRESS=(PROTOCOL=tcp) (HOST=inst1_host) (port1521))
    (CONNECT_DATA=(SERVICE_NAME=inst1) (SERVER=shared))
  )
```

To use RMAN with a shared server:

1. Create a net service name in the `tnsnames.ora` file that connects to the non-shared SID. For example, enter:

```
inst1_ded =
  (DESCRIPTION=
    (ADDRESS=(PROTOCOL=tcp) (HOST=inst1_host) (port1521))
    (CONNECT_DATA=(SERVICE_NAME=inst1) (SERVER=dedicated))
  )
```

2. Start SQL*Plus and then connect using both the shared server and dedicated server service names to confirm the mode of each session. For example, to connect to a dedicated session you can issue:

```
CONNECT SYS/oracle@inst1_ded
SELECT SERVER
  FROM V$SESSION
  WHERE SID = (SELECT DISTINCT SID FROM V$MYSTAT);
```

```
SERVER
-----
DEDICATED
1 row selected.
```

To connect to a shared server session, you can issue:

```
CONNECT SYS/oracle@inst1_shs AS SYSDBA
SELECT SERVER
  FROM V$SESSION
  WHERE SID = (SELECT DISTINCT SID FROM V$MYSTAT);
```

```
SERVER
-----
SHARED
1 row selected.
```

3. Connect to the target database (and optionally the recovery catalog) with the dedicated service name. For example, enter:


```
% rman TARGET SYS/oracle@inst1_ded CATALOG rman/cat@catdb
```

See Also: Your platform-specific Oracle documentation and your *Oracle Database Net Services Reference* for a complete description of Oracle Net connect string syntax

Making Backups with RMAN: Advanced Topics

This chapter describes how to use RMAN to make backups. This chapter contains these topics:

- [Configuring and Allocating Channels for Use in Backups](#)
- [Making Split Mirror Backups with RMAN](#)
- [Backing Up Backup Sets with RMAN](#)
- [Backing Up Existing Image Copy Backups with RMAN](#)
- [RMAN Encrypted Backups](#)
- [Restarting and Optimizing RMAN Backups](#)
- [Validating Backups with RMAN](#)
- [RMAN Backup Examples](#)

Configuring and Allocating Channels for Use in Backups

You have the following options for executing backups:

- Configure automatic channels with the `CONFIGURE` command, and then issue `BACKUP` commands at the RMAN prompt or within a `RUN` block
- Within a `RUN` block, allocate channels manually with the `ALLOCATE CHANNEL` command, and then issue `BACKUP` commands using those channels

The easiest way to make backups is to configure automatic channels. For example, so long as you have already configured an `sbt` device type, you can configure a default `sbt` channel as follows (note that the `PARMS` value is vendor-specific) and then back up the database using these defaults:

```
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
RMAN> CONFIGURE CHANNEL DEVICE TYPE sbt PARMS 'ENV=(NSR_SERVER=bksvr1)';
RMAN> BACKUP DATABASE;
```

RMAN preconfigures a `DISK` channel for you, so you can make disk backups using automatic channels without performing any configuration whatsoever.

The other method is to allocate channels manually within a `RUN` block before running the `BACKUP` command. For example, this command allocates multiple disk channels and then backs up the database and archived redo logs:

```
RMAN> RUN
```

```
{
  ALLOCATE CHANNEL ch1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL ch2 DEVICE TYPE DISK;
  ALLOCATE CHANNEL ch3 DEVICE TYPE DISK;
  BACKUP DATABASE PLUS ARCHIVELOG;
}
```

The following example manually allocates an `sbt` channel (with a vendor-specific `PARMS` value) and backs up a datafile copy:

```
RMAN> RUN
{
  ALLOCATE CHANNEL ch1 DEVICE TYPE sbt PARMS 'ENV=(NSR_SERVER=bksvr1)';
  BACKUP DATAFILECOPY '/tmp/system01.dbf';
}
```

Most examples in this chapter assume that you have configured automatic channels.

Duplexing Backup Sets

RMAN can make up to four copies of a backup set simultaneously, each an exact duplicate of the others. A copy of a backup set is a copy of each backup piece in the backup set, with each copy getting a unique copy number (for example, `0tcm8u2s_1_1` and `0tcm8u2s_1_2`).

In most cases, the easiest method of duplexing backup sets is to use `BACKUP . . . COPIES` or `CONFIGURE . . . BACKUP COPIES` to duplex backup sets. For `DISK` channels, specify multiple values in the `FORMAT` option to direct the multiple copies to different physical disks. For `sbt` channels, if you use a media manager that supports Version 2 of the SBT API, then the media manager will automatically put each copy onto a separate medium (for example, a separate tape).

Note that it is not possible to duplex backup sets to the flash recovery area, and that duplexing only applies to backup sets, not image copies. It is an error to specify the `BACKUP . . . COPIES` when creating image copy backups, and the `CONFIGURE . . . BACKUP COPIES` setting is ignored for image copy backups.

Duplexing Backup Sets with CONFIGURE BACKUP COPIES

The `CONFIGURE . . . BACKUP COPIES` command specifies the number of identical backup sets that you want to create on the specified device type. This setting applies to all backups except control file autobackups (because the autobackup of a control file always produces one copy) and backup sets when backed up with the `BACKUP BACKUPSET` command. You must have automatic channels configured.

To duplex a backup with CONFIGURE BACKUP COPIES:

1. Configure the number of copies on the desired device type for datafiles and archived redo logs on the desired device types. This example configures duplexing for datafiles and archived logs on tape as well as duplexing for datafiles (but not archived logs) on disk:

```
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/save1/%U', '/save2/%U';
RMAN> CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE sbt TO 2;
RMAN> CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE sbt TO 2;
RMAN> CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 2;
```

2. Execute the `BACKUP` command. The following command backs up the database and archived logs to tape, making two copies of each datafile and archived redo log:

```

RMAN> BACKUP DATABASE PLUS ARCHIVELOG; # uses default sbt channel

```

Because of the configured formats for the disk channel, the following command backs up the database to disk, placing one copy of the backupsets produced in the `/save1` directory and the other in the `/save2` directory:

```

RMAN> BACKUP DEVICE TYPE DISK AS COPY DATABASE;

```

3. Issue a `LIST BACKUP` command to see a listing of backup sets and pieces. For example, enter:

```

RMAN> LIST BACKUP SUMMARY;

```

The `#Copies` column shows the number of backupsets, which may have been produced by duplexing or by multiple backup commands.

Duplexing Backupsets with `BACKUP... COPIES`

The `COPIES` option of the `BACKUP` command overrides every other `COPIES` or `DUPLEX` setting to control duplexing of backupsets.

To duplex a backup with `BACKUP COPIES`:

1. Specify the number of identical copies with the `COPIES` option of the `BACKUP` command. For example, run the following to make three copies of each backup set in the default `DISK` location:

```

RMAN>
BACKUP AS BACKUPSET DEVICE TYPE DISK
  COPIES 3
  INCREMENTAL LEVEL 0
  DATABASE;

```

Because you specified `COPIES` on the `BACKUP` command, RMAN makes three backupsets of each datafile regardless of the `CONFIGURE DATAFILE COPIES` setting.

2. Issue a `LIST BACKUP` command to see a listing of backup sets and pieces (the `#Copies` column shows the number of copies, which may have been produced through duplexing or through multiple invocations of the `BACKUP` command). For example, enter:

```

RMAN>
LIST BACKUP SUMMARY;

```

Making Split Mirror Backups with RMAN

Many sites keep an backup of the database stored on disk in case a failure occurs on the primary database or an incorrect user action such as a `DROP TABLE` requires point-in-time recovery. A datafile backup on disk simplifies the restore step of recovery, making recovery much quicker and more reliable.

Caution: Never make backups, split mirror or otherwise, of online redo logs. Restoring online redo log backups can create two archived logs with the same sequence number but different contents. Also, it is best to use the `BACKUP CONTROLFILE` command rather than a split mirror to make control file backups.

One way of creating a datafile backup on disk is to use disk mirroring. For example, you can use the operating system to maintain three identical copies of each file in the database. In this configuration, you can split off a mirrored copy of the database to use as a backup.

RMAN does not automate the splitting of mirrors, but can make use of split mirrors in backup and recovery operations. For example, RMAN can treat a split mirror of a datafile as a datafile copy, and can also back up this copy to disk or tape.

The following procedure shows how to make a split mirror backup with the `SUSPEND/RESUME` functionality. The `SUSPEND/RESUME` feature is not required for split mirror backups in most cases, although it is necessary if your system requires the database cache to be free of dirty buffers before the volume can be split.

To make a split mirror backup of a tablespace by using `SUSPEND/RESUME`:

1. Start RMAN and then place the tablespaces that you want to back up into backup mode with the `ALTER TABLESPACE . . . BEGIN BACKUP` statement. (To place all tablespaces in backup mode, you can use `ALTER DATABASE BEGIN BACKUP` instead.)

For example, to place tablespace `users` in backup mode, start RMAN and run the following commands:

```
RMAN> CONNECT TARGET SYS/oracle@trgt
RMAN> CONNECT CATALOG rman/cat@catdb

RMAN> SQL 'ALTER TABLESPACE users BEGIN BACKUP';
```

2. Suspend the I/Os if your mirroring software or hardware requires it. For example, enter the following SQL statement:

```
RMAN> SQL 'ALTER SYSTEM SUSPEND';
```

3. Split the mirrors for the underlying datafiles contained in these tablespaces.
4. Take the database out of the suspended state:

```
RMAN> SQL 'ALTER SYSTEM RESUME';
```

5. Take the tablespaces out of backup mode. For example, enter:

```
RMAN> SQL 'ALTER TABLESPACE users END BACKUP';
```

You could also use `ALTER DATABASE END BACKUP` to take all tablespaces out of backup mode.

6. Start an RMAN session and then catalog the user-managed mirror copies as datafile copies with the `CATALOG` command. For example, enter:

```
RMAN> CATALOG DATAFILECOPY '/dk2/oradata/trgt/users01.dbf'; # catalog split
mirror
```

7. Back up the datafile copies. For example, assuming that you have configured automatic channels, run the `BACKUP DATAFILECOPY` command at the prompt:

```

RMAN> BACKUP DATAFILECOPY '/dk2/oradata/trgt/users01.dbf';

```

8. When you are ready to resilver the split mirror, first use the `CHANGE . . . UNCATALOG` command to uncatalog the datafile copies you cataloged in step 6. For example, enter:

```

RMAN> CHANGE DATAFILECOPY '/dk2/oradata/trgt/users01.dbf' UNCATALOG;

```

9. Resilver the split mirror for the affected datafiles.

See Also: *Oracle Database SQL Reference* for `ALTER SYSTEM SUSPEND` syntax

Backing Up Backup Sets with RMAN

Use the `BACKUP BACKUPSET` command to back up backup sets produced by other backup jobs. This command is especially useful in the following scenarios:

- Ensuring that all backups exist both on disk and on tape
- Moving backups from disk to tape and then freeing the space on disk

Note: You cannot duplex backups when running `BACKUP BACKUPSET`. RMAN always makes one and only one copy on the specified media when performing `BACKUP BACKUPSET`.

To back up backup sets from disk to tape:

1. Assuming that you have configured an automatic `sbt` channel, issue the `BACKUP BACKUPSET` command at the RMAN prompt. This example backs up all disk backup sets to tape:

```

RMAN> BACKUP DEVICE TYPE sbt BACKUPSET ALL;

```

This example backs up all disk backup sets to tape and then deletes the input disk backups:

```

RMAN>
BACKUP DEVICE TYPE sbt BACKUPSET ALL DELETE INPUT;

```

2. Issue a `LIST` command to see a listing of backup sets and pieces.

BACKUP BACKUPSET and Backup Retention Policy

`BACKUP BACKUPSET` creates additional copies of backup pieces in a backup set by copying backup pieces from disk to tape. If you use the `LIST BACKUPS` command, the output contains the list of all copies (including copies created by `BACKUP BACKUPSET` command).

However, `BACKUP BACKUPSET` does not create a new backup set. In this way, `BACKUP BACKUPSET` is similar to using the `DUPLEX` or `MAXCOPIES` option of `BACKUP`. The extra copy of a backup set created by `BACKUP BACKUPSET` is not considered to be a new backup set, as the several copies of a backup set produced by these other forms of the `BACKUP` command are not separate backup sets.

For the purposes of a redundancy-based backup retention policy, a backup set is counted as one instance of backup, even if there are multiple copies of the backup pieces that make up backup set, such as when a backup set has been backed up from

disk to tape. For the purposes of a recovery window retention policy, either all of the copies of a backup set are obsolete, or none of them are.

This is easier to understand if you look at the output of the LIST and REPORT commands. For example, perform the following backup:

```
RMAN> backup as backupset datafile 5
RMAN> backup backupset <previous backupset>;
```

Now, run the LIST command. The output contains the following:

Notice that the set_stamp and set_count values remain the same, but the copy# is incremented for the new backup. To see the effect of these copies under a redundancy-based backup retention policy, use the following command:

```
report obsolete redundancy 1;
```

None of the copies is reported as obsolete because both copies of the backup set have the same values for set_stamp & set_count.

To see the effect of these copies under a recovery window-based retention policy, use the following command:

```
report obsolete recovery window 1 day;
```

None of the copies of the backup set is reported as obsolete or based on the checkpoint_change# of this backupset, with respect to the current time and the availability of other backups.

Backing Up Existing Image Copy Backups with RMAN

Use the BACKUP COPY OF command to back up image copies of datafiles, control files, and archived logs. The output of the BACKUP COPY OF command can be either backup sets or image copies, so you can generate backup sets from your image copies. This form of backup is used to back up a database backup created as image copies on disk to tape. You can use the MAXSETSIZE parameter of the BACKUP command to set a maximum size for each backup set.

Assuming that you have configured an automatic sbt channel, issue this command to back up the latest image copy backup of the database to tape:

```
RMAN>
BACKUP DEVICE TYPE sbt COPY OF DATABASE;
```

This example backs up the latest image copy backup of a database to backup sets on tape, and then deletes the input disk backups:

```
RMAN>
BACKUP DEVICE TYPE sbt COPY OF DATABASE DELETE INPUT;
```

When backing up your image copies, identifying the image copy to back up is easier if you provide tags for your backups. Image copies of datafiles have associated tags (if you do not provide one, one is automatically generated). The tag of an image copy is inherited by default when the image copy is backed up as a new image copy. You can also specify your own tag.

To take advantage of the tags associated with your image copy backups, use the WITH TAG selector. The new backup will be tagged with the same tag as the image copy backup being backed up. When multiple image copies have the same tag, the most recent image copy of a file with the specified tag will be backed up.

RMAN Encrypted Backups

For improved security, RMAN backups created as backup sets can be encrypted. Encrypted backups cannot be read if they are obtained by unauthorized people.

Any RMAN backups as backup sets can be encrypted. However, image copy backups cannot be encrypted.

Encrypted backups are decrypted automatically during restore and recover operations, as long as the required decryption keys are available, by means of either a user-supplied password or the Oracle Encryption Wallet.

To use RMAN encryption, the COMPATIBLE initialization parameter at the target database must be set to at least 10.2.0.

When the `BACKUP BACKUPSET` command is used with encrypted backup sets, the backup sets are backed up in their encrypted form. Because `BACKUP BACKUPSET` just copies an already-encrypted backup set to disk or tape, no decryption key is needed during a `BACKUP BACKUPSET` operation, and the data is never decrypted during any part of the operation. The `BACKUP BACKUPSET` command can neither encrypt nor decrypt backup sets.

If some columns in the database are encrypted using Transparent Data Encryption, and those columns are backed up using backup encryption, then those columns will be encrypted a second time during the backup. When the backup sets are decrypted during a restore, the encrypted columns are returned to their original encrypted form.

The `V$RMAN_ENCRYPTION_ALGORITHMS` view contains a list of encryption algorithms supported by RMAN. If no encryption algorithm is specified, the default encryption algorithm is 128-bit AES.

RMAN Backup Encryption Modes

RMAN offers three encryption modes: transparent mode, password mode, and dual mode.

Note: Both transparent mode and dual mode depend upon the Oracle Encryption Wallet. See *Oracle Advanced Security Administrator's Guide* for details about configuring the Oracle Encryption Wallet.

Transparent Encryption of Backups

Transparent encryption can create and restore encrypted backups with no DBA intervention, as long as the required Oracle key management infrastructure is available. Transparent encryption is best suited for day-to-day backup operations, where backups will be restored at the same database that they were backed up from. Transparent encryption is the default mode for RMAN encryption.

When using transparent encryption, you must first configure the Oracle Encryption Wallet, as described in the documentation for Oracle's Transparent Data Encryption feature. After the Oracle Encryption Wallet is configured, encrypted backups can be created and restored with no further DBA intervention.

Note: Because the Oracle key management infrastructure archives all previous master keys in the Oracle Encryption Wallet, changing or resetting the current database master key will not affect your ability to restore encrypted backups performed using an older master key. You may reset the database master key at any time, and RMAN will always be able to restore all encrypted backups that were ever created by this database.

Transparent backup encryption supports both the encrypted and autologin forms of the Oracle Encryption Wallet. When using the encrypted wallet, the wallet must be opened before any backup encryption operations, either backups or restores, can be done. When using the autologin wallet, encrypted backup operations can be done at any time, because the autologin wallet is always open.

Caution: If you use an autologin wallet, do not back up the autologin wallet along with your encrypted backup data, because anybody will be able to read the encrypted backups if they obtain both the backups and the autologin wallet. It is safe to back up the encrypted wallet, because that form of the wallet cannot be used without the wallet password.

Caution: If you lose your Oracle Encryption Wallet then you will be unable to restore any transparently-encrypted backups.

Password Encryption of Backups

Password encryption requires that the DBA provide a password when creating and restoring encrypted backups. Restoring a password-encrypted backup requires the same password that was used to create the backup. Password encryption is useful for backups that will be restored at remote locations, but which must remain secure in transit. Password encryption cannot be persistently configured. The Oracle Encryption Wallet need not be configured if password encryption is to be used exclusively.

Caution: If you forget, or lose, the password that you used to encrypt a password-encrypted backup, you will be unable to restore that backup.

To use password encryption, use the `SET ENCRYPTION ON IDENTIFIED BY password ONLY` command in your RMAN scripts.

Dual Mode Encryption of Backups

Dual-mode encrypted backups can be restored either transparently or by specifying a password. Dual-mode encrypted backups are useful when you create backups that are normally restored on-site using the Oracle Encryption Wallet, but which occasionally need to be restored off-site, where the Oracle Encryption Wallet is not available.

When restoring a dual-mode encrypted backup, you can use either the Oracle Encryption Wallet or a password for decryption.

Caution: If you forget, or lose, the password that you used to encrypt a dual-mode encrypted backup and you also lose your Oracle Encryption Wallet, then you will be unable to restore that backup.

To create dual-mode encrypted backup sets, specify the `SET ENCRYPTION ON IDENTIFIED BY password` command in your RMAN scripts.

Using CONFIGURE and SET to Control RMAN Backup Encryption

The `CONFIGURE` command is used to make persistent settings controlling RMAN encryption of backups, including:

- Whether to encrypt backups of all database files.
- Whether to encrypt backups of specific tablespaces.
- Which algorithm to use for encrypting backups.

The `SET ENCRYPTION` command is used to:

- Override the encryption settings specified by the `CONFIGURE ENCRYPTION` command. For example, you can use `SET ENCRYPTION OFF` to create an unencrypted backup, even though your database is configured to create encrypted backups.
- Set a password for backup encryption, persisting until the RMAN client exits. Due to the sensitive nature of passwords, RMAN does not allow configuration of passwords that persist between RMAN sessions.

If you wish to modify your existing backup environment so that all RMAN backups are encrypted, perform the following steps:

- Set up the Oracle Encryption Wallet
- Issue the following RMAN command:

```
RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON
```

After these steps, all RMAN backup sets created by this database will be encrypted, unless you explicitly override this behavior for an RMAN session with:

```
RMAN> SET ENCRYPTION ON
```

This remains in effect until you issue the `SET ENCRYPTION OFF` command during an RMAN session, or change the persistent setting again with:

```
RMAN> CONFIGURE ENCRYPTION FOR DATABASE OFF
```

Creating Encrypted Backups

The `BACKUP` command arguments do not change for creating encrypted backups. Encryption is performed based on the encryption settings specified with `CONFIGURE ENCRYPTION` or `SET ENCRYPTION`. Any backup to backup sets is encrypted with the specified settings.

Oracle uses a new encryption key for every encrypted backup. The backup encryption key is then encrypted with either the password, the database master key, or both, depending on the chosen encryption mode. Individual backup encryption keys or passwords are never stored in the clear.

Restoring Data from Encrypted Backups

RMAN automatically decrypts backup sets when their contents are restored. Transparently encrypted backups require no intervention to restore, as long as the Oracle Encryption Wallet is open and available. Password encrypted backups require the correct password to be entered before they can be restored.

You must enter the encryption password using the `SET DECRYPTION` command. Restore passwords cannot be persistently configured. If restoring from a set of backups that were created with different passwords, then specify all of the required passwords on the `SET DECRYPTION` command. RMAN will automatically use the correct password with each backup set.

Encryption of Archived Log Backups

There is no persistent configuration that controls whether archive log backups as backupsets are encrypted. Backup sets containing archived logs are encrypted if any of the following are true:

- `SET ENCRYPTION ON` is in effect at the time that the archive log backup is being created.
- Encryption is configured for backups of the whole database or at least one tablespace.

This behavior ensures that the redo associated with any encrypted backup of a datafile is also encrypted.

Performance Impact of Encrypting RMAN Backups

Encryption can have a negative effect upon backup performance. Because encrypted backups will use more CPU resource than non-encrypted backups, you can improve performance of encrypted backups by using more RMAN channels.

Restarting and Optimizing RMAN Backups

RMAN supports two distinct features by which it can back up only those files that require backups: **restartable backups** and **backup optimization**.

With the restartable backup feature, RMAN backs up only those files that were not backed up after a specified date. For example, by specifying the `NOT BACKED UP SINCE TIME` clause, you can direct RMAN to back up only those files that were not backed up within the last day.

With backup optimization, the `BACKUP` command skips the backup of a file if the identical file has already been backed up to the allocated device type. To override this behavior and back up all files whether or not they have changed, specify the `FORCE` option on the `BACKUP` command. To enable or disable backup optimization, specify `ON` or `OFF` on the `CONFIGURE BACKUP OPTIMIZATION` command.

Additionally, `BACKUP . . . PLUS ARCHIVELOG` can archive unarchived online logs as well as back up archived logs.

See Also: ["Backup Optimization"](#) on page 2-35 for a conceptual overview of optimization, and ["Restartable Backups"](#) on page 2-39 for a conceptual overview of restartable backups

Backing Up Files Using Backup Optimization

For backup optimization to be enabled, you must `CONFIGURE BACKUP OPTIMIZATION` to `ON`. Backup optimization is `OFF` by default.

To use backup optimization with a backup operation:

1. If you have not already enabled backup optimization, then enable it by running the `CONFIGURE BACKUP OPTIMIZATION` command. For example, enter:

```

RMAN>
CONFIGURE BACKUP OPTIMIZATION ON;

```

2. Back up the desired files. The following example backs up to an `sbt` device any archived redo logs that either have not been already backed up, or where the existing backups do not satisfy the current duplexing setting:

```

RMAN>
BACKUP DEVICE TYPE sbt ARCHIVELOG ALL;

```

RMAN does not signal an error when it skips backing up files due to backup optimization.

See Also: ["Backup Optimization"](#) on page 2-35 for a conceptual overview of optimization and backup retention policies

Restarting a Backup After It Partially Completes

Use the `SINCE TIME` parameter of the `BACKUP` command to specify a date after which a new backup is required. If you do not specify the `SINCE` parameter, then RMAN only backs up files that have never been backed up.

To only back up files that were not backed up after a specified date:

Specify a valid date in the `SINCE TIME` parameter. For example, this command uses the default configured channel to back up all database files and archived redo logs that have not been backed up in the last two weeks:

```

RMAN> BACKUP NOT BACKED UP SINCE TIME 'SYSDATE-14'
      DATABASE PLUS ARCHIVELOG;

```

Backups to CD, DVD and Other Disk Devices with Large Block Sizes

When backing up database files to `DISK`, the logical block size of the database file to be backed up must be an even multiple of the physical block size of the destination device. For example, a `DISK` device with a block size of 2K can only be used as a destination for backups of database files with logical block sizes of 2K, 4K, 6K and so on. In practice, most disk drives have physical block sizes of 512 bytes, so this limitation rarely affects backup to disk drives. However, you can encounter this limitation when using `BACKUP ... DEVICE TYPE DISK` to back your database up to a writeable CD or DVD, or some other device which has a larger physical block size.

Validating Backups with RMAN

You can use the `VALIDATE` keyword of the `BACKUP` command to do the following:

- Check datafiles for physical and logical corruption
- Confirm that all database files exist and are in the correct locations

RMAN does not actually produce backup sets, but rather reads the specified files in their entirety, to determine whether they can be backed up and are not corrupted. In this sense, the `BACKUP... VALIDATE` command is similar to the `RESTORE... VALIDATE` command, except for backups rather than restore jobs.

If the backup validation discovers corrupt blocks, then RMAN updates the `V$DATABASE_BLOCK_CORRUPTION` view with rows describing the corruptions. After a corrupt block is repaired, the row identifying this block is deleted from the view.

For example, you can validate that all database files and archived redo logs can be backed up by running a command as follows:

```
RMAN> BACKUP VALIDATE DATABASE ARCHIVELOG ALL;
```

This form of the command would check for physical corruption. To check for logical corruption,

```
RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
```

RMAN displays the same output that it would if it were really backing up the files. If RMAN cannot validate the backup of one or more of the files, then it displays an error message. For example, RMAN may show output similar to the following:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of backup command at 08/29/2001 14:33:47
ORA-19625: error identifying file /oracle/oradata/trgt/arch/archive1_6.dbf
ORA-27037: unable to obtain file status
SVR4 Error: 2: No such file or directory
Additional information: 3
```

You cannot use the `MAXCORRUPT` or `PROXY` parameters with the `VALIDATE` option.

See Also:

- *Oracle Database Backup and Recovery Reference* for `BACKUP` syntax
- "[Block Media Recovery of Blocks Listed in V\\$DATABASE_BLOCK_CORRUPTION](#)" on page 7-14 to learn how to repair corrupt blocks discovered by `BACKUP . . . VALIDATE`

RMAN Backup Examples

This section contains these topics:

- [Skipping Tablespaces when Backing Up a Database: Example](#)
- [Restarting a Backup: Example](#)
- [Spreading a Backup Across Multiple Disk Drives: Example](#)
- [Specifying the Size of Backup Sets: Example](#)
- [Limiting the Size of Backup Pieces: Example](#)
- [Backing Up Archived Redo Logs in a Failover Scenario: Example](#)
- [Backing Up Archived Logs Needed to Recover an Online Backup: Example](#)
- [Backing Up and Deleting Multiple Copies of an Archived Redo Log: Example](#)
- [Determining How Channels Distribute a Backup Workload: Example](#)
- [Backing Up in NOARCHIVELOG Mode: Example](#)

- [Keeping a Long-Term Backup: Example](#)
- [Using Backup Optimization: Examples](#)
- [Handling Corruption During Backups: Example](#)

Skipping Tablespaces when Backing Up a Database: Example

The following example assumes that the database is running in ARCHIVELOG mode and that you have an automatic sbt channel configured as follows:

```
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
RMAN> CONFIGURE CHANNEL DEVICE TYPE sbt PARMS='ENV=(NSR_DATA_VOLUME_
POOL=BackupPool)';
```

To back up the database while skipping offline and read-only tablespaces, you can run the following command:

```
RMAN> BACKUP DATABASE
      SKIP READONLY
      SKIP OFFLINE;
```

You only need to back up a read-only tablespace once after it has been made read-only. You can use the `SKIP READONLY` option to skip read-only datafiles. If you use the `SKIP OFFLINE` option, then the `BACKUP` command does not attempt to access offline datafiles. Use this option if the offline datafiles are not available.

Another way to persistently skip tablespaces across RMAN sessions is to issue the `CONFIGURE EXCLUDE` command for each tablespace that you always want to skip. For example, you may always want to skip the `example` tablespace, which has been made read-only. You can then issue:

```
RMAN> CONFIGURE EXCLUDE FOR TABLESPACE example;
```

Then, whenever you run `BACKUP DATABASE`, RMAN skips this tablespace. You do not have to specify a `SKIP` clause on the `BACKUP` command. You can override this behavior and include the `example` tablespace as follows:

```
RMAN> BACKUP DATABASE NOEXCLUDE;
```

Restarting a Backup: Example

Assume that you back up the database and archived logs every night to tape by running this command:

```
RMAN> BACKUP
      MAXSETSIZE 10G
      DATABASE PLUS ARCHIVELOG;
```

The preceding command sets an upper limit to the size of each backup set so that RMAN produces multiple backup sets. Assume that the media management device fails halfway through the backup and is then restarted. The next day you discover that only half the backup sets completed. In this case, you can run this command in the evening:

```
RMAN> BACKUP
      # Note that the NOT BACKED UP SINCE clause should be placed immediately after
the BACKUP
      # keyword or after each individual backupSpec clause
      NOT BACKED UP SINCE TIME 'SYSDATE-1'
      MAXSETSIZE 10M
```

```
DATABASE PLUS ARCHIVELOG;
```

With this form of the command, RMAN backs up only files that were not backed up during in the previous 24 hours. When RMAN finds out that a backup from the specified time window is already available, it displays output similar to the following:

```
RMAN-06501: skipping datafile 1; already backed up on NOV 02 2003 18:10:00
RMAN-06501: skipping datafile 2; already backed up on NOV 02 2003 18:09:45
RMAN-06501: skipping datafile 3; already backed up on NOV 02 2003 18:09:45
```

If the NOT BACKED UP SINCE clause is placed immediately after the backup command, it affects all objects to be backed up. It can also be placed after individual backupSpec clauses, to cause only backups for those objects described by the backupSpec to be subject to the limitation.

Spreading a Backup Across Multiple Disk Drives: Example

Typically, you do not need to specify a format when backing up to tape because the default %U variable generates a unique filename for tape backups. When backing up to disk, however, you can specify a format if you need to spread the backup across several drives for improved performance. In this case, allocate one DISK channel for each disk drive and specify the format string on the ALLOCATE CHANNEL command so that the filenames are on different disks. For example, issue:

```
RUN
{
  ALLOCATE CHANNEL disk1 DEVICE TYPE DISK FORMAT '/disk1/%d_backups/%U';
  ALLOCATE CHANNEL disk2 DEVICE TYPE DISK FORMAT '/disk2/%d_backups/%U';
  ALLOCATE CHANNEL disk3 DEVICE TYPE DISK FORMAT '/disk3/%d_backups/%U';
  BACKUP AS COPY DATABASE;
}
```

You can distribute backups in this manner by default in the future, by configuring channels as follows:

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 3;
CONFIGURE DEFAULT DEVICE TYPE TO DISK;
CONFIGURE CHANNEL 1 DEVICE TYPE DISK FORMAT '/disk1/%d_backups/%U';
CONFIGURE CHANNEL 2 DEVICE TYPE DISK FORMAT '/disk2/%d_backups/%U';
CONFIGURE CHANNEL 3 DEVICE TYPE DISK FORMAT '/disk3/%d_backups/%U';
BACKUP AS COPY DATABASE;
```

If you specify a nonexistent directory, RMAN displays output such as the following:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03009: failure of backup command on ORA_DISK_1 channel at 08/29/2001
           14:36:04
ORA-19504: failed to create file "/nosuchdisk/0cd2momi_1_1"
ORA-27040: skgfrcre: create error, unable to create file
SVR4 Error: 2: No such file or directory
```

Specifying the Size of Backup Sets: Example

When making backups, RMAN divides the total number of files requiring backups by the number of allocated channels to calculate the number of files to place in each backup set. Use the MAXSETSIZE parameter to override this calculation and specify how many files should go in each backup set.

The `MAXSETSIZE` parameter specifies a maximum size for a backup set in units of bytes (default), kilobytes, megabytes, or gigabytes. Thus, to limit a backup set to 305 MB, specify `MAXSETSIZE=305M`. RMAN attempts to limit all sets to this size.

You can use `MAXSETSIZE` to limit the size of backup sets so that the database is divided among more than one backup set. If you configure `MAXSETSIZE` so that you generate multiple backup sets, however, then if the backup fails partway through, you can use the restartable backup feature to back up only those files that were not backed up during the previous attempt. See ["Restartable Backups"](#) on page 2-39 for a conceptual overview of restartable backups.

The following example configures a tape device, then backs up archived redo logs to tape, limiting the size to 100 MB so that if the backup fails partway through, it can be restarted:

```
RMAN> CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
RMAN> BACKUP MAXSETSIZE = 100M ARCHIVELOG ALL;
```

This example accomplishes the same result with `CONFIGURE MAXSETSIZE`:

```
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO sbt;
RMAN> CONFIGURE MAXSETSIZE = 100M;
RMAN> BACKUP ARCHIVELOG ALL;
```

Note that if you specify a `MAXSETSIZE` value that is too small to contain the biggest file that you are backing up (either the actual size of that file, or if binary compression is specified, then the size of the file after compression), then RMAN displays an error stack such as the following:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of backup command at 11/03/03 14:40:33
RMAN-06182: archive log larger than MAXSETSIZE: thread 1 seq 1
              /oracle/oradata/trgt/arch/archive1_1.dbf
```

Limiting the Size of Backup Pieces: Example

Backup piece size is an issue in those situations where it exceeds the maximum file size of the file system or media management software. Use the `MAXPIECESIZE` parameter of the `CONFIGURE CHANNEL` or `ALLOCATE CHANNEL` command to limit the size of backup pieces.

For example, to always limit the backup piece size to 2GB or less, you can configure the automatic DISK channel as follows and then run `BACKUP DATABASE`:

```
# max file size for backup pieces is 2GB
RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2GB;
RMAN> BACKUP DATABASE;
```

Note that in version 2.0 of the media management API, media management vendors can specify the maximum size of a backup piece that can be written to their media manager. RMAN will respect this limit regardless of the settings you configure for `MAXPIECESIZE`.

Backing Up Archived Redo Logs in a Failover Scenario: Example

Assume that you set your initialization parameters so that you archive to the following local destinations:

```
LOG_ARCHIVE_DEST_1 = 'LOCATION=/disk1/arch/'
LOG_ARCHIVE_DEST_2 = 'LOCATION=/disk2/arch/'
LOG_ARCHIVE_DEST_3 = 'LOCATION=/disk3/arch/'
```

Each directory contains the same set of logs, starting with log sequence 1 and ending at log sequence 400. Unknown to you, a user inadvertently deletes logs 300 through 400 from /disk1/arch and logs 350 through 400 from /disk2/arch. You run this backup command:

```
RMAN> BACKUP ARCHIVELOG
FROM SEQUENCE 288 UNTIL SEQUENCE 388
THREAD 1
DELETE INPUT;
```

RMAN begins backing up logs starting with log sequence 288. If the copy of log 300 that was deleted from /disk1/arch is the one that RMAN attempts to back up, then RMAN checks the repository to determine whether other copies of this log sequence exist, and backs up the log in either /disk2/arch or /disk3/arch. Hence, because there is at least one intact copy of each log from sequence 288 through sequence 388, RMAN can back up all the specified logs.

Backing Up Archived Logs Needed to Recover an Online Backup: Example

Assume that you back up database `trgt` while it is open. You want to back up only those archived redo logs required to recover this online backup.

The recommended solution to this problem is to add the `PLUS ARCHIVELOG` clause to your database backup command, as shown here:

```
RMAN> BACKUP DATABASE PLUS ARCHIVELOG;
```

See *Oracle Database Backup and Recovery Basics* and *Oracle Database Backup and Recovery Reference* for details on using `BACKUP . . . PLUS ARCHIVELOG`.

You can also manually determine which archived logs are required and back them up, using the following procedure.

1. Start SQL*Plus and archive all unarchived logs, including the current log:

```
SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;
```

2. Query `V$LOG` to determine the log sequence number of the current redo log, as in the following example (which includes output):

```
SQL> SELECT SEQUENCE# FROM V$LOG WHERE STATUS = 'CURRENT';
```

```
SEQUENCE#
-----
          9100
```

3. Start RMAN and make an online backup of the database. For example, enter:

```
RMAN>
BACKUP DATABASE;
```

4. Archive all unarchived logs, including the current log:

```
RMAN>
SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';
```

5. In SQL*Plus, query `V$LOG` to determine the log sequence number of the current redo log:

```
SQL> SELECT SEQUENCE# FROM V$LOG WHERE STATUS = 'CURRENT';
```

```
SEQUENCE#
-----
          9112
```

6. Back up the logs beginning with the first sequence number that you queried, and ending with the last sequence number minus 1. The log before the current log is the most recent archived log. For example, if the first query returned 9100, then start at 9100. If the second query returned 9112, then end at 9111.

For example, issue the following to back up the necessary archived logs:

```
RMAN>
BACKUP ARCHIVELOG FROM SEQUENCE 9100 UNTIL SEQUENCE 9111;
```

Backing Up and Deleting Multiple Copies of an Archived Redo Log: Example

In this scenario, you set initialization parameters so that you automatically archive redo logs to two directories: `*/oradata/trgt/arch/dest_1` and `*/oradata/trgt/arch/dest_2`. Therefore, you have two identical copies of the archived redo log for each log sequence number. You decide to back up each copy of the archived redo logs and then delete the originals. (Note that the degree of backup duplexing configured or specified in the `BACKUP` command determines the number of output files, independent of the number of input files. See [Duplexing Backup Sets](#) on page 6-2 for details.)

The easiest solution in this case is to use the `DELETE ALL INPUT` option means that RMAN deletes all logs that match the `ARCHIVELOG` criteria. Hence, it can remove all logs from both `*/oradata/trgt/arch/dest_1` and `*/oradata/trgt/arch/dest_2`.

For example, run the following command to back up all logs that could be used to recover from a point 10 days ago, and then delete all logs within the specified time range from disk:

```
RMAN> BACKUP DEVICE TYPE sbt
ARCHIVELOG ALL FROM TIME 'SYSDATE-10'
DELETE ALL INPUT;
```

Determining How Channels Distribute a Backup Workload: Example

When you create multiple backup sets and allocate multiple channels, RMAN automatically writes multiple backup sets in parallel. The allocated server sessions share the work of backing up the specified datafiles, control files, and archived redo logs. Note that you cannot stripe a single backup set across multiple channels.

RMAN automatically assigns a backup set to a device. You can use the `CHANNEL` parameter so that RMAN writes all backup sets for a *backupSpec* to a specific channel.

For example, this example parallelizes the backup operation by specifying which channels RMAN should back up to disk and which to `sbt`:

```
RMAN> RUN
{
  ALLOCATE CHANNEL ch1 DEVICE TYPE DISK FORMAT = '/backup/df/%U';
  ALLOCATE CHANNEL ch2 DEVICE TYPE DISK FORMAT = '/backup/cf/%U';
  ALLOCATE CHANNEL ch3 DEVICE TYPE sbt;
  BACKUP AS BACKUPSET # all output files are in backup sets
  # channel ch1 backs up datafiles to /backup/df directory
```

```

DATAFILE 1,2,3,4
CHANNEL ch1
# channel ch2 backs up control file copy to /backup/cf directory
CONTROLFILECOPY '/tmp/control01.ct1'
CHANNEL ch2;
BACKUP AS BACKUPSET
# channel ch3 backs up archived redo logs to tape
ARCHIVELOG FROM TIME 'SYSDATE-14'
CHANNEL ch3;
}

```

Note: You cannot back up to DISK and sbt at the same time using configured channels: you must manually allocate them.

Backing Up in NOARCHIVELOG Mode: Example

This script puts the database into the correct mode for a consistent, whole database backup and then backs up the database. The script performs a shutdown, startup, shutdown, and then startup again before creating multiple copies of the backup:

```

# Shut down database cleanly with immediate option. This type of shutdown lets
# current calls to the database complete, but prevents further logons or calls.
# If the database is not up, you receive a message saying so but RMAN will not
# treat this situation as an error.
SHUTDOWN IMMEDIATE;

# Start up the database in case it suffered instance failure or was
# closed with SHUTDOWN ABORT before starting this script.
# The script performs instance recovery if
# needed. Oracle uses the default init.ora file. Alternatively, use this form:
# STARTUP FORCE DBA pfile=filename.
# Use the DBA option because you are going to shut down again
# and do not want to let users in during the short interval. Use the FORCE
# option because it cannot hurt and might help in certain situations.
STARTUP FORCE DBA;
SHUTDOWN IMMEDIATE;

# The database is cleanly closed and ready for a consistent backup. RMAN
# requires that the database be started and mounted to perform a backup.
RMAN> STARTUP MOUNT;

# this example uses automatic channels to make the backup
BACKUP
  COPIES 2
  INCREMENTAL LEVEL 0
  MAXSETSIZE 10M
  DATABASE
  INCLUDE CURRENT CONTROLFILE
  TAG 'BACKUP_1'
;

# Now that the backup is complete, open the database.
ALTER DATABASE OPEN;

```

Note the inclusion of the current control file with the backup, and the use of the tag to identify the backup. To use this backup of the database, the control file must be restored from the same backup as the rest of the database. Adding INCLUDE CURRENT CONTROLFILE ensures that a usable backup of the control file is included

with the backup and tagged in order to simplify restoring the control file with the rest of the database.

You can skip tablespaces, such as read-only tablespaces, but any skipped tablespace that has not been offline or read-only since its last backup is lost if the database has to be restored from a backup.

Keeping a Long-Term Backup: Example

If you configure a retention policy, then you may want to exclude specified backups from this policy. For example, you may want to archive a consistent backup of the database once a year to serve as a historical record. This long-term backup does not function as a backup that you may perform recovery on, but an archived snapshot of data at a particular time.

To exempt a backup from the retention policy, specify the `KEEP` option on the `BACKUP` command. You can also specify `LOGS` or `NOLOGS` to indicate whether RMAN should save archived logs for possible recovery of this backup. If you specify `NOLOGS`, then the backup must be consistent.

This example keeps the backup of the database indefinitely and does not save archived logs needed to recover it:

```
RMAN> SHUTDOWN IMMEDIATE;
RMAN> STARTUP MOUNT; # put database in consistent state
RMAN> BACKUP DATABASE KEEP FOREVER NOLOGS
        TAG 'db_archive_1'; # make long-term consistent backup

# mark backup as unavailable in the repository so that RMAN does not attempt to
# restore it unless explicitly specified on the RESTORE command
RMAN> CHANGE BACKUP TAG 'db_archive_1' UNAVAILABLE;
RMAN> SQL 'ALTER DATABASE OPEN';
```

Using Backup Optimization: Examples

Run the `CONFIGURE BACKUP OPTIMIZATION` command to enable backup optimization. When specific conditions are met (described in "[Backup Optimization Algorithm](#)" on page 2-35), RMAN skips backups of files that are identical to files that are already backed up.

Assume that you configure optimization and a retention policy as follows:

```
CONFIGURE DEFAULT DEVICE TYPE TO sbt;
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 4 DAYS;
```

Optimizing a Database Backup: Example

Then, you run this command every night to back up the database to tape:

```
BACKUP DATABASE;
```

Because backup optimization is configured, RMAN skips backups of offline and read-only datafiles only if the most recent backups were made on or after the earliest point in the recovery window. RMAN does not skip backups when the most recent backups are older than the window. For example, optimization ensures you do not end up with a new backup of the read-only datafile `/?/oradata/trgt/history01.dbf` every night, so long as one backup set containing this file exists within the recovery window.

For example, if the most recent backup of the datafiles was on Sunday, and the point of recoverability (that is, the earliest date in the recovery window) is on Saturday, then RMAN skips the datafiles when you run the Wednesday backup. On Friday, the point of recoverability is now Monday, so the Sunday backup is now outside the window. Hence, the Friday backup does not skip the datafiles.

Optimizing a Daily Archived Log Backup to a Single Tape: Example

Assume that you want to back up all the archived logs every night. However, you do not want to have multiple copies of each log sequence number. So, you configure backup optimization to ON, then run this command in a script every night at 1 a.m.:

```
BACKUP DEVICE TYPE sbt ARCHIVELOG ALL;
```

RMAN skips all logs except those produced in the last 24 hours. In this way, you keep only one copy of each archived log on tape.

Optimizing a Daily Archived Log Backup to Multiple Tapes: Example

In this example, you back up logs that are not already on tape to one tape pool, then back up the same logs to a second tape pool. Finally, you delete old logs.

For the first step, perform the one-time configuration:

```
# configure backup optimization
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE DEFAULT DEVICE TYPE TO sbt;
```

Then, run the following script at the same time every night to back up the logs generated during the previous day to two separate tape pools:

```
# The following command backs up just the logs that are not on tape. The
# first copies are saved to the tapes from the pool "archivelog_pool_1"
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt
    PARMS='NSR_DATA_VOLUME_POOL=ARCHIVELOG_POOL_1';
  BACKUP ARCHIVELOG ALL;
}
# Make one more copy of the archived logs and save them to tapes from a
# different pool
RUN
{
  ALLOCATE CHANNEL c2 DEVICE TYPE sbt
    PARMS='NSR_DATA_VOLUME_POOL=ARCHIVELOG_POOL_2';
  BACKUP ARCHIVELOG
    FROM TIME 'SYSDATE-1'
    UNTIL TIME 'SYSDATE';
  # specify UNTIL so RMAN does not archive current log
}
# Delete old logs - for example, delete logs created within the last week.
DELETE ARCHIVELOG ALL COMPLETED AFTER 'SYSDATE-7';
```

Note: Specifying the UNTIL TIME clause causes RMAN to not use backup optimization in deciding which archive logs to back up. In this case, logs are backed up even if other usable backups are already available.

Creating a Weekly Secondary Backup of Archived Logs: Example

Assume a more sophisticated scenario in which your goal is to back up the archived logs to tape every day. However, you are worried about tape failure, so you want to ensure that you have more than copy of each log sequence number on an separate tape before you perform your weekly deletion of logs from disk.

First, perform a one-time configuration:

```
# configure backup optimization
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
CONFIGURE default DEVICE TYPE TO sbt;
# configure a default channel that sends backups to tape pool "first_copy"
CONFIGURE CHANNEL DEVICE TYPE sbt PARMS='ENV=(NSR_DATA_VOLUME_POOL=first_copy);
```

Because you have optimization enabled, you can run the following command every evening to back up all archived logs to the "first_copy" pool that have not already been backed up:

```
BACKUP ARCHIVELOG ALL TAG first_copy;
```

Every Friday evening you create an additional backup of all archived logs in a different tape pool. Also, at the end of the backup, you want to delete all archived logs that already have at least two copies on tape. So you run the following script:

```
RUN
{
  # manually allocate a channel, in order to specify that the backup run by this
  # channel should go to both pools "first_copy" and "second_copy"
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt
    PARMS='ENV=(NSR_DATA_VOLUME_POOL=second_copy)';
  ALLOCATE CHANNEL c2 DEVICE TYPE sbt
    PARMS='ENV=(NSR_DATA_VOLUME_POOL=first_copy)';
  BACKUP CHANNEL C1 ARCHIVELOG UNTIL TIME 'SYSDATE'
    NOT BACKED UP 2 TIMES # back up only logs without 2 backups on tape
    TAG SECOND_COPY;
  BACKUP CHANNEL C2 ARCHIVELOG UNTIL TIME 'SYSDATE'
    NOT BACKED UP 2 TIMES # back up only logs without 2 backups on tape
    TAG FIRST_COPY;
}

# now delete from disk all logs that have been backed up to tape at least twice
DELETE ARCHIVELOG ALL
  BACKED UP 2 TIMES TO DEVICE TYPE sbt;
```

The Friday script creates a second copy of all archived logs in the "second_copy" tape pool. After the backup, you can send the tape from the pool "second_copy" to secure storage. You should use this tape backup only if the primary tape from pool "first_copy" is damaged. Because the secondary tape is in a secure place, you do not want RMAN to use it for recovery, so you can mark the backup as unavailable:

```
CHANGE BACKUP OF ARCHIVELOG TAG SECOND_COPY UNAVAILABLE;
```

Handling Corruption During Backups: Example

By default a checksum is calculated for every block read from a datafile and stored in the backup or image copy. If you use the `NOCHECKSUM` option, then checksums are not calculated. If the block already contains a checksum, however, then the checksum is validated and stored in the backup. If the validation fails, then the block is marked corrupt in the backup.

The `SET MAXCORRUPT FOR DATAFILE` command sets how many corrupt blocks in a datafile that `BACKUP` will tolerate. If a datafile has more corrupt blocks than specified by the `MAXCORRUPT` parameter, the command terminates. If you specify the `CHECK LOGICAL` option, RMAN checks for logical and physical corruption.

By default, the `BACKUP` command terminates when it cannot access a datafile. You can specify parameters to prevent termination, as listed in the following table.

If you specify the option ... Then RMAN skips...

<code>SKIP INACCESSIBLE</code>	Inaccessible datafiles. A datafile is only considered inaccessible if it cannot be read. Some offline datafiles can still be read because they exist on disk. Others have been deleted or moved and so cannot be read, making them inaccessible.
<code>SKIP OFFLINE</code>	Offline datafiles.
<code>SKIP READONLY</code>	Datafiles in read-only tablespaces.

The following example uses an automatic channel to back up the database, and sets the corruption level for the datafile in the `SYSTEM` tablespace so that up to 10 errors will be accepted:

```
RMAN> RUN
{
  SET MAXCORRUPT FOR DATAFILE 1 TO 10;
  BACKUP DATABASE
    SKIP INACCESSIBLE
    SKIP READONLY
    SKIP OFFLINE;
}
```

Advanced RMAN Recovery Techniques

This chapter describes how to perform restore and recovery using RMAN in a number of advanced scenarios.

This chapter contains the following topics:

- [Restore and Recovery of NOARCHIVELOG Databases](#)
- [Restore and Recovery of the Database on a New Host](#)
- [Performing Recovery with a Backup Control File](#)
- [Performing Disaster Recovery](#)
- [Performing Block Media Recovery with RMAN](#)
- [RMAN Restore and Recovery Examples](#)

Restore and Recovery of NOARCHIVELOG Databases

Restore of a database running in NOARCHIVELOG mode is very similar to restore of a database in ARCHIVELOG mode. The main differences are:

- Only cold backups (that is, backups created when the database was shut down normally) can be used in restoring a database in NOARCHIVELOG mode
- Media recovery is not possible, because there are no archived logs

A limited form of restore and recovery is possible for NOARCHIVELOG databases if the backup strategy for the database includes incremental backups. The incremental backups (which, like the full backup of a NOARCHIVELOG database must be created when the database is shut down) can be applied to a full database backup to apply recent changes up to the time of the incremental backup.

Recovering a Database in NOARCHIVELOG Mode Using Incremental Backups

You can perform limited recovery of changes to a database running in NOARCHIVELOG mode by applying incremental backups. Note that the incremental backups must be consistent, like all backups of a database run in NOARCHIVELOG mode, so you cannot make backups of the database when it is open.

Assume the following scenario:

- You run database `trgt` in NOARCHIVELOG mode.
- You use a recovery catalog.
- You shut down the database consistently and make a level 0 backup of database `trgt` to tape on Sunday afternoon.

- You shut down the database consistently and make a level 1 differential incremental backup to tape at 3:00 a.m. on Wednesday and Friday.
- The database has a media failure on Saturday, destroying half of the datafiles as well as the online redo logs.

In this case, you must perform an incomplete media recovery until Friday, the date of the most recent incremental backup. RMAN uses the level 0 Sunday backup as well as the Wednesday and Friday level 1 backups.

Because the online redo logs are lost, you must specify the `NOREDO` option in the `RECOVER` command.

You must also specify `NOREDO` if the online logs are available but the redo cannot be applied to the incrementals.

If you do not specify `NOREDO`, then RMAN searches for redo logs after applying the Friday incremental backup, and issues an error message when it does not find them.

After connecting to `trgt` and the catalog database, recover the database with the following command:

```
STARTUP FORCE MOUNT;
RESTORE CONTROLFILE; # restore control file from consistent backup
ALTER DATABASE MOUNT;
RESTORE DATABASE; # restore datafiles from consistent backup
RECOVER DATABASE NOREDO; # specify NOREDO because online redo logs are lost
ALTER DATABASE OPEN RESETLOGS;
```

The recovered database reflects only changes up through the time of the Friday incremental backup. Because there are no archived redo logs, there is no way to recover changes made after the incremental backup.

Note: If the current online logs contain all changes since the last incremental, then you can run `RECOVER DATABASE` without specifying `NOREDO`. In such a case, the changes in the online logs are applied.

Restore and Recovery of the Database on a New Host

The procedure described in this section can be used to perform test restores, or to move a production database to a new host.

Note: If your goal is to perform a test run of the disaster recovery procedures you would use following a real disaster, or to permanently move the target database to the new host, then use the procedure described in this section, which uses the `RESTORE` and `RECOVER` commands.

Note, however, that the DBID for the restored test database will be the same as the DBID for the original database. If, after the restore and recovery process is complete, you connect to the test database and the recovery catalog, the recovery catalog is updated with information about the test database that can interfere with RMAN's ability to restore and recover the source database.

If your goal is to create a new copy of your target database for ongoing use on a new host, then use the `RMAN DUPLICATE` command instead of this procedure. `DUPLICATE` assigns a new DBID to the duplicate database it creates, allowing it to be registered in the same recovery catalog as the original target database. See "[Creating a Duplicate Database with RMAN: Overview](#)" on page 13-1 for details about duplicating a database.

Preparing for Restore of a Database to a New Host

To prepare for the restore of the database to a new host, take the following steps:

- Record the DBID for your source database. If you do not know the DBID for your database, see *Oracle Database Backup and Recovery Basics* for details on ways to determine the DBID.
- Make the source database initialization parameter file accessible on the new host. Copy the file from the old host to a new host using an operating system utility.
- Make sure backups used for the restore are accessible on the restore host. For example, if the backups were made with a media manager, then make sure the tape device is connected to the new host.

Note: If you perform a test restore only, then **do not** connect to the recovery catalog when restoring the datafiles. Otherwise, RMAN records information about the restored datafiles to the recovery catalog. This interferes with future attempts to restore and recover the primary database. If you must use a recovery catalog because the control file is not large enough to contain the RMAN repository data on all of the backups that you need to restore, then export the catalog and import it into a different schema or database and use the copied recovery catalog for the test restore. Otherwise, the catalog considers the restored database as the current target database.

Testing the Restore of a Database to a New Host: Scenario

This scenario assumes the following:

- Two networked machines, `hosta` and `hostb`, are running Linux
- A target database named `trgta` is on `hosta` and uses a recovery catalog `catdb`
- You want to test the restore and recovery of `trgta` on `hostb`, while keeping database `trgta` up and running on `hosta`

- The directory structure of `hostb` is different from `hosta`, so that `trgta` is located in `/net/hosta/dev3/oracle/dbs`, but you want to restore the database to `/net/hostb/oracle/oradata/test`
- Database `trgta` uses a server parameter file (not a client-side initialization parameter file)
- The `ORACLE_SID` for the `trgta` database is `trgta` and will not change for the restored database
- You have a record of the `DBID` for `trgta`
- A media manager is accessible by both machines
- You have recoverable backups on tape of all datafiles
- You have backups of the archived logs required to recover the datafiles
- You have control file and server parameter file autobackups on tape

Use the following steps to perform the restore process:

1. Make backups of the target database available to `hostb`. To test disaster recovery, you need to have a recoverable backup of the target database. When preparing your disaster recovery strategy, ensure that the backups of the datafiles, control files, and server parameter file are restorable on `hostb`. Hence, you must configure the media management software so that `hostb` is a media manager client and can read the backup sets created on `hosta`. Consult the media management vendor for support on this issue.
2. Configure the `ORACLE_SID` on `hostb`. This scenario assumes that you want to authenticate yourself through the operating system, which is much faster than configuring Oracle Net and creating a password file. However, you must be connected to `hostb` either locally or through a SQLNet alias.

While connected to `hostb` with administrator privileges, edit the `/etc/group` file so that you are included: in the DBA group:

```
dba:*:614:<your_user_name>
```

Set the `ORACLE_SID` environment variable on `hostb` to the same value used on `hosta`:

```
% setenv ORACLE_SID trgta
```

Start RMAN and connect to the target instance **without** connecting to the recovery catalog.

```
% rman TARGET / NOCATALOG
```

3. Start the instance without mounting it. To start the instance, you first need to set the `DBID`. (If you do not know the `DBID` for your database, see *Oracle Database Backup and Recovery Basics* for details on how to determine the `DBID`.)

Run `SET DBID` to set the `DBID`, then run `STARTUP NOMOUNT`:

```
SET DBID 1340752057;  
STARTUP NOMOUNT
```

RMAN will fail to find the server parameter file, which has not yet been restored, but will start the instance with a "dummy" file. Sample output follows:

```
startup failed: ORA-01078: failure in processing system parameters  
LRM-00109: could not open parameter file '/net/hostb/oracle/dbs/inittrgta.ora'
```

```
trying to start the Oracle instance without parameter files ...
Oracle instance started
```

4. Restore and edit the server parameter file.

Because you enabled the control file autobackup feature when making your backups, the server parameter file is included in the backup sets.

Allocate a channel to the media manager, then restore the server parameter file (SPFILE) as a client-side parameter file (PFILE).

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS='...';
  RESTORE SPFILE TO PFILE '?/oradata/test/inittrgta.ora' FROM AUTOBACKUP;
  SHUTDOWN ABORT;
}
```

Next, edit the restored PFILE. Change any location-specific parameters, for example, those ending in `_DEST` and `_PATH`, to reflect the new directory structure. For example, edit the following parameters:

```
- IFILE
- *_DUMP_DEST
- LOG_ARCHIVE_DEST*
- CONTROL_FILES
```

Then restart the instance, using the edited PFILE:

```
STARTUP FORCE NOMOUNT PFILE='?/oradata/test/inittrgta.ora';
```

5. Restore the control file from an autobackup and then mount the database. RMAN restores the control file to whatever location you specified in the `CONTROL_FILES` initialization parameter. For example:

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS='...';
  RESTORE CONTROLFILE FROM AUTOBACKUP;
  ALTER DATABASE MOUNT;
}
```

6. Query the database filenames recorded in the control file on the new host (hostb). Because the control file is from the `trgta` database, the recorded filenames use the original `hosta` filenames. You can query `V$` views to obtain this information. Start a new `SQL*Plus` session and connect to the newly created instance on `hostb`:

```
% sqlplus '/ AS SYSDBA'
```

Run the following query in `SQL*Plus`:

```
SQL> COLUMN NAME FORMAT a60
SQL> SPOOL LOG 'db_filenames.out'
SQL> SELECT FILE# AS "File/Grp#", NAME FROM V$DATAFILE
        UNION
        SELECT GROUP#,MEMBER FROM V$LOGFILE;
SQL> SPOOL OFF
SQL> EXIT
```

7. Write the RMAN recovery script. The script must include the following steps:

- For each datafile on the destination host that is restored to a different path than it had on the source host, use a `SET NEWNAME` command to specify the

new path on the destination host. (If the file systems on the destination system are set up to have the same paths as the source host, then do not use SET NEWNAME for those files restored to the same path as on the source host.)

- For each online redo log that is to be created at a different location than it had on the source host, use SQL ALTER DATABASE RENAME FILE commands to specify the pathname on the destination host. (If the file systems on the destination system are set up to have the same paths as the source host, then do not use ALTER DATABASE RENAME FILE for those files restored to the same path as on the source host.)
- Perform a SET UNTIL to limit media recovery to the end of the archived redo logs.
- Run SWITCH so that the control file recognizes the new path names as the official new names of the datafiles
- Restore and recover the database

For example, consider the following RMAN script to perform these steps, which is contained in text file reco_test.rman:

```
RUN
{
  # allocate a channel to the tape device
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS='...';

  # rename the datafiles and online redo logs
  SET NEWNAME FOR DATAFILE 1 TO '?/oradata/test/system01.dbf';
  SET NEWNAME FOR DATAFILE 2 TO '?/oradata/test/undotbs01.dbf';
  SET NEWNAME FOR DATAFILE 3 TO '?/oradata/test/cwmlite01.dbf';
  SET NEWNAME FOR DATAFILE 4 TO '?/oradata/test/drsys01.dbf';
  SET NEWNAME FOR DATAFILE 5 TO '?/oradata/test/example01.dbf';
  SET NEWNAME FOR DATAFILE 6 TO '?/oradata/test/indx01.dbf';
  SET NEWNAME FOR DATAFILE 7 TO '?/oradata/test/tools01.dbf';
  SET NEWNAME FOR DATAFILE 8 TO '?/oradata/test/users01.dbf';
  SQL "ALTER DATABASE RENAME FILE '/dev3/oracle/dbs/redo01.log'
      TO '?/oradata/test/redo01.log' ";
  SQL "ALTER DATABASE RENAME FILE '/dev3/oracle/dbs/redo02.log'
      TO '?/oradata/test/redo02.log' ";

  # Do a SET UNTIL to prevent recovery of the online logs
  SET UNTIL SCN 123456;
  # restore the database and switch the datafile names
  RESTORE DATABASE;
  SWITCH DATAFILE ALL;

  # recover the database
  RECOVER DATABASE;
}
EXIT
```

Online logs and datafiles are relocated as specified,

For example, connect and execute the script as shown here:

```
% rman TARGET / NOCATALOG
RMAN> @reco_test.rman
```

RMAN will apply as many of the archived redo logs as it can and leave the database in a state in which it can be opened.

8. Now perform an OPEN RESETLOGS at the restored database.

Caution: When you re-open your database in the next step, **do not connect to the recovery catalog**. Otherwise, the new database incarnation created is registered automatically in the recovery catalog, and the filenames of the production database are replaced by the new filenames specified in the script.

If this is a test restore, **never connect RMAN to the test-restore database and the recovery catalog**.

From the RMAN prompt, open the database with the RESETLOGS option:

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

9. If this was a test restore, and it was successful, then you can shut down the test database instance, and delete the test database with all of its files. Use the DROP DATABASE command to delete all files associated with the database automatically.

Note: If you used an ASM disk group, then DROP DATABASE is the only way to safely remove the files of the test database. If you restored to non-ASM storage then you can also use operating system commands to remove the database.

```
RMAN> SHUTDOWN ABORT
RMAN> EXIT
```

Remove all test files. You can do this with an operating system utility or in RMAN. For example, in Unix you could perform the procedure this way:

```
% rm $ORACLE_HOME/oradata/test/*
```

You can also use RMAN for a procedure that works on all platforms. For example:

```
RMAN> STARTUP FORCE NOMOUNT PFILE='?/oradata/test/inittrgta.ora';
RMAN> DROP DATABASE;
```

Because you did not perform the restore and recovery when connected to the recovery catalog, the recovery catalog contains no records for any of the restored files or the procedures performed during the test. Likewise, the control file of the trgta database is completely unaffected by the test.

Performing Recovery with a Backup Control File

If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file before you can perform recovery. When using a backup control file, and using a recovery catalog, the process is identical to recovery with a current control file, as the RMAN repository information missing from the backup control file is available from the recovery catalog. There are special considerations when using a backup controlfile and not using a recovery catalog.

The following notes and restrictions apply regardless of whether you use a recovery catalog:

- You must run the RECOVER command after restoring a backup control file, even if no datafiles have been restored.
- You must open the database with the RESETLOGS option after performing either complete *or* point-in-time recovery with a backup control file.
- If the online redo logs are inaccessible, then you must perform incomplete recovery to an SCN before the earliest SCN in the online redo logs. This limitation is necessary because RMAN does not back up online logs.
- During recovery, RMAN automatically searches for online and archived redo logs that are not recorded in the RMAN repository, and catalogs any that it finds so that it can use them in recovery.

RMAN attempts to find a valid archived log in any of the current archiving destinations with the current log format. The current format is specified in the initialization parameter file used to start the instance (or all instances in a Real Application Clusters installation). Similarly, RMAN attempts to find the online redo logs by using the filenames as specified in the control file.

If you changed the archiving destination or format during recovery, or if you added new online log members after the backup of the control file, then RMAN may not be able to automatically catalog a needed online or archived log. In this situation, RMAN reports errors similar to the following:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of recover command at 08/29/2001 14:23:09
RMAN-06054: media recovery requesting unknown log: thread 1 scn 86945
```

In this case, you must use the CATALOG command to manually add the required logs to the repository so that recovery can proceed. The cataloging procedure is described in *Oracle Database Backup and Recovery Basics*.

Performing Recovery with a Backup Control File and No Recovery Catalog: Scenario

This section assumes that you have RMAN backups of the control file, but do not use a recovery catalog.

Assuming that you enabled the control file autobackup feature for the target database, you can restore an autobackup of the control file. Because the autobackup uses a default format, RMAN can restore it even though it does not have a repository available that lists the available backups. You can restore the autobackup to the default or a new location. RMAN replicates the control file to all CONTROL_FILES locations automatically.

Note: If you know the backup piece name (for example, from the media manager or because the piece is on disk), then you can specify the piece name using the RESTORE CONTROLFILE FROM '*filename*' command. The server records the location of every autobackup in the alert log.

Because you are not connected to a recovery catalog, the RMAN repository contains only information about available backups at the time of the control file backup. If you know the location of other usable backup sets or image copies, add them to the control file RMAN repository with the CATALOG command.

Because the repository is not available when you restore the control file, you must use the `SET DBID` command to identify the target database. The DBID is used to determine the location of control file autobackups. Use `SET DBID` command only in the following special circumstances:

- You are not connected to a recovery catalog and want to restore the control file or server parameter file.
- You are connected to a recovery catalog and want to restore the control file, but the database name is not unique in the recovery catalog.
- The server parameter file is lost and you want to restore it.

To recover the database with an autobackup of the control file without a recovery catalog:

1. Start RMAN and connect to the target database. For example, run:

```
CONNECT TARGET /
```

2. Start the target instance without mounting the database. For example:

```
STARTUP NOMOUNT;
```

3. Set the database identifier for the target database with `SET DBID`. RMAN displays the DBID whenever you connect to the target. You can also obtain it by inspecting saved RMAN log files, querying the catalog, or looking at the filenames of control file autobackup. (refer to ["Restoring Control File When Databases in the Catalog Have the Same Name: Example"](#) on page 7-15). For example, run:

```
SET DBID 676549873;
```

4. Restore the autobackup control file, then perform recovery. Do the following:
 - a. Optionally, specify the most recent backup time stamp that RMAN can use when searching for a control file autobackup to restore.
 - b. If you know that a different control file autobackup format was in effect when the control file autobackup was created, then specify a nondefault format for the restore of the control file.
 - c. If the channel that created the control file autobackup was device type `sbt`, then you must allocate one or more `sbt` channels. Because no repository is available, you cannot use preconfigured channels.
 - d. Restore the autobackup of the control file, optionally setting the maximum number of days backward that RMAN can search (up to 366) and the initial sequence number that it should use in its search for the first day.
 - e. If you know that your control file contained information about configured channels that will be useful to you in the rest of the restore process, you can exit the RMAN client at this point, to clear manually allocated channels from step "c". If you then restart the RMAN client and mount the database those configured channels become available for your use in the rest of the restore and recovery process.

If you do not care about using configured channels from your control file, then you can simply mount the database at this point.

- f. If the online logs are usable, then perform a complete restore and recovery as described in *Oracle Database Backup and Recovery Basics*.

Otherwise, restore and perform incomplete recovery of the database, as described in *Oracle Database Backup and Recovery Basics*. Use an `UNTIL` clause

to specify a target time , SCN or log sequence number for the recovery prior to the first SCN of the online redo logs.

In this example, the online redo logs have been lost, and the most recent archived log sequence number is 13243. This example shows how to restore the control file autobackup, then performs recovery of the database to log sequence 13243.

```
RUN
{
  # Optionally, set upper limit for eligible time stamps of control file
  # backups
  # SET UNTIL TIME '09/10/2000 13:45:00';
  # Specify a nondefault autobackup format only if required
  # SET CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK
  # TO '?/oradata/%F.bck';
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARS='...'; # allocate manually
  RESTORE CONTROLFILE FROM AUTOBACKUP
  MAXSEQ 100          # start at sequence 100 and count down
  MAXDAYS 180;       # start at UNTIL TIME and search back 6 months
  ALTER DATABASE MOUNT DATABASE;
}
# uses automatic channels configured in restored control file
RESTORE DATABASE UNTIL SEQUENCE 13243;
RECOVER DATABASE UNTIL SEQUENCE 13243; # recovers to latest archived log
```

5. If recovery was successful, then open the database and reset the online logs:

```
ALTER DATABASE OPEN RESETLOGS;
```

Performing Disaster Recovery

Disaster recovery includes the restore of and recovery of the target database after the loss of the entire target database, the recovery catalog database, all current control files, all online redo log files, and all parameter files.

To perform a disaster recovery, the minimum required set of backups is backups of some datafiles, some archived redo logs generated after the time of the backup, and at least one autobackup of the control file.

The basic procedure for disaster recovery begins with restoring an autobackup of the server parameter file, as described in *Oracle Database Backup and Recovery Basics*.

Once you have an SPFILE, you can start the target database instance, restore the control file from autobackup and mount it.

With the control file mounted, then follow the instructions found in "[Performing Recovery with a Backup Control File](#)" on page 7-7 to restore and recover your datafiles.

Note: If you are restoring to a new host, you should review the considerations described in "[Restore and Recovery of the Database on a New Host](#)" on page 7-2.

The following scenario restores and recovers the database to the most recently available archived log, which in this example is log 1124 in thread 1. It assumes that:

- You are restoring the database to a new host with the same directory structure.
- You have one tape drive containing backups of all the datafiles and archived redo logs through log 1124, as well as autobackups of the control file and server parameter file.

- You do not use a recovery catalog.

In this scenario, perform the following steps:

1. If possible, restore all relevant network files such as `tnsnames.ora` and `listener.ora` by means of operating system utilities.
2. Start RMAN and connect to the target database. If you do not have the Oracle Net files, then connect using operating system authentication.
3. Specify the DBID for the target database with the `SET DBID` command, as described in ["Performing Recovery with a Backup Control File and No Recovery Catalog: Scenario"](#) on page 7-8.
4. Run the `STARTUP NOMOUNT` command. RMAN attempts to start the instance with a dummy server parameter file.
5. Allocate a channel to the media manager and then run the `RESTORE SPFILE FROM AUTOBACKUP` command.
6. Run `STARTUP FORCE NOMOUNT` mode so that the instance is restarted with the restored server parameter file.
7. Allocate a channel to the media manager and then restore a control file autobackup (refer to ["Performing Recovery with a Backup Control File and No Recovery Catalog: Scenario"](#) on page 7-8).
8. Mount the restored control file.
9. Catalog any backups not recorded in the repository with the `CATALOG` command (refer to ["Removing DELETED Records From the Recovery Catalog After Upgrade"](#) on page 10-9).
10. Restore the datafiles to their original locations. If volume names have changed, then run `SET NEWNAME` commands before the restore and perform a switch after the restore to update the control file with the new locations for the datafiles (refer to ["Performing Disaster Recovery"](#) on page 7-10).
11. Recover the datafiles. RMAN stops recovery when it reaches the log sequence number specified.
12. Open the database in `RESETLOGS` mode. Only complete this last step if you are certain that no other archived logs can be applied.

```
# Start RMAN and connect to the target database
% rman TARGET SYS/oracle@trgt

# Set the DBID for the target database
RMAN> SET DBID 676549873;
RMAN> STARTUP FORCE NOMOUNT; # rman starts instance with dummy parameter file
RUN
{
  ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
  RESTORE SPFILE FROM AUTOBACKUP;
}
# Restart instance with restored server parameter file
RMAN> STARTUP FORCE NOMOUNT;

RMAN> RUN
{
  # Manually allocate a channel to the media manager
  ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
  # Restore autobackup of the control file. This example assumes that you have
  # accepted the default format for the autobackup name.
```

```

RESTORE CONTROLFILE FROM AUTOBACKUP;
# The set until command is used in case the database
# structure has changed in the most recent backups, and you wish to
# recover to that point-in-time. In this way RMAN restores the database
# to the same structure that the database had at the specified time.
ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 1124 THREAD 1;
RESTORE DATABASE;
RECOVER DATABASE;
}
RMAN> ALTER DATABASE OPEN RESETLOGS; # Reset the online logs after recovery
completes

```

The following example of the RUN command shows the same scenario except with new filenames for the restored datafiles:

```

RMAN> RUN
{
# If you need to restore the files to new locations, tell Recovery Manager
# to do this using SET NEWNAME commands:
SET NEWNAME FOR DATAFILE 1 TO '/dev/vgd_1_0/rlvt5_500M_1';
SET NEWNAME FOR DATAFILE 2 TO '/dev/vgd_1_0/rlvt5_500M_2';
SET NEWNAME FOR DATAFILE 3 TO '/dev/vgd_1_0/rlvt5_500M_3';
ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
RESTORE CONTROLFILE FROM AUTOBACKUP;
ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 124 THREAD 1;
RESTORE DATABASE;
SWITCH DATAFILE ALL; # Update control file with new location of datafiles.
RECOVER DATABASE;
}
RMAN> ALTER DATABASE OPEN RESETLOGS;

```

Performing Block Media Recovery with RMAN

The BLOCKRECOVER command can restore and recover individual datablocks within a datafile. This procedure is useful when a trace file or standard output reveals that a small number of blocks within a datafile are corrupt.

Block media recovery is not useful in cases where the extent of data loss or corruption is not known; in this case, use datafile recovery instead.

See Also:

- ["Block Media Recovery with RMAN"](#) on page 3-7 for an overview of block media recovery,
- *Oracle Database Backup and Recovery Reference* for BLOCKRECOVER syntax
- *Oracle Database Reference* for details about the \$DATABASE_BLOCK_CORRUPTION view

Block Media Recovery Using All Available Backups

In this scenario, you identify the blocks that require recovery and then use any available backup to perform the restore and recovery of these blocks.

To recover datablocks by using all available backups:

1. Obtain the datafile numbers and block numbers for the corrupted blocks. Typically, you obtain this output from the standard output, the `alert.log`, trace files, or a media management interface. For example, you may see the following in a trace file:

```
ORA-01578: ORACLE data block corrupted (file # 8, block # 13)
ORA-01110: data file 8: '/oracle/oradata/trgt/users01.dbf'
ORA-01578: ORACLE data block corrupted (file # 2, block # 19)
ORA-01110: data file 2: '/oracle/oradata/trgt/undotbs01.dbf'
```

2. Assuming that you have preallocated automatic channels, run the `BLOCKRECOVER` command at the RMAN prompt, specifying the file and block numbers for the corrupted blocks as in the following example:

```
RMAN> BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19;
```

Block Media Recovery Using Specific Backups

In this scenario, you identify the blocks that require recovery, and then use only selected backups to perform the restore and recovery of these blocks.

To recover datablocks while limiting the type of backup:

1. Obtain the datafile numbers and block numbers for the corrupted blocks. Typically, you obtain this output from the standard output, the `alert.log`, trace files, or a media management interface. For example, you may see the following in a trace file:

```
ORA-01578: ORACLE data block corrupted (file # 8, block # 13)
ORA-01110: data file 8: '/oracle/oradata/trgt/users01.dbf'
ORA-01578: ORACLE data block corrupted (file # 2, block # 19)
ORA-01110: data file 2: '/oracle/oradata/trgt/undotbs01.dbf'
```

2. Assuming that you have preallocated automatic channels, execute the `BLOCKRECOVER` command at the RMAN prompt, specifying the file and block numbers for the corrupted blocks and limiting the backup candidates by means of the available options. For example, you can specify what type of backup should be used to restore the blocks:

```
# restore from backupset
RMAN>
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19 FROM BACKUPSET;
# restore from datafile image copy
RMAN>
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19
        FROM DATAFILECOPY;
```

You can indicate the backup by specifying a tag:

```
# restore from backupset with tag "mondayam"
RMAN>
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 199
        FROM TAG = mondayam;
```

You can limit the backup candidates to those made before a certain point:

```
# restore using backups made before one week ago
RMAN>
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19
        RESTORE UNTIL 'SYSDATE-7';
# restore using backups made before SCN 100
RMAN>
```

```
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19
  RESTORE UNTIL SCN 100;
# restore using backups made before log sequence 7024
RMAN>
BLOCKRECOVER DATAFILE 8 BLOCK 13 DATAFILE 2 BLOCK 19
  RESTORE UNTIL SEQUENCE 7024;
```

Note that if you limit the restore of datablocks with the UNTIL clause, then RMAN must perform more recovery on the blocks, and the recovery phase must scan all logs for changes to the specified blocks.

Block Media Recovery of Blocks Listed in V\$DATABASE_BLOCK_CORRUPTION

The V\$DATABASE_BLOCK_CORRUPTION view indicates which blocks in a datafile were marked corrupt since the most recent BACKUP or BACKUP VALIDATE command was run. After a corrupt block is repaired, the row identifying this block is deleted from the view.

You can check for logical corruption in the database by running the BACKUP (with or without VALIDATE option) with the CHECK LOGICAL command. If RMAN finds corrupt blocks, then it populates V\$DATABASE_BLOCK_CORRUPTION. The backup will stop if the number of corrupt blocks exceeds MAXCORRUPT. A historical record of block corruptions in RMAN backups is kept in V\$BACKUP_CORRUPTION and V\$COPY_CORRUPTION.

In this scenario, you identify the blocks that require recovery by querying V\$DATABASE_BLOCK_CORRUPTION, and then instruct RMAN to recover all blocks listed in this view by means of the CORRUPTION LIST keyword.

To recover datablocks while limiting the type of backup:

1. Query V\$DATABASE_BLOCK_CORRUPTION to determine whether corrupt blocks exist in the most recent backups of the datafiles:

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

2. Assuming that you have preallocated automatic channels, recover all blocks marked corrupt in V\$DATABASE_BLOCK_CORRUPTION by running the BLOCKRECOVER CORRUPTION LIST command. For example, this command restores blocks from backups created more than 10 days ago:

```
BLOCKRECOVER CORRUPTION LIST
  RESTORE UNTIL TIME 'SYSDATE-10';
```

See *Oracle Database Backup and Recovery Reference* for more details on block media recovery in RMAN.

RMAN Restore and Recovery Examples

The following sections illustrate the use of RMAN restore and recovery techniques in advanced scenarios.

Restoring Datafile Copies to a New Host: Example

To move the database to a new host by means of datafile copies, you must transfer the copies manually to the new machine. This example assumes that you are using a recovery catalog.

1. After connecting to the target database and recovery catalog, run a `LIST` command to see a listing of datafile copies and their associated primary keys, as in the following example:

```
LIST COPY;
```

2. Copy the datafile copies to the new host with an operating system utility. For example, in UNIX:

```
% cp -r /tmp/*dbf /net/new_host/oracle/oradata/trgt
```

3. Start RMAN and then uncatalog the datafile copies on the old host. For example, enter:

```
CHANGE COPY OF DATAFILE 1,2,3,4,5,6,7,8 UNCATALOG;
```

4. Catalog the datafile copies, using their new filenames or `CATALOG START WITH` (if you know all the files are in directories with a common prefix easily addressed with a `CATALOG START WITH`). For example, run:

```
CATALOG START WITH '?/oradata/trgt/';
```

Or this example specifies files individually:

```
CATALOG DATAFILECOPY
  '?/oradata/trgt/system01.dbf', '?/oradata/trgt/undotbs01.dbf',
  '?/oradata/trgt/cwmlite01.dbf', '?/oradata/trgt/drsys01.dbf',
  '?/oradata/trgt/example01.dbf', '?/oradata/trgt/indx01.dbf',
  '?/oradata/trgt/tools01.dbf', '?/oradata/trgt/users01.dbf';
```

5. Perform the restore and recovery operation described in ["Performing Disaster Recovery"](#) on page 7-10.

Restoring Control File When Databases in the Catalog Have the Same Name: Example

When using a recovery catalog and attempting to restore a lost control file, you encounter an error if there are other databases registered in the recovery catalog with the same name as your target database.

To resolve this error, you must uniquely identify the database by DBID for the restore operation. This requires determining the correct DBID for your database, and then using the `SET DBID` command to identify the target database before the `RESTORE CONTROLFILE` command, as shown in the following example:

1. Start RMAN and connect to the target database.
2. Run the `STARTUP FORCE NOMOUNT` command.
3. Run the `SET DBID` command to distinguish this connected target database from other target databases that have the same name.
4. Run the `RESTORE CONTROLFILE` command. After restoring the control file, you can mount the database to restore the rest of the database.

See Also: *Oracle Database Backup and Recovery Reference* for more details on the use of `SET DBID`.

Restoring a Backup Control File By Using the DBID

To set the DBID, connect RMAN to the target database and run the following `SET` command, where *target_dbid* is the value you obtained from the previous step:

```
SET DBID = target_dbid;
```

To restore the control file to its default location and then mount it, run:

```
RESTORE CONTROLFILE;  
ALTER DATABASE MOUNT;
```

To restore and recover the database, run:

```
RESTORE DATABASE;  
RECOVER DATABASE  
# optionally, delete logs restored for recovery and limit disk space used  
DELETE ARCHIVELOG MAXSIZE 2M;
```

Recovering a Lost Datafile Without a Backup: Example

RMAN can handle lost datafiles without user intervention during restore and recovery. When a datafile is lost, the possible cases can be classified as follows:

- The control file knows about the datafile, that is, the user backed up the control file after datafile creation, but the datafile itself is not backed up. If the datafile record is in the control file, then `RESTORE` creates the datafile in the original location or in a user-specified location (for example, with `SET NEWNAME`). The `RECOVER` command can then apply the necessary logs to the datafile.
- The control file does not have the datafile record, that is, the user did not back up the control file after datafile creation. During recovery, the database will detect the missing datafile and report it to RMAN, which will create a new datafile and continue recovery by applying the remaining logs. If the datafile was created in a parent incarnation, it will be created during restore or recover as appropriate.

In this example, the following sequence of events occurs:

1. You make a whole database backup of your `ARCHIVELOG` mode database.
2. You create a tablespace `history` containing a single datafile called `/mydb/history01.dbf`.
3. You populate the newly created datafile with data.
4. You archive all the active online redo logs.
5. A user accidentally deletes the datafile `history01.dbf` from the operating system before you have a chance to back it up.

In this case, the current control file knows about the datafile. To restore and recover the datafile, start RMAN, connect to the target database, and then enter the following commands at the RMAN prompt:

```
# take the tablespace with the missing datafile offline  
SQL "ALTER TABLESPACE history OFFLINE IMMEDIATE";  
# restore the tablespace even though you have no backup  
RESTORE TABLESPACE history;  
# recover tablespace  
RECOVER TABLESPACE history;  
# bring the recovered tablespace back online  
SQL "ALTER TABLESPACE history ONLINE";
```

RMAN Tablespace Point-in-Time Recovery (TSPITR)

Recovery Manager (RMAN) automatic **tablespace point-in-time recovery** (commonly abbreviated **TSPITR**) enables you to quickly recover one or more tablespaces in an Oracle database to an earlier time, without affecting the state of the rest of the tablespaces and other objects in the database.

This chapter explains when you can and cannot use TSPITR, what RMAN actually does to your database during TSPITR, how to prepare a database for TSPITR, how to run TSPITR, and options for controlling the TSPITR process.

This chapter contains the following sections:

- [Understanding RMAN TSPITR](#)
- [Planning and Preparing for TSPITR](#)
- [Performing Basic RMAN TSPITR](#)
- [Performing Customized RMAN TSPITR with an RMAN-Managed Auxiliary Instance](#)
- [Performing RMAN TSPITR Using Your Own Auxiliary Instance](#)
- [Troubleshooting RMAN TSPITR](#)


Understanding RMAN TSPITR

In order to use TSPITR effectively, you need to understand what problems it can solve for you, what the major elements used in TSPITR are, what RMAN does during TSPITR, and limitations on when and how it can be applied.

RMAN TSPITR Concepts

[Figure 8–1, "Tablespace Point-in-Time Recovery \(TSPITR\) Architecture"](#) illustrates the context within which TSPITR takes place, and a general outline of the process.

Figure 8–1 Tablespace Point-in-Time Recovery (TSPITR) Architecture



The figure contains the following entities:

- The **target instance**, containing the tablespace to be recovered
- The **Recovery Manager** client
- The **control file** and (optional) **recovery catalog**, used for the RMAN repository records of backup activity
- **Archived redo logs** and **backup sets** from the target database, which are the source of the reconstructed tablespace.
- The **auxiliary instance**, an Oracle database instance used in the recovery process to perform the actual work of recovery.

There are four other important terms related to TSPITR, which will be used in the rest of this discussion:

- The **target time**, the point in time or SCN that the tablespace will be left at after TSPITR
- The **recovery set**, which consists of the datafiles containing the tablespaces to be recovered;
- The **auxiliary set**, which includes datafiles required for TSPITR of the recovery set which are not themselves part of the recovery set. The auxiliary set typically includes:
 - A copy of the SYSTEM tablespace

- Datafiles containing rollback or undo segments from the target instance
- In some cases, a temporary tablespace, used during the export of database objects from the auxiliary instance

The auxiliary instance has other files associated with it, such as a control file, parameter file, and online logs, but they are not part of the auxiliary set.

- The **auxiliary destination**, an optional location on disk which can be used to store any of the auxiliary set datafiles, control files and online logs of the auxiliary instance during TSPITR. Files stored here can be deleted after TSPITR is complete.

All of these terms will be referenced throughout the remainder of this chapter.

How TSPITR Works With an RMAN-Managed Auxiliary Instance

To perform TSPITR of the recovery set using RMAN and an automated auxiliary instance, you carry out the preparations for TSPITR described in "[Planning and Preparing for TSPITR](#)" on page 8-5, and then issue the `RECOVER TABLESPACE` command, specifying, at a minimum, the tablespaces of the recovery set and the target time for the point-in-time recovery, and, if desired, an auxiliary destination as well.

RMAN then carries out the following steps:

1. If there is no connection to an auxiliary instance, RMAN creates the auxiliary instance, starts it up and connects to it.
2. Takes the tablespaces to be recovered offline in the target database
3. Restores a backup control file from a point in time before the target time to the auxiliary instance
4. Restores the datafiles from the recovery set and the auxiliary set to the auxiliary instance. Files are restored either in locations you specify for each file, or the original location of the file (for recovery set files) or in the auxiliary destination (for auxiliary set files, if you used the `AUXILIARY DESTINATION` argument of `RECOVER TABLESPACE`)
5. Recovers the restored datafiles in the auxiliary instance to the specified time
6. Opens the auxiliary database with the `RESETLOGS` option
7. Exports the dictionary metadata about objects in the recovered tablespaces to the target database
8. Shuts down the auxiliary instance
9. Issues `SWITCH` commands on the target instance, so that the target database control file now points to the datafiles in the recovery set that were just recovered at the auxiliary instance.
10. Imports the dictionary metadata from the auxiliary instance to the target instance, allowing the recovered objects to be accessed.
11. Deletes all auxiliary set files.

At that point the TSPITR process is complete. The recovery set datafiles are returned to their contents at the specified point in time, and belong to the target database.

Deciding When to Use TSPITR

Like a table import, RMAN TSPITR enables you to recover a consistent data set; however, the data set recovered includes an entire tablespace rather than one object.

RMAN TSPITR is most useful for situations such as these:

- Recovering data lost after an erroneous `TRUNCATE TABLE` statement;
- Recovering from logical corruption of a table;
- Undoing the effects of an incorrect batch job or other DML statement that has affected only a subset of the database;
- Recovering a logical schema to a point different from the rest of the physical database, when multiple schemas exist in separate tablespaces of one physical database.

Note that, as with database point-in-time recovery (DBPITR), you cannot perform TSPITR if you do not have your archived redo logs. For databases running in `NOARCHIVELOG` mode, you cannot perform TSPITR. You can only restore your entire database from a consistent backup.

Limitations of TSPITR

There are a number of situations which you cannot resolve by using TSPITR.

- You cannot recover dropped tablespaces.
- You cannot recover a renamed tablespace to a point in time before it was renamed. If you try to perform a TSPITR to an SCN earlier than the rename operation, RMAN cannot find the new tablespace name in the repository as of that earlier SCN (because the tablespace did not have that name at that SCN).

In this situation, you must recover the entire database to a point in time before the tablespace was renamed. The tablespace will be found under the name it had at that earlier time.

- You cannot recover tables without their associated constraints, or constraints without the associated tables.
- You cannot use TSPITR to recover any of the following:
 - Replicated master tables
 - Partial tables (for example, if you perform RMAN TSPITR on partitioned tables and spread partitions across multiple tablespaces, then you must recover all tablespaces which include partitions of the table.)
 - Tables with `VARRAY` columns, nested tables, or external files
 - Snapshot logs and snapshot tables
 - Tablespaces containing undo or rollback segments
 - Tablespaces that contain objects owned by `SYS`, including rollback segments

TSPITR has some other limitations:

- If a datafile was added after the point to which RMAN is recovering, an empty datafile by the same name will be included in the tablespace after RMAN TSPITR.
- TSPITR will not recover query optimizer statistics for recovered objects. You must gather new statistics after the TSPITR.
- Assume that you run TSPITR on a tablespace, and then bring the tablespace online at time t . Backups of the tablespace created before time t are no longer usable for recovery with a current control file. You cannot run TSPITR again on this tablespace to recover it to any time less than or equal to time t , nor can you use the current control file to recover the database to any time less than or equal to t . Therefore, you must back up the recovered tablespace as soon as TSPITR is complete.

Limitations of TSPITR Without a Recovery Catalog If you do not use a recovery catalog when performing TSPITR, then note the following special restrictions:

- The undo segments at the time of the TSPITR must be part of the auxiliary set. Because RMAN has no historical record of the undo in the control file, RMAN assumes that the current rollback or undo segments were the same segments present at the time to which recovery is performed. If the undo segments have changed since that time, then TSPITR will fail.
- TSPITR to a time that is too old may not succeed if Oracle has reused the control file records for needed backups. (In planning your database, set the `CONTROL_FILE_RECORD_KEEP_TIME` initialization parameter to a value large enough to ensure that control file records needed for TSPITR are kept.)
- Assume that you run TSPITR on a tablespace, and then bring the tablespace online at time *t*. When not using a recovery catalog, the current control file has no record of the older incarnation of the recovered tablespace. Thus, recovery with a current control file that involves this tablespace can no longer use a backup taken prior to time *t*. You can, however, perform incomplete recovery of the whole database to any time less than or equal to *t*, if you can restore a backup control file from before time *t*.

Planning and Preparing for TSPITR

There are several steps to be carried out in preparing for TSPITR:

- [Choosing the Right Target Time for TSPITR](#)
- [Determining the Recovery Set: Analyzing Data Relationships](#)
- [Identifying and Preserving Objects That Will Be Lost After TSPITR](#)

Choosing the Right Target Time for TSPITR

It is extremely important that you choose the right target time or SCN for your TSPITR. As noted already, once you bring a tablespace online after TSPITR, you cannot use any backup from a time earlier than the moment you brought the tablespace online. In practice, this means that you cannot make a second attempt at TSPITR if you choose the wrong target time the first time, unless you are using a recovery catalog. (If you have a recovery catalog, however, you can perform repeated TSPITRs to different target times.)

For example, assume that you are not using a recovery catalog, and you run TSPITR on a tablespace, and then bring the tablespace online at 5PM on Friday. Backups of the tablespace created before 5PM Friday are no longer usable for recovery with a current control file. You cannot run TSPITR again on this tablespace with a target time earlier than 5PM Friday, nor can you use the current control file to recover the database to any time earlier than 5PM Friday. Your only option will be point-in-time recovery of your entire database using a restored control file.

To investigate past states of your data to identify the target time for TSPITR, you can use features of Oracle such as Oracle Flashback Query, Oracle Transaction Query and Oracle Flashback Version Query to find the point in time when unwanted database changes occurred. See *Oracle Database Backup and Recovery Basics* for more details on Flashback Query, and *Oracle Database Application Developer's Guide - Fundamentals* for more information on Flashback Transaction Query and Flashback Version Query.

Determining the Recovery Set: Analyzing Data Relationships

Your recovery set starts out including the datafiles for the tablespaces you wish to recover. If, however, objects in the tablespaces you need have relationships (such as constraints) to objects in other tablespaces, you will have to account for this relationship before you can perform TSPITR. You have three choices when faced with such a relationship:

- Add the tablespace including the related objects to your recovery set
- Remove the relationship
- Suspend the relationship for the duration of TSPITR

Identifying and Resolving Dependencies on the Primary Database

The `TS_PITR_CHECK` view lets you identify relationships between objects that span the recovery set boundaries. If this view returns rows when queried, then investigate and correct the problem. Proceed with TSPITR *only* when `TS_PITR_CHECK` view returns no rows for the tablespaces not in the recovery set. Record all actions performed during this step so that you can re-create any suspended or removed relationships after completing TSPITR.

The following query illustrates how to use the `TS_PITR_CHECK` view. For an example with an initial recovery set consisting of `tools` and `users`, the `SELECT` statement against `TS_PITR_CHECK` would be as follows:

```
SELECT *
FROM SYS.TS_PITR_CHECK
WHERE (
    TS1_NAME IN ('USERS', 'TOOLS')
    AND TS2_NAME NOT IN ('USERS', 'TOOLS')
)
OR (
    TS1_NAME NOT IN ('USERS', 'TOOLS')
    AND TS2_NAME IN ('USERS', 'TOOLS')
);
```

To run a complete TSPITR check on all the tablespaces in the database (not just the tablespaces in the recovery set), you can run the following query:

```
SELECT *
FROM SYS.TS_PITR_CHECK
WHERE (
    'SYSTEM' IN (TS1_NAME, TS2_NAME)
    AND TS1_NAME <> TS2_NAME
    AND TS2_NAME <> '-1'
)
OR (
    TS1_NAME <> 'SYSTEM'
    AND TS2_NAME = '-1'
);
```

Because of the number and width of the columns in the `TS_PITR_CHECK` view, you may want to format the columns as follows when running the query:

```
SET LINESIZE 120
COLUMN OBJ1_OWNER HEADING "own1"
COLUMN OBJ1_OWNER FORMAT a6
COLUMN OBJ1_NAME HEADING "name1"
COLUMN OBJ1_NAME FORMAT a5
COLUMN OBJ1_SUBNAME HEADING "subname1"
```

```

COLUMN OBJ1_SUBNAME FORMAT a8
COLUMN OBJ1_TYPE HEADING "obj1type"
COLUMN OBJ1_TYPE FORMAT a8 word_wrapped
COLUMN TS1_NAME HEADING "ts1_name"
COLUMN TS1_NAME FORMAT a6
COLUMN OBJ2_NAME HEADING "name2"
COLUMN OBJ2_NAME FORMAT a5
COLUMN OBJ2_SUBNAME HEADING "subname2"
COLUMN OBJ2_SUBNAME FORMAT a8
COLUMN OBJ2_TYPE HEADING "obj2type"
COLUMN OBJ2_TYPE FORMAT a8 word_wrapped
COLUMN OBJ2_OWNER HEADING "own2"
COLUMN OBJ2_OWNER FORMAT a6
COLUMN TS2_NAME HEADING "ts2_name"
COLUMN TS2_NAME FORMAT a6
COLUMN CONSTRAINT_NAME HEADING "cname"
COLUMN CONSTRAINT_NAME FORMAT a5
COLUMN REASON HEADING "reason"
COLUMN REASON FORMAT a25 word_wrapped
    
```

Assume a case in which the partitioned table `tp` has two partitions, `p1` and `p2`, that exist in tablespaces `users` and `tools` respectively. Also assume that a partitioned index called `tpind` is defined on `tp`, and that the index has two partitions `id1` and `id2` (that exist in tablespaces `id1` and `id2` respectively). In this case, you would get the following output when `TS_PITR_CHECK` is queried against tablespaces `users` and `tools` (assuming appropriate formatting):

```

own1  name1 subname1 obj1type ts1_name name2 subname2 obj2type own2      ts2_name  cname reason
---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
SYSTEM TP  P1      TABLE  USER    TPIND IP1      INDEX   PARTITION PARTITION SYS   ID1 Partitioned
Objects not fully contained in the recovery set
SYSTEM TP  P2      TABLE  TOOLS   TPIND IP2      INDEX   PARTITION PARTITION SYS   ID2 Partitioned
Objects not fully contained in the recovery set
    
```

The table `SYSTEM.tp` has a partitioned index `tpind` that consists of two partitions, `ip1` in tablespace `id1` and `ip2` in tablespace `id2`. To perform TSPITR, you must either drop `tpind` or include `id1` and `id2` in the recovery set.

See Also: *Oracle Database Reference* for more information about the `TS_PITR_CHECK` view

Identifying and Preserving Objects That Will Be Lost After TSPITR

When RMAN TSPITR is performed on a tablespace, any objects created after the target recovery time are lost. You can preserve such objects, once they are identified, by exporting them before TSPITR using an Oracle export utility (Data Pump Export or Original Export) and re-importing them afterwards using the corresponding import utility.

To see which objects will be lost in TSPITR, query the `TS_PITR_OBJECTS_TO_BE_DROPPED` view on the primary database. The contents of the view are described in [Table 8–1](#).

Table 8–1 *TS_PITR_OBJECTS_TO_BE_DROPPED View*

Column Name	Meaning
OWNER	Owner of the object to be dropped.
NAME	The name of the object that will be lost as a result of undergoing TSPITR

Table 8–1 (Cont.) TS_PITR_OBJECTS_TO_BE_DROPPED View

Column Name	Meaning
CREATION_TIME	Creation timestamp for the object.
TABLESPACE_NAME	Name of the tablespace containing the object.

Filter the view for objects whose CREATION_TIME is after the target time for TSPITR. For example, with a recovery set consisting of `users` and `tools`, and a recovery point in time of November 2, 2002, 7:03:11 AM, issue the following statement:

```
SELECT OWNER, NAME, TABLESPACE_NAME,
       TO_CHAR(CREATION_TIME, 'YYYY-MM-DD:HH24:MI:SS')
FROM TS_PITR_OBJECTS_TO_BE_DROPPED
WHERE TABLESPACE_NAME IN ('USERS', 'TOOLS')
AND CREATION_TIME > TO_DATE('02-NOV-02:07:03:11', 'YY-MON-DD:HH24:MI:SS')
ORDER BY TABLESPACE_NAME, CREATION_TIME;
(The TO_CHAR and TO_DATE functions are used to avoid issues with different national
date formats. You can, of course, use local date formats in your own work.)
```

See Also: *Oracle Database Reference* for more information about the TS_PITR_OBJECTS_TO_BE_DROPPED view

Performing Basic RMAN TSPITR

Having selected your tablespaces to recover and your target time, you are now ready to perform RMAN TSPITR. You have a few different options available to you:

- **Fully automated TSPITR**--in which you specify an auxiliary destination and let RMAN manage all aspects of the TSPITR. This is the simplest way to perform TSPITR, and is recommended unless you specifically need more control over the location of recovery set files after TSPITR or auxiliary set files during TSPITR, or control over the channel configurations or some other aspect of your auxiliary instance.
- **Customized TSPITR with an automatic auxiliary instance**--in which you base your TSPITR on the behavior of fully automated TSPITR, possibly still using an auxiliary destination, but customize one or more aspects of the behavior, such as the location of auxiliary set or recovery set files, or specifying initialization parameters or channel configurations for the auxiliary instance created and managed by RMAN.
- **TSPITR with your own auxiliary instance**--in which you take responsibility for setting up, starting, stopping and cleaning up the auxiliary instance used in TSPITR, and possibly also manage the TSPITR process using some of the methods available in customized TSPITR with an automatic auxiliary instance.

Fully Automated RMAN TSPITR

When performing fully automated TSPITR, letting RMAN manage the entire process, there are only two requirements beyond the preparations in ["Planning and Preparing for TSPITR"](#) on page 8-5:

- You must specify the auxiliary destination for RMAN to use for the auxiliary set datafiles and other files for the auxiliary instance.
- You must configure any channels required for the TSPITR on the target instance. (The auxiliary instance will use the same channel configuration as the target instance when performing the TSPITR.)

RMAN bases as much of the configuration for TSPITR as possible on your target database. During TSPITR, the recovery set datafiles are written in their current locations on the target database. The same channel configurations in effect on the target database are used on the auxiliary instance when restoring files from backup. Auxiliary set datafiles and other auxiliary instance files, however, are stored in the auxiliary destination.

Using an Auxiliary Destination for Automated RMAN TSPITR

Oracle recommends that you use an auxiliary destination with your auxiliary instance. Even if you use other methods to rename some or all of the auxiliary set datafiles, specifying an `AUXILIARY DESTINATION` parameter provides a default location for auxiliary set datafiles for which names are not specified. This way, TSPITR will not fail if you inadvertently do not provide names for all auxiliary set datafiles.

To specify an auxiliary destination, find a location on disk with enough space to hold your auxiliary set datafiles. Then, use the `AUXILIARY DESTINATION` parameter in your `RECOVER TABLESPACE` command to specify the auxiliary destination location, as shown in the next section.

Performing Fully Automated RMAN TSPITR

To actually perform automated RMAN TSPITR, start the RMAN client, connecting to the target database and, if applicable, a recovery catalog.

Note: Do not connect to an auxiliary instance when starting the RMAN client for automated TSPITR. If RMAN is connected to an auxiliary instance when you run `RECOVER TABLESPACE`, RMAN will assume that you are trying to manage your own auxiliary instance, as described in ["Performing RMAN TSPITR Using Your Own Auxiliary Instance"](#) on page 8-18, and try to use the connected auxiliary for TSPITR.

If you have configured channels that RMAN can use to restore from backup on the primary instance, then you are ready to perform TSPITR now, by running the `RECOVER TABLESPACE... UNTIL...` command.

This example returns the users and tools tablespaces to the end of log sequence number 1300, and stores the auxiliary instance files (including auxiliary set datafiles) in the destination `/disk1/auxdest`:

```
RMAN> RECOVER TABLESPACE users, tools
      UNTIL LOGSEQ 1300 THREAD 1
      AUXILIARY DESTINATION '/disk1/auxdest';
```

Assuming the TSPITR process completes without error, the tablespaces are taken offline by RMAN, restored from backup and recovered to the desired point in time on the auxiliary instance, and then re-imported to the target database. The tablespaces are left offline at the end of the process. All auxiliary set datafiles and other auxiliary instance files are cleaned up from the auxiliary destination.

Tasks to Perform After Successful TSPITR

If TSPITR completes successfully, you must back up the recovered tablespaces, and then you can bring them online.

Backing Up Recovered Tablespaces After TSPITR It is very important that you backup recovered tablespaces immediately after TSPITR is completed.

After you perform TSPITR on a tablespace, you cannot use backups of that tablespace from before the TSPITR was completed and the tablespace put back on line. If you start using the recovered tablespaces without taking a backup, you are running your database without a usable backup of those tablespaces. For this example, the `users` and `tools` tablespaces must be backed up, as follows:

```
RMAN> BACKUP TABLESPACE users, tools;
```

You can then safely bring the tablespaces online, as follows:

```
RMAN> SQL "ALTER TABLESPACE users, tools ONLINE";
```

Your recovered tablespaces are now ready for use.

Handling Errors in Automated TSPITR

In the event of an error during automated TSPITR, you should refer to "[Troubleshooting RMAN TSPITR](#)" on page 8-23. The auxiliary set datafiles and other auxiliary instance files will be left in place in the auxiliary destination as an aid to troubleshooting. The state of the recovery set files is determined by the type of failure. Once you resolve the problem, you can try your TSPITR operation again.

Performing Customized RMAN TSPITR with an RMAN-Managed Auxiliary Instance

There are several aspects of RMAN TSPITR which you can customize while still mostly following the basic procedure described in "[Fully Automated RMAN TSPITR](#)" on page 8-8:

- Renaming or relocating your recovery set datafiles, so that the datafiles making up the recovered tablespaces are not stored in the original locations after TSPITR. (You might do this, for example, if the disk that originally contained the tablespace is unusable.)
- Specifying a location other than the auxiliary destination for some or all auxiliary set datafiles (or not using an auxiliary destination at all). You might do this if there is no single location on disk with enough space for all auxiliary set files.
- Setting up image copy backups of your datafiles in advance, to speed up TSPITR by avoiding restores from backup
- Using a different channel configuration for the auxiliary instance
- Specifying different initialization parameters for your RMAN-managed auxiliary instance

Renaming TSPITR Recovery Set Datafiles with SET NEWNAME

You may not want the recovery set datafiles restored and recovered in their original locations. The `SET NEWNAME` command, used in a `RUN` block, lets you specify a new destination for the restore from backup and recovery of a datafile.

Note: CONFIGURE AUXNAME can be used to rename recovery set datafiles as well, but the effects of doing so are quite different and the two commands cannot be used interchangeably. (They do interact, in that if you use SET NEWNAME to rename a file, this takes precedence over any renaming performed with CONFIGURE AUXNAME.) Refer to the discussion of ["Using Image Copies for Faster RMAN TSPITR Performance"](#) on page 8-15 for details.

Create a RUN block and use SET NEWNAME commands within it to specify new recovery set filenames, as shown here:

```

RUN {
...
  SET NEWNAME FOR DATAFILE 'ORACLE_HOME/oradata/trgt/users01.dbf'
    TO '/newfs/users01.dbf';
  ...other setup commands...
  RECOVER TABLESPACE users, tools UNTIL SEQUENCE 1300 THREAD 1;
}

```

RMAN restores the specified datafile from backup to the new location during TSPITR and recovers it in the new location, and updates the control file so that the newly recovered datafile replaces the old one in the control file. Any existing image copy backup of a datafile found at the new specified location is overwritten.

If the name specified with SET NEWNAME conflicts with the name of a valid datafile in the target database, then RMAN reports an error while executing the RECOVER command. The valid datafile is not overwritten.

Note that RMAN does not detect conflicts between names set with SET NEWNAME and current datafile names on the target database until the actual RECOVER TABLESPACE . . . UNTIL operation. At that point, the conflict is detected, TSPITR fails and RMAN reports an error. If you rename your recovery set datafiles, be sure to assign them names that do not conflict with each other, or with the names of your current datafiles.

Renaming TSPITR Auxiliary Set Datafiles

Unlike the recovery set datafiles, which can be and usually are stored in their original locations, the auxiliary set datafiles must not overwrite the corresponding original files in the target database. If you do not specify a location for an auxiliary set file that is different from its original location, then TSPITR will fail when RMAN attempts to overwrite the corresponding file in the original database and discover that the file is in use.

The simplest way to provide locations for your auxiliary set datafiles is to specify an auxiliary destination for TSPITR. However, RMAN supports two other methods of controlling the location of your auxiliary set datafiles: specifying new names for individual files with SET NEWNAME, and using DB_FILE_NAME_CONVERT to provide rules for converting datafile names in the target database to datafile names for the auxiliary database.

Even if you intend to use either of these methods to provide locations for specific files, it is still suggested that you provide an AUXILIARY DESTINATION argument to RECOVER TABLESPACE. This will ensure that, if you overlook renaming some auxiliary set datafiles, your TSPITR will still succeed. Any files not otherwise renamed will be placed in the auxiliary destination.

Renaming TSPITR Auxiliary Set Datafiles with SET NEWNAME

To use the SET NEWNAME command to specify a new name for an auxiliary set datafile, enclose your RECOVER TABLESPACE command in a RUN block, and use a SET NEWNAME command within the RUN block to rename the file. For example:

```

RMAN> RUN
{
  SET NEWNAME FOR DATAFILE '?/oradata/prod/system01.f'
  TO '/disk1/auxdest/system01.f'
  RECOVER TABLESPACE users, tools
  UNTIL LOGSEQ 1300 THREAD 1
  AUXILIARY DESTINATION '/disk1/auxdest';
}

```

The resulting behavior depends upon whether there is a file at /disk1/auxdest/system01.f when the RECOVER TABLESPACE command is executed. If there is an image copy backup of the file ?/oradata/system01.f at the specified location, created at an SCN prior to the target time for TSPITR, then the behavior is as described in ["SET NEWNAME and CONFIGURE AUXNAME With Auxiliary Set Image Copies"](#) on page 8-16. Otherwise, the auxiliary set datafile will be restored to the NEWNAME specified instead of the default location. If your intention is only to control where the auxiliary set datafiles are stored, you should make sure that there is no file stored at the location specified by SET NEWNAME before performing your TSPITR.

Using DB_FILE_NAME_CONVERT to Name Auxiliary Set Datafiles

If you do not want to use an auxiliary destination for all of your auxiliary set datafiles, but you also do not want to name every file individually, you can include a DB_FILE_NAME_CONVERT initialization parameter in the initialization parameter file used by your auxiliary instance. You can only use this method in two circumstances:

- If you are creating your own initialization parameter file for RMAN's automatically managed auxiliary instance, as described in ["Customizing Initialization Parameters for the Automatic Auxiliary Instance in TSPITR"](#) on page 8-17;
- If you are creating your own auxiliary instance, as described in ["Performing RMAN TSPITR Using Your Own Auxiliary Instance"](#) on page 8-18.

Refer to the appropriate discussion for your circumstance, to see how to add a parameter to your initialization parameter file.

The DB_FILE_NAME_CONVERT parameter in the auxiliary instance specifies how to derive names for files in the auxiliary instance from the original names of the corresponding files in the target instance. The parameter consists of a list of pairs of strings, such that for any filename that contains the first string of a pair as a substring, the name of the corresponding file in the auxiliary instance is generated by substituting the second string of the pair into the original filename.

For example, assume that the target instance contains the following files:

- ?/oradata/trgt/system01.dbf of the SYSTEM tablespace
- ?/oradata/trgt/undotbs01.dbf of the undotbs tablespace

and you need to locate the corresponding files in the auxiliary instance in '/bigtmp', then you would add the following line to the auxiliary instance parameter file:

```
DB_FILE_NAME_CONVERT=('?/oradata/trgt', '/bigtmp')
```

New filenames for the corresponding auxiliary instance files are
 /bigtmp/trgt/system01.dbf and /bigtmp/trgt/undotbs01.dbf.

The most important thing to remember is that `DB_FILE_NAME_CONVERT` **must** be present in the auxiliary instance parameter file.

If the auxiliary instance was manually created, add `DB_FILE_NAME_CONVERT` to the auxiliary instance parameter file.

Note that you can still rename individual auxiliary set datafiles using `SET NEWNAME` or `CONFIGURE AUXNAME`. Also, files that do not match the patterns provided in `DB_FILE_NAME_CONVERT` will not be renamed. You may wish to use the `AUXILIARY DESTINATION` parameter of `RECOVER TABLESPACE` to ensure that all auxiliary set datafiles are sent to some destination. If none of the renaming methods used provide a new name for a file at the auxiliary instance, TSPITR will fail.

Renaming ASM OMF Datafiles Using `DB_FILE_NAME_CONVERT` in TSPITR The `DB_FILE_NAME_CONVERT` initialization parameter cannot be used to control generation of new names for files at the auxiliary instance which are Oracle Managed Files (OMF) at the target instance. When using Oracle Managed Files at the target instance, it is not possible to generate valid OMF filenames for the auxiliary instance by replacing a substring of the target instance OMF filename. When using ASM Oracle Managed Files, RMAN will change these invalid names to valid filenames in the converted disk group name.

To avoid this issue, use one of the other supported options for generating new names for OMF files (including files stored in ASM):

- Use an auxiliary destination, as described in ["Using an Auxiliary Destination for Automated RMAN TSPITR"](#) on page 8-9.
- Use the `DB_CREATE_FILE_DEST` initialization parameter in the auxiliary instance parameter file to specify a location for all auxiliary instance files for which no new name is specified using `SET NEWNAME` or `CONFIGURE AUXNAME`.
- For ASM files, you can use `SET NEWNAME` to redirect individual files to a specific disk group accessible from the auxiliary instance (and allow the database to generate the filename within the disk group). For example:

```

RUN {
  SET NEWNAME FOR DATAFILE 1 TO "+DISK2";
  SET NEWNAME FOR DATAFILE 2 TO "+DISK3";
  RECOVER TABLESPACE users, tools
    UNTIL LOGSEQ 1300 THREAD 1
    AUXILIARY DESTINATION '/disk1/auxdest';
}

```

Renaming of Tempfiles During TSPITR Tempfiles are considered part of the auxiliary set for your database. When creating the auxiliary instance, tempfiles can be renamed using `SET NEWNAME FOR TEMPFILE`, `DB_FILE_NAME_CONVERT` or `AUXILIARY DESTINATION`. When the auxiliary instance is opened, the tempfiles are recreated according to the applicable renaming rule. When the auxiliary instance is cleaned up, the tempfiles are deleted along with the rest of the auxiliary instance files.

Order of Precedence Among File Renaming Methods in TSPITR

The different methods of renaming files follow an order of precedence, as follows:

- `SET NEWNAME`

- `CONFIGURE AUXNAME`
- `DB_FILE_NAME_CONVERT`
- `AUXILIARY DESTINATION` argument to `RECOVER TABLESPACE`

Settings higher on the list override settings lower on the list, in situations where both have been applied (by, for example, running `RECOVER TABLESPACE . . . AUXILIARY DESTINATION` on a target database where some auxiliary set datafiles also have auxnames configured with `CONFIGURE AUXNAME`).

Note: You can view any current `CONFIGURE AUXNAME` settings using the `SHOW AUXNAME` command, described in *Oracle Database Backup and Recovery Reference*.

Specifying Auxiliary Instance Control File Location in TSPITR

You can specify your own location for the control file of your auxiliary instance, if you use a client-side initialization parameter file. Set the `CONTROL_FILES` initialization parameter to specify any location you wish for the control files.

If you do not explicitly specify a location for the control file, RMAN will locate it in the auxiliary destination if you use the `AUXILIARY DESTINATION` parameter when performing TSPITR. If you do not use an `AUXILIARY DESTINATION` parameter, the auxiliary instance control files are stored in an operating system-specific location. (on Unix, `ORACLE_HOME/rdbms/admin/params_auxinit.ora`).

No matter where you store your auxiliary instance control file, it is removed at the end of a successful TSPITR operation. Because control files are relatively small, it is rare that RMAN will encounter a problem creating an auxiliary control file, but if there is not enough space to create the control file, TSPITR will fail.

Specifying Auxiliary Instance Online Log Location in TSPITR

If you specify the `LOG_FILE_NAME_CONVERT` initialization parameter in your auxiliary instance parameter file, this parameter will determine the online redo log location. Otherwise, if RMAN is using an auxiliary destination and managing the auxiliary instance for you, it creates the online redo log in the auxiliary destination.

Note: If you do not specify a location for the online redo logs using `LOG_FILE_NAME_CONVERT` or `AUXILIARY DESTINATION`, your TSPITR will fail trying to create the online redo logs. Even if `DB_FILE_CREATE_DEST` or `LOG_FILE_CREATE_DEST` are specified in the initialization parameter file, in TSPITR they do not control the creation of the online redo logs of the auxiliary instance.

Renaming ASM OMF Redo Logfiles with `LOG_FILE_NAME_CONVERT` in TSPITR `LOG_FILE_NAME_CONVERT` cannot be used to control generation of new names for redo log files at the auxiliary instance which are Oracle Managed Files (OMF) at the target instance. When using Oracle Managed Files at the target instance, it is not possible to generate valid OMF filenames for the auxiliary instance by replacing a substring of the target instance OMF filename. When using ASM Oracle Managed Files, RMAN uses the pattern to convert the disk group name only, and generates a valid filename in the converted disk group.

To avoid this issue, use one of the other supported methods of generating new names for OMF redo log files (including files stored in ASM):

- Use an auxiliary destination, as described in ["Using an Auxiliary Destination for Automated RMAN TSPITR"](#) on page 8-9
- Use the `DB_CREATE_FILE_DEST`, `DB_RECOVERY_FILE_DEST` or `DB_CREATE_ONLINE_LOG_DEST_n` initialization parameter in the auxiliary instance parameter file to specify a location

Using Image Copies for Faster RMAN TSPITR Performance

TSPITR performance can be greatly enhanced by redirecting RMAN to use existing image copies of the recovery set and auxiliary set datafiles on disk in TSPITR, rather than restoring them from backup. You can use the `CONFIGURE AUXNAME` command with image copies of recovery set datafiles or auxiliary set datafiles, or the `SET NEWNAME` command with image copies of auxiliary set datafiles, to tell RMAN about the possible existence of an image copy of a datafile.

While exact details vary depending on the command used and whether the file is an auxiliary set or recovery set file, in general, if a suitable image copy is available in the specified location, then during TSPITR, RMAN uncatalogs the image copy from the RMAN repository of the target instance, and catalogs it in the control file of the auxiliary instance. The auxiliary instance then performs point-in-time recovery using the image copy.

Details of using image copies with each type of file are explained in the following sections.

Using CONFIGURE AUXNAME With Recovery Set Image Copies in TSPITR

During TSPITR, RMAN looks in the specified `AUXNAME` location for the datafile, to see whether the file there is an image copy backup of the datafile, with a checkpoint SCN early enough that it can be recovered to the target time for TSPITR. If such an image copy is found, it is used in TSPITR. Otherwise, the datafile is restored and recovered in its original location, and any file in the location specified by the `AUXNAME` is not changed or deleted.

```
RMAN> CONFIGURE AUXNAME FOR DATAFILE 'ORACLE_HOME/oradata/trgt/users01.dbf'
      TO '/newfs/users1.dbf';
...other RMAN commands, if any...
RMAN> RECOVER TABLESPACE users, tools UNTIL SEQUENCE 1300 THREAD 1;
```

The primary use of `CONFIGURE AUXNAME` is as the basis of a technique to make TSPITR faster by eliminating restore times. If you have tablespaces on which you anticipate performing TSPITR, you can include in your backup routine the maintenance of a set of image copies of the affected datafiles, and update these periodically to the earliest point to which you expect to perform TSPITR. The expected usage model is:

- Configure the `AUXNAME` for the files once, when setting up this strategy
- Perform `"BACKUP AS COPY DATAFILE n FORMAT auxname"` regularly to maintain the updated image copy, or, for better performance, use an incrementally updated backups strategy as described in *Oracle Database Backup and Recovery Basics* to keep the image copies up to date without performing full backups of the datafiles
- When TSPITR is needed, specify a target time since the last update of the image copy.

In planning for TSPITR with image copies, remember that you may not know which tablespaces will require image copies in advance. As discussed in ["Determining the](#)

[Recovery Set: Analyzing Data Relationships](#)" on page 8-6, relationships between the tablespaces you wish to TSPITR and other tablespaces may require that you add tablespaces to your final recovery set, and still other tablespaces may wind up in the auxiliary set. You should configure an AUXNAME for each datafile that is likely to be part of your recovery set, and update image copies of all datafiles often.

If you do not correctly anticipate all tablespaces to include in the recovery set, or if for reasons of overhead you do not want to maintain copies of all possible recovery set tablespaces, you can prepare only a subset of the datafiles using this strategy. The TSPITR process is still the same if only a subset of datafiles are prepared. The process takes longer, because RMAN must recover recovery set datafiles for which there are no image copies in their original locations.

Note that the order of precedence of naming methods is still respected when you use `CONFIGURE AUXNAME` to rename a recovery set file. A `SET NEWNAME` for a recovery set file set as part of the TSPITR command overrides the effect of the persistent `CONFIGURE AUXNAME` command for the same file. Behavior in this instance will be as described in ["Renaming TSPITR Recovery Set Datafiles with SET NEWNAME"](#) on page 8-10. `SET NEWNAME` used with a recovery set file never refers to an image copy file.

SET NEWNAME and CONFIGURE AUXNAME With Auxiliary Set Image Copies

As with recovery set datafiles, `CONFIGURE AUXNAME` sets a persistent alternative location for an auxiliary set datafile image copy, and `SET NEWNAME` sets an alternative location for the duration of a RUN block. However, RMAN handles values for auxiliary set datafiles differently from recovery set datafiles.

If `SET NEWNAME` is used to specify a new location for an auxiliary set datafile, and there is an image copy at that location with an SCN such that it can be used in TSPITR, then the image copy will be used. If there is no usable image copy at that location, however, RMAN will restore a usable copy from backup. (If an image copy is present but the SCN is after the target time for TSPITR, then the datafile is overwritten by the restored file.)

If `CONFIGURE AUXNAME` is used to specify a new location for an auxiliary set datafile, and there is an image copy at that location with an SCN such that it can be used in TSPITR, then the image copy will be used. If there is no usable copy at the specified location, the file is restored to this location from backup.

As with all auxiliary set files, the file is deleted after successful TSPITR, or left for use in troubleshooting if TSPITR fails, regardless of whether it was an image copy created before TSPITR or restored by RMAN from backup during TSPITR.

Note: You can view any current `CONFIGURE AUXNAME` settings using the `SHOW AUXNAME` command, described in *Oracle Database Backup and Recovery Reference*.

Performing TSPITR With CONFIGURE AUXNAME and Image Copies: Scenario

You have enough disk space to save image copies of your entire database for use in TSPITR. In preparation for the possibility that you need perform TSPITR, you perform the following tasks:

- Configure an AUXNAME for each datafile in your database using:

```
CONFIGURE AUXNAME FOR DATAFILE n TO auxname_n;
```


- Every week on Sunday night you take an image copy of the database which is placed in the files referenced by the configured AUXNAMES:

```
BACKUP AS COPY DATAFILE n FORMAT auxname_n
```

Note that if the image copies are all in the same location on disk and named similarly to the original datafiles, it is possible to use `FORMAT` or `DB_FILE_NAME_CONVERT` options of the `BACKUP` command and use `BACKUP AS COPY DATABASE` instead of performing individual backups of every datafile. For example if the configured AUXNAMES are a simple translation of the location 'maindisk' to 'auxdisk', you could use the following backup command:

```
BACKUP AS COPY DATABASE DB_FILE_NAME_CONVERT=(maindisk, auxdisk);
```

Note: Because OMF filenames cannot generally be translated using a simple substitution, the technique of using `DB_FILE_NAME_CONVERT` to generate names for these image copies cannot generally be used if the image copies to be used are stored in OMF.

You are then prepared for TSPITR without restoring from backup. If, for example, an erroneous batch job started at November 15 2003, 19:00:00 updates incorrectly the tables in the tablespace `parts`, you could use the following command to perform TSPITR on tablespace `parts`:

```
RECOVER TABLESPACE parts UNTIL TIME 'November 15 2003, 19:00:00';
```

Because AUXNAMES are configured and refer to datafile copies from an SCN before the TSPITR target time, the auxiliary set and recovery set datafiles are not restored from backup. Instead the datafile copies are directly used in recovery, eliminating the restore overhead.

Note that at the end of the TSPITR, the datafiles for tablespace `parts` will not be located in the original datafile locations, but in the `auxname` locations. If only the AUXNAMES for the auxiliary set should be used (so that the recovery set is left in its original locations), then `CONFIGURE AUXNAME ... CLEAR` should be used before TSPITR is started. In such a case, though, note that the datafiles will have to be restored.

Customizing Initialization Parameters for the Automatic Auxiliary Instance in TSPITR

The automatic auxiliary instance looks for initialization parameters in a file that is operating system dependent (for Unix this location is `?/rdbms/admin/params_auxint.ora`, where '?' stands for `ORACLE_HOME`, and the file is located on the node running the RMAN client, not necessarily the same node as the one running the database instances). RMAN always looks for this default parameter file for the automatic auxiliary instance when performing TSPITR. If the file is not found RMAN does not generate an error.

Another way to specify parameters for the automatic auxiliary instance is to place the initialization parameters in a file, and then provide the location of these file with the `SET AUXILIARY INSTANCE PARAMETER FILE` command before executing TSPITR. (Note that the path specified when using `SET AUXILIARY INSTANCE PARAMETER` is a path on the system running the RMAN client, not the target or auxiliary instances.)

RMAN defines the following basic parameters for the automatic auxiliary instance:

- `DB_NAME` - Same as `db_name` of the target database

- `DB_UNIQUE_NAME` - Generated, based on the `DB_NAME`, to be unique
- `DB_BLOCK_SIZE` - Same as the `DB_BLOCK_SIZE` of the target database
- `COMPATIBLE` - Same as the compatible setting of the target database

If `AUXILIARY DESTINATION` is used, RMAN also defines:

- `DB_CREATE_FILE_DEST` - Set to the auxiliary destination
- `CONTROL_FILES` - Generated filename in the auxiliary destination

When an auxiliary destination is specified, RMAN uses these two parameters in creating the auxiliary instance online logs and control files in the auxiliary destination.

If `AUXILIARY DESTINATION` is not used, then you must use `LOG_FILE_NAME_CONVERT` in an auxiliary instance parameter file to specify the online log file names. Otherwise, TSPITR fails when attempting to create the online logs for the automatic instance.

If `AUXILIARY DESTINATION` is not used and you do not use `CONTROL_FILES` in an auxiliary instance parameter file, the auxiliary instance will create one control file with an operating system-dependent name in an operating system dependent location. (In Unix, it defaults to `?/dbs/cntrl_@.dbf`, where '?' stands for `ORACLE_HOME` and '@' stands for `ORACLE_SID`. For an automatic auxiliary instance, `ORACLE_SID` is randomly generated by RMAN).

It is rarely necessary, however, to alter the parameter file, especially if you provide an `AUXILIARY DESTINATION` argument to `RECOVER TABLESPACE`. If you override one of the six basic initialization parameters in the auxiliary instance parameter file with an inappropriate value, TSPITR may fail due to problems with the auxiliary instance. However, other parameters besides these basic parameters can be added if needed. For example, you can use `DB_FILE_NAME_CONVERT` to specify the names of the datafiles in the auxiliary set.

Performing RMAN TSPITR Using Your Own Auxiliary Instance

Oracle recommends that you allow RMAN to manage the creation and destruction of the auxiliary instance used during RMAN TSPITR. However, creating and using your own auxiliary instance is also supported. One reason you might want to do this is to exercise control of channels used in TSPITR. RMAN's automatic auxiliary instance uses the configured channels of the target database as the basis for the channels to configure on the auxiliary instance and use during restore. If you need different channel settings, and you do not want to use `CONFIGURE` to change the settings on the target database, you can operate your own auxiliary instance.

Preparing Your Own Auxiliary Instance for RMAN TSPITR

Creating an Oracle instance suitable for use as an auxiliary instance requires that you carry out all of the following steps:

- [Step 1: Create an Oracle Password File for the Auxiliary Instance](#)
- [Step 2: Create an Initialization Parameter File for the Auxiliary Instance](#)
- [Step 3: Check Oracle Net Connectivity to the Auxiliary Instance](#)

Step 1: Create an Oracle Password File for the Auxiliary Instance

For instructions on how to create and maintain Oracle password files, refer to the *Oracle Database Administrator's Guide*.

Step 2: Create an Initialization Parameter File for the Auxiliary Instance

Create a client-side initialization parameter file for the auxiliary instance on the machine where you will be running SQL*Plus to control the auxiliary instance. For this example, assume your parameter file is placed at `/tmp/initAux.ora`. Set the parameters described in the following table.

Note: Paths used in parameters such as `DB_FILE_NAME_CONVERT`, `LOG_FILE_NAME_CONVERT` and `CONTROL_FILES` all refer to server-side paths, not client-side.

Table 8–2 Initialization Parameters in the Auxiliary Instance

Parameter	Mandatory?	Value
<code>DB_NAME</code>	YES	The same name as the target database.
<code>DB_UNIQUE_NAME</code>	YES	A value different from any database in the same Oracle home. For simplicity, specify <code>_dbname</code> . For example, if the target database name is <code>trgt</code> , then specify <code>_trgt</code> .
<code>REMOTE_LOGIN_PASSWORDFILE</code>	YES	Set to <code>EXCLUSIVE</code> when connecting to the auxiliary instance by means of a password file. Otherwise, set to <code>NONE</code> .
<code>COMPATIBLE</code>	YES	The same value as the parameter in the target database.
<code>DB_BLOCK_SIZE</code>	YES	If this initialization parameter is set in the target database, then it must be set to the same value in the auxiliary instance.
<code>LOG_FILE_NAME_CONVERT</code>	YES	<p>Patterns to generate filenames for the online redo logs of the auxiliary database based on the online redo log names of the target database. Query <code>V\$LOGFILE.MEMBER</code>, to obtain target instance online log names, and ensure that the conversion pattern matches the format of the filename displayed in the view.</p> <p>This parameter is the only way to name the online redo logs for the auxiliary instance. Without it, TSPITR will fail when trying to open the auxiliary instance because the online logs cannot be created.</p> <p>Note: Some platforms do not support ending patterns in a forward or backward slash (<code>\</code> or <code>/</code>).</p> <p>Note: Some platforms do not support ending patterns in a forward or backward slash (<code>\</code> or <code>/</code>).</p> <p>See Also: "Specifying Auxiliary Instance Online Log Location in TSPITR" on page 8-14 for restrictions on possible values for <code>LOG_FILE_NAME_CONVERT</code> with OMF filenames</p>

Table 8–2 (Cont.) Initialization Parameters in the Auxiliary Instance

Parameter	Mandatory?	Value
DB_FILE_NAME_CONVERT	NO	Patterns to convert filenames for the datafiles of the auxiliary database. You can use this parameter to generate filenames for those files that you did not name with SET NEWNAME or CONFIGURE AUXNAME. Obtain the datafile filenames by querying V\$DATAFILE.NAME, and ensure that the conversion pattern matches the format of the filename displayed in the view. You can also specify this parameter on the RECOVER command itself. Note: Some platforms do not support ending patterns in a forward or backward slash (\ or /). See Also: "Using DB_FILE_NAME_CONVERT to Name Auxiliary Set Datafiles" on page 8-12
CONTROL_FILES	NO	Filenames that do not conflict with the control file names of the target instance (or any other existing file).

Set other parameters as needed, including the parameters that allow you to connect as SYSDBA through Oracle Net.

The following example shows possible initialization parameter settings for an auxiliary instance for TSPITR:

```
DB_NAME=trgt
DB_UNIQUE_NAME=_trgt
CONTROL_FILES=/tmp/control01.ctl
DB_FILE_NAME_CONVERT=('/oracle/oradata/trgt/', '/tmp/')
LOG_FILE_NAME_CONVERT=('/oracle/oradata/trgt/redo', '/tmp/redo')
REMOTE_LOGIN_PASSWORDFILE=exclusive
COMPATIBLE =10.1.0
DB_BLOCK_SIZE=8192
```

Note: After setting these initialization parameters, ensure that you do not overwrite the initialization settings for the production files at the target database.

See Also: *Oracle Database Net Services Administrator's Guide* for more information about Oracle Net

Step 3: Check Oracle Net Connectivity to the Auxiliary Instance

The auxiliary instance must have a valid net service name. Before proceeding, use SQL*Plus to ensure that you can establish a connection to the auxiliary instance.

Preparing RMAN Commands for TSPITR with Your Own Auxiliary Instance

If you are running your own auxiliary instance, then you may find that the sequence of commands required for TSPITR is quite long, if you allocate a complex channel configuration for restoring from backup, or if you are not using DB_FILE_NAME_CONVERT to control file naming.

You may wish to store the sequence of commands for TSPITR in a command file, a text file under the host operating system. This command file can be read into RMAN using

the @ command (or the `CMDFILE` command line argument when starting RMAN) to execute the series of commands in the command file.

See ["Using RMAN with Command Files"](#) on page 1-3 for more details.

Planning Channels for TSPITR with Your Own Auxiliary Instance

When you run your own auxiliary instance, the default behavior is to use the automatic channel configuration of the target instance. However, if you decide to allocate your own channel configuration, you can do so by including the `ALLOCATE AUXILIARY CHANNEL` commands in a `RUN` block along with the `RECOVER TABLESPACE` command for TSPITR. Plan out these commands, if necessary, and add them to the sequence of commands you will run to perform your TSPITR.

See the example in ["Executing TSPITR With Your Own Auxiliary Instance: Scenario"](#) on page 8-22 for details of how to include channel allocation in your TSPITR script.

Planning Datafile Names with Your Own Auxiliary Instance: SET NEWNAME

You may wish to use `SET NEWNAME` commands, either to refer to existing image copies of auxiliary set files to improve TSPITR performance, or to assign new names to the recovery set files for after TSPITR. Plan out these commands, if necessary, and add them to the sequence of commands you will run to perform your TSPITR.

Executing TSPITR with Your Own Auxiliary Instance

With the preparations complete and your TSPITR commands completely planned, you are now ready to carry out your TSPITR. The following steps are required:

- [Step 1: Start the Auxiliary Instance in NOMOUNT Mode](#)
- [Step 2: Connect the RMAN Client to Target and Auxiliary Instances](#)
- [Step 3: Execute the RECOVER TABLESPACE Command](#)

Step 1: Start the Auxiliary Instance in NOMOUNT Mode

Before beginning RMAN TSPITR, use SQL*Plus to connect to the auxiliary instance and start it in `NOMOUNT` mode, specifying a parameter file if necessary. For example:

```
SQL> CONNECT SYS/oracle@aux AS SYSDBA
SQL> STARTUP NOMOUNT PFILE='/tmp/initAux.ora'
```

Remember that the path for the `PFILE` will be a client-side path, on the machine from which you run SQL*Plus, not a server-side path.

Because the auxiliary instance does not yet have a control file, you can only start the instance in `NOMOUNT` mode. Do not create a control file or try to mount or open the auxiliary instance for TSPITR.

Step 2: Connect the RMAN Client to Target and Auxiliary Instances

Start RMAN connecting to the target and the manually created auxiliary instance:

```
% rman target / auxiliary sysuser/syspwd@auxiliary_service_name
```

Step 3: Execute the RECOVER TABLESPACE Command

Now you are ready to run your TSPITR commands. In the simplest case, just execute the `RECOVER TABLESPACE . . . UNTIL` command at the RMAN prompt:

```
RMAN> RECOVER TABLESPACE ts1, ts2... UNTIL TIME 'time'
```

If you want to use `ALLOCATE CHANNEL` or `SET NEWNAME` then create a `RUN` block which includes those commands before the `RECOVER TABLESPACE` command.

```
RUN {
  ALLOCATE CHANNEL c1 DEVICE TYPE DISK;
  ALLOCATE CHANNEL c2 DEVICE TYPE SBT;
  # and so on...
  RECOVER TABLESPACE ts1, ts2 UNTIL TIME 'time';
}
```

Using a Command File for TSPITR Entering a lengthy series of commands in a `RUN` block can be error-prone. To avoid making mistakes entering the sequence of commands, create a command file (called, for example, `/tmp/tspitr.rman`) to store the whole sequence of commands for your TSPITR. Review it carefully to catch any errors. Then run the command file at the RMAN prompt, using this command:

```
RMAN> @/tmp/tspitr.rman ;
```

The results will be the same as in the previous example.

Executing TSPITR With Your Own Auxiliary Instance: Scenario

The following example shows the execution of a `RECOVER TABLESPACE... UNTIL` operation using the following features of RMAN TSPITR:

- Managing your own auxiliary instance
- Configuring channels for restore of backups from disk and sbt
- Using recoverable image copies for some auxiliary set datafiles using `SET NEWNAME`
- Specifying new names for recovery set datafiles using `SET NEWNAME`

The process used is as follows:

1. Prepare the auxiliary instance as described in "[Preparing Your Own Auxiliary Instance for RMAN TSPITR](#)" on page 8-18. Specify "tspitr" as the password for the auxiliary instance in the password file, and set up the auxiliary instance parameter file `/bigtmp/init_tspitr_prod.ora` with the following settings:

```
db_name=PROD
db_unique_name=tspitr_PROD
control_files=/bigtmp/tspitr_cntrl.f'
db_file_name_convert=('?/oradata/prod', '/bigtmp')
log_file_name_convert=('?/oradata/prod', '/bigtmp')
compatible=10.1.0
block_size=8192
remote_login_password=exclusive
```

2. Create service name `pitprod` for the auxiliary instance, and check for connectivity.
3. Start the auxiliary instance in `NOMOUNT` state, as shown:

```
$ sqlplus
SQL> connect sys/tspitr@pit_prod as sysdba
SQL> startup nomount pfile=/bigtmp/init_tspitr_prod.ora
```

4. Start up RMAN, connecting to the auxiliary instance:

```
% rman target / auxiliary sys/tspitr@pit_prod
```

5. Enter the following commands, in a RUN block, to set up and execute the TSPITR:

```
run {
# Specify NEWNAMES for recovery set datafiles
  SET NEWNAME FOR DATAFILE '?/oradata/prod/clients01.f'
                                TO '?/oradata/prod/clients01_rec.f';
  SET NEWNAME FOR DATAFILE '?/oradata/prod/clients02.f'
                                TO '?/oradata/prod/clients02_rec.f';
  SET NEWNAME FOR DATAFILE '?/oradata/prod/clients03.f'
                                TO '?/oradata/prod/clients03_rec.f';
  SET NEWNAME FOR DATAFILE '?/oradata/prod/clients04.f'
                                TO '?/oradata/prod/clients04_rec.f';

# Specified newnames for some of the auxiliary set
# datafiles that have a valid image copy to avoid restores:
  SET NEWNAME FOR DATAFILE '?/oradata/prod/system01.f'
                                TO '/backups/prod/system01_monday_noon.f';
  SET NEWNAME FOR DATAFILE '?/oradata/prod/system02.f'
                                TO '/backups/prod/system02_monday_noon.f';
  SET NEWNAME FOR DATAFILE '?/oradata/prod/undo01.f'
                                TO '/backups/prod/undo01_monday_noon.f';

# Specified the disk and SBT channels to use
  allocate auxiliary channel c1 device type disk;
  allocate auxiliary channel c2 device type disk;
  allocate auxiliary channel t1 device type sbt;
  allocate auxiliary channel t2 device type sbt;

# Recovered the clients tablespace to 24 hours ago:
  RECOVER TABLESPACE clients UNTIL TIME 'sysdate-1';
}
```

Consider storing this command sequence in a command file and executing the command file, to avoid errors.

If the TSPITR operation is successful, then the results are:

- The recovery set datafiles are registered in the target database control file under the names specified with SET NEWNAME, with their contents as of the time specified time for the TSPITR.
- The auxiliary files are removed by RMAN, including the control files, online logs and auxiliary set datafiles of the auxiliary instance
- The auxiliary instance is shut down

If the TSPITR operation fails, the auxiliary files are left on disk for troubleshooting purposes. If RMAN created the auxiliary instance, it is shut down; otherwise it is left in whatever state it was in when the TSPITR operation failed.

Troubleshooting RMAN TSPITR

A variety of problems can cause TSPITR to fail before the process is complete.

- There can be name conflicts between files already in the target database, filenames assigned by the SET NEWNAME or CONFIGURE AUXNAME commands, and filenames generated by the effect of the DB_FILE_NAME_CONVERT parameter.
- When RMAN exports the metadata about recovered objects from the auxiliary instance, it uses space in the temporary tablespace for sorting. If there is

insufficient space in the temporary tablespace for the sorting operation, you need to increase the amount of sort space available.

Troubleshooting TSPITR: Filename Conflicts

If your uses of SET NEWNAME, CONFIGURE AUXNAME and DB_FILE_NAME_CONVERT cause multiple files in the auxiliary or recovery sets to have the same name, RMAN will report an error during TSPITR. To correct the problem, use different values for these parameters to eliminate the duplicate name.

Troubleshooting TSPITR: Insufficient Sort Space during Export

In this case, you need to edit the `recover.bsq` file, wherever it resides on your host platform. For instance, on UNIX, it is located in `$ORACLE_HOME/rdbms/admin`. This file contains the following:

```
#
# tspitr_7: do the incomplete recovery and resetlogs. This member is used once.
#
define tspitr_7
<<<
# make the control file point at the restored datafiles, then recover them
RECOVER CLONE DATABASE TABLESPACE &1&;
ALTER CLONE DATABASE OPEN RESETLOGS;
# PLUG HERE the creation of a temporary tablespace if export fails due to lack
# of temporary space.
# For example in Unix these two lines would do that:
# sql clone "create tablespace aux_tspitr_tmp
#          datafile '/tmp/aux_tspitr_tmp.dbf' size 500K";
}
>>>
```

Remove the '#' symbols from the last two lines of comments and modify the statement to create a temporary tablespace. Retry the TSPITR operation, increasing the size of the tablespace until the export operation succeeds.

Troubleshooting TSPITR: RMAN Cannot Identify Tablespaces with Undo Segments

During TSPITR, RMAN needs information about which tablespaces had undo segments at the TSPITR target time. This information is usually available in the recovery catalog, if one is used.

If there is no recovery catalog, or if the information is not found in the recovery catalog, RMAN proceeds assuming that the set of tablespaces with undo segments at the target time is the same as the set of tablespaces with undo segments at the present time. If this assumption is not correct, TSPITR will fail with an error. In such a case, use the UNDO TABLESPACE clause to provide a list of tablespaces with undo segments at the target time.

Troubleshooting: Restarting Manual Auxiliary Instance After TSPITR Failure

If you are managing your own auxiliary instance and there is a failure in TSPITR, then before you can try TSPITR again, you must shut down the auxiliary instance, correct the problem which interfered with TSPITR, and then bring the auxiliary instance back to NOMOUNT before trying TSPITR again.

RMAN Backup and Repository Maintenance

This chapter provides conceptual and procedural information related to maintenance tasks associated with using the Recovery Manager (RMAN) utility, including maintaining the RMAN repository and deleting backups.

This chapter contains the following topics:

- [RMAN Reporting](#)
- [Crosschecks of RMAN Backups](#)
- [Deleting RMAN Backups](#)
- [CHANGE... AVAILABLE and UNAVAILABLE with RMAN Backups](#)
- [Changing Retention Policy Status of RMAN Backups](#)
- [Monitoring RMAN Through V\\$ Views](#)

RMAN Reporting

The RMAN repository for a database contains extensive records of backups of the database, as well as other useful information such as database schema and configuration settings. You can use RMAN commands `LIST`, `REPORT`, and `SHOW` to access this repository information.

In addition to these general reporting commands, you can also make use of the `RESTORE... PREVIEW` command to see which backup files are required to restore specific database objects from backup. See *Oracle Database Backup and Recovery Basics* for more details on `RESTORE... PREVIEW`.

Using the RMAN LIST Command

The `LIST` command is used to query the RMAN repository and obtain data about:

- Backup sets and image copies generated by the `RMAN BACKUP` command;
- Specified objects contained in the `BACKUP`-generated files, that is, archived logs, datafiles, control files, and server parameter files;
- Incarnations of a specified database, or of all databases known to a recovery catalog.

RMAN `LIST` output is sent either to standard output or to the message log (though not to both at the same time). You can also control how the output is organized as well as the level of detail in the output.

You can also list backups by querying `V$BACKUP_FILES` and the `RC_BACKUP_FILES` recovery catalog view. These views provide access to the same information as the `LIST BACKUPSET` command.

The `LIST` command displays the same files that the `CROSSCHECK` and `DELETE` commands operate on. Consequently, you can issue `LIST` to see what is in the repository, and then run `CROSSCHECK` to ensure that these files exist on disk or tape.

See Also:

- *Oracle Database Backup and Recovery Basics* to learn how to use `LIST`
- ["Querying Recovery Catalog Views"](#) on page 10-22 to learn how to use views as an alternative to `LIST`
- *Oracle Database Backup and Recovery Reference* for `LIST` command syntax
- *Oracle Database Backup and Recovery Reference* for `LOG` command-line syntax

RMAN Reports

RMAN reports are intended to provide analysis of your backup and recovery situation. An RMAN report can answer questions such as:

- Which datafiles need a backup?
- Which backups are obsolete because they are redundant or because they are not needed for recovery within a recovery window?
- Are any datafiles now unrecoverable because they have been the target of unrecoverable operations?
- What is the current physical schema of the database, or what was it at some previous time?

RMAN's reporting can be used to monitor and validate your ongoing backup strategy. The `REPORT NEED BACKUP` and `REPORT UNRECOVERABLE` commands let you ensure that the necessary backups are available for media recovery, and that you can perform media recovery within a reasonable amount of time.

Note: A datafile that does not have a backup is still considered recoverable by RMAN, as long as a complete set of archived redo logs is available, from the time the datafile was created to the present. During recovery, an empty datafile is created, and then all of the changes to the datafile from the archived redo logs are applied to reconstruct the full contents of the file.

Reports of Obsolete Backups

The `REPORT OBSOLETE` command displays backups of datafiles, control files, and archived redo logs that can be deleted because they are no longer needed. Backups are determined to be obsolete with respect to a retention policy, based either on redundancy or a required recovery window, as described in the following table:

Retention Policy	Meaning
REDUNDANCY <i>integer</i>	At least <i>integer</i> more recent backups of this file already exist.
RECOVERY WINDOW <i>integer</i>	The backup is not needed for recovery to any point within the recovery window of <i>integer</i> days. For each datafile, one backup that is older than the recovery window must exist, because point-in-time recovery to the beginning of the recovery window requires must begin with restoring the database from such a backup. In other words, one backup of each datafile must satisfy the condition <code>SYSDATE - CHECKPOINT_TIME >= RECOVERY WINDOW</code> . All backups older than the most recent backup that satisfies this condition are obsolete.

In addition to obsolete datafile backups, RMAN reports obsolete archived logs and archived log backups. When backups of datafiles from a point in time are obsolete, then archived logs that can only be applied to those datafiles are also no longer needed and become obsolete according to the specified retention policy.

Note: Note that if a datafile has never been backed up, then all archived redo logs back to the creation time of the file are required, and none of them are considered obsolete. With a full set of logs, the file can be completely re-created during media recovery. An empty datafile is automatically created during recovery, and the same changes applied to the original datafile after it was created are re-applied to the newly created file.

By default, the `REPORT OBSOLETE` command reports which files are obsolete under the currently configured retention policy. To generate reports of which files are obsolete according to different retention policies by using `REDUNDANCY` or `RECOVERY WINDOW` retention policy options with the `REPORT OBSOLETE` command. For example, if you run any of these commands:

```
REPORT OBSOLETE REDUNDANCY 2;
REPORT OBSOLETE RECOVERY WINDOW OF 5 DAYS;
```

RMAN displays backups that are obsolete according to those retention policies, regardless of the actual configured retention policy.

If you disable the retention policy completely (that is, if you run `CONFIGURE RETENTION POLICY TO NONE`), then RMAN does not consider any backups to be obsolete. If you run `REPORT OBSOLETE` with no options and no retention policy is configured, then RMAN issues an error message.

You can also query `V$BACKUP_FILES` and `RC_BACKUP_FILES`, using the `OBSOLETE` column to identify backup sets, datafile copies, and archived logs that are obsolete according to the configured retention policy.

If you are managing backup storage yourself instead of using a flash recovery area, then you should run `REPORT OBSOLETE` regularly to identify backups no longer needed to meet your retention policy, and delete these backups with `DELETE OBSOLETE`. If you use a flash recovery area, then backups stored there that are obsolete according to the configured retention policy are deleted automatically as space is needed. You do not have to perform `DELETE OBSOLETE` to reclaim space used for such backups.

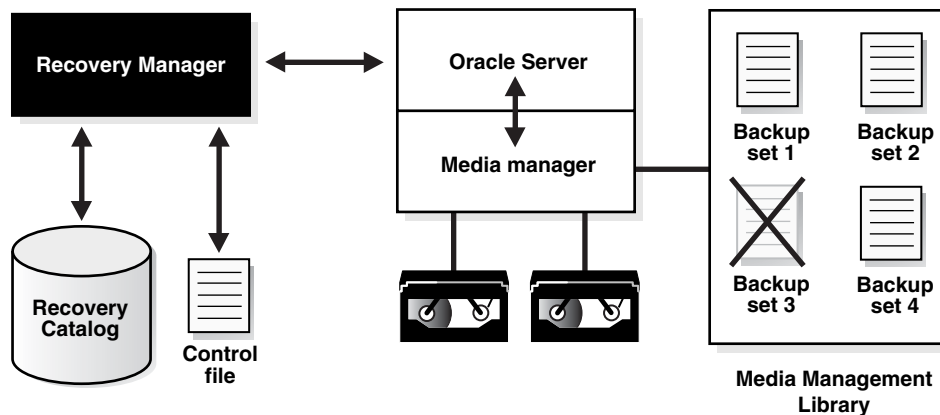
Crosschecks of RMAN Backups

RMAN's record of backups can become inconsistent with the actual backups that exist on tape or disk. For example, a user may delete backup pieces from disk using operating system commands, or a tape containing backups may be lost or may become corrupted.

To ensure that data about backups in the recovery catalog or control file is synchronized with actual files on disk or in the media management catalog, perform a **crosscheck**. The `CROSSCHECK` command operates only on files that are recorded in the RMAN repository.

Figure 9–1 illustrates a crosscheck of the media manager. RMAN queries the RMAN repository for the names and locations of the four backup pieces to be checked. RMAN sends this information to the target database server, which queries the media management software about the backups. The media management software then checks its media catalog and reports back to the server that backup set 3 is missing. RMAN updates the status of backup set 3 to `EXPIRED` in the repository. The record for backup set 3 will now be deleted if you run `DELETE EXPIRED`.

Figure 9–1 Crosschecking the Media Manager



Crosschecks are useful because they can

- Update outdated information about backups that disappeared from disk or tape or became corrupted
- Update the repository if you delete archived redo logs or other files with operating system commands

Use the crosscheck feature to check the status of a backup on disk or tape. If the backup is on disk, then `CROSSCHECK` checks whether the header of the file is valid. If a backup is on tape, then the command checks that the backups exist in the media management software's catalog.

Backup pieces and image copies can have the status `AVAILABLE`, `EXPIRED`, or `UNAVAILABLE`. You can view the status information in the output of the `LIST` command and the recovery catalog views.

You can issue the `DELETE EXPIRED` command to delete all expired backups. RMAN removes the record for the expired file from the repository. If for some reason the file still exists on the media, then RMAN issues warnings and lists the mismatched objects that cannot be deleted.

Note: The CROSSCHECK command does *not* delete operating system files or remove repository records. You must use the DELETE command for these operations.

See Also:

- *Oracle Database Backup and Recovery Basics* to learn how to perform crosschecks
- *Oracle Database Backup and Recovery Reference* for CROSSCHECK syntax and a description of the possible status values
- *Oracle Database Backup and Recovery Reference* for DELETE syntax

Deleting RMAN Backups

Several RMAN commands can cause the deletion of backups or archived logs, either directly or as a consequence of a backup of those files.

This section contains the following topics:

- [Summary of RMAN Methods for Deleting Backups](#)
- [Removal of Backups with the DELETE Command](#)
- [Removal of Backups with the DELETE Command](#)
- [How RMAN Deletes Backup Records from the RMAN Repository](#)
- [Behavior of DELETE Command When the Repository and Media Do Not Correspond](#)

Summary of RMAN Methods for Deleting Backups

Table 9–1 describes the functionality of the various RMAN commands that can cause backups to be deleted.

Table 9–1 Maintenance Commands and Scripts

Command or Script	Purpose
DELETE	<p>To delete backups, update the control file records to status DELETED, and remove their records from the recovery catalog (if a recovery catalog is used).</p> <p>You can specify that DELETE should remove backups that are EXPIRED or OBSOLETE. If you run DELETE EXPIRED on a backup that exists, RMAN issues a warning and does not delete the backup. You can override this behavior and delete the backup by using DELETE... FORCE .</p>
BACKUP . . . DELETE [ALL] INPUT	<p>To back up archived logs, datafile copies, or backup sets, then delete the input files from the operating system after the successful completion of the backup. RMAN also deletes and updates repository records for the deleted input files.</p> <p>If you specify DELETE INPUT (without ALL), then RMAN deletes only the specific files that it backs up. If you specify ALL INPUT, then RMAN deletes all copies of the files recorded in the RMAN repository.</p>

Table 9–1 (Cont.) Maintenance Commands and Scripts

Command or Script	Purpose
CHANGE . . . UNCATALOG	To delete recovery catalog records for specified backups and change their control file records to status DELETED. Note that the CHANGE . . . UNCATALOG command only changes the RMAN repository record of backups, and does not actually delete backups.

See Also: ["Crosschecks of RMAN Backups"](#) on page 9-4

Removal of Backups with the DELETE Command

The DELETE command can remove any file that the LIST and CROSSCHECK commands can operate on. For example, you can delete backup sets, archived redo logs, and datafile copies. The DELETE command removes both the physical file and the catalog record for the file.

Advantage of Using DELETE Instead of Operating System Commands

Always use DELETE command within RMAN to remove RMAN backups, rather than an operating system or media manager utility or command. Otherwise, the RMAN repository can contain records of backups that are no longer available for use in restore operations.

If you delete backups without using RMAN, you can use one of the following methods within RMAN to update the RMAN repository directly without performing a crosscheck:

- Use CROSSCHECK to change the status of these files to EXPIRED and then run DELETE EXPIRED to delete the records from the RMAN repository
- Use CHANGE . . . UNCATALOG to remove the catalog records

Deletion of Obsolete Backups

The DELETE OBSOLETE command provides a convenient way to delete backups that are no longer needed. It uses the same REDUNDANCY and RECOVERY WINDOW options as the REPORT OBSOLETE command.

If you have configured a retention policy, then you can run DELETE OBSOLETE periodically to delete all backups considered obsolete by this policy. For example, you can run DELETE OBSOLETE in a script every night with a scheduling utility, freeing disk and tape space used by backups that are no longer needed.

Note that using a flash recovery area as the destination for all backups eliminates the need to manage obsolete backups. Obsolete backups will be deleted from the flash recovery area automatically as disk space is needed to store backup-related files.

See Also: ["Reports of Obsolete Backups"](#) on page 9-2

Deletion of Expired Backups

The CROSSCHECK command updates the repository status for a backup to EXPIRED when it cannot locate it at the location to which it was backed up. This condition could occur if, for example, a backup was deleted from disk at the operating system level. You can identify expired backups by running the CROSSCHECK command as in the following example:

```
RMAN> CROSSCHECK BACKUP;
```

```

crosschecked backup piece: found to be 'AVAILABLE'
backup piece handle=0ad8d32i_1_1 recid=10 stamp=445025363
crosschecked backup piece: found to be 'AVAILABLE'
backup piece handle=c-1334876723-20011105-00 recid=11 stamp=445025367
crosschecked backup piece: found to be 'EXPIRED'
backup piece handle=0cd8d36i_1_1 recid=12 stamp=445025473
crosschecked backup piece: found to be 'AVAILABLE'
backup piece handle=c-1334876723-20011105-01 recid=13 stamp=445025475
Crosschecked 4 objects

```

If you run `CROSSCHECK` while some backup device is temporarily not accessible, this can mark backups as expired when in fact they are still present on the backup device, and are available again once the device is accessible. For example, this can happen if a disk is unmounted or if RMAN does not correctly connect to a media manager. In such a case, fix the problem that prevented RMAN from finding the backups and rerun `CROSSCHECK`.

The `DELETE EXPIRED` command removes the record of expired backups from the RMAN repository, by actually deleting the recovery catalog records for expired backups, and updates their control file records to status `DELETED`.

This command is especially useful if a user inadvertently deletes RMAN backups or archived logs from disk with an operating system utility. In such a case, the RMAN repository is not synchronized with the actual contents of disk. By running the `CROSSCHECK` command, RMAN marks the backups that it cannot find as expired. Then, you can run `DELETE EXPIRED` to remove the records for these files.

Deletion of Archived Redo Logs That Are Already Backed Up

You may want to delete files such as archived logs only if they have been backed up a specified number of times to tape. The `DELETE` command supports this behavior. The following example deletes all archived redo logs that have already been backed up at least two times to tape:

```
RMAN> DELETE ARCHIVELOG ALL BACKED UP 2 TIMES TO DEVICE TYPE sbt;
```

Deletion of Files After Backup Using `BACKUP ... DELETE INPUT`

The `BACKUP ... DELETE INPUT` command can delete archived redo logs, datafile copies, and backup sets after backing them up. This functionality is especially useful when backing up archived logs on disk to tape. RMAN backs up one copy of each log sequence number, and then deletes the file that it backs up. For example, assume that you issue:

```
RMAN> BACKUP ARCHIVELOG ALL DELETE INPUT;
```

In this command, RMAN backs up one copy of each log for each available sequence number, and then deletes only the archived redo log file that it actually backs up. If you have multiple redo log archiving destinations, the other copies of the same log sequence number are not deleted.

If you specify the `DELETE ALL INPUT` option, then RMAN deletes whichever files match the criteria that you specify, even if there are several files of the same log sequence number in different archiving destinations. For example, assume that you archive to three different directories. Then, you issue this command:

```
RMAN> BACKUP ARCHIVELOG ALL FROM SEQUENCE 1200 DELETE ALL INPUT;
```

In this case, RMAN backs up only one copy of each log sequence between 1200 and the most recent sequence, but deletes all logs with these sequence numbers from all three archive destinations.

Note: If you want to delete only a subset of the logs being backed up, you can use `BACKUP` without the `DELETE INPUT` clause, and then use a separate `DELETE` command with a `BACKED UP n TIMES TO DEVICE TYPE` clause.

In particular, in a Data Guard environment, it is not recommended to use `BACKUP... DELETE INPUT` to delete archived logs, as RMAN cannot delete the logs generated on another node. A separate `DELETE` command with a `BACKED UP n TIMES TO DEVICE TYPE SBT` command provides more direct control over deleting logs in this environment.

See Also: *Oracle Database Backup and Recovery Reference* for `DELETE` syntax

BACKUP... DELETE INPUT With Corrupt Archive Logs

During backup of archived redo logs, RMAN checks the file being backed up for corruption. If corruption is found, RMAN automatically switches to reading another copy of the same archived redo log, if one exists. For example, assume that `/log1` and `/log2` are the only enabled archiving destinations, and that they contain logs with sequence number up through 150. You run this command:

```
RMAN> BACKUP ARCHIVELOG FROM SEQUENCE 123 DELETE ALL INPUT;
```

RMAN can start reading from any copy of a given log, and switch to reading a different copy if it is unable to read the first copy in its entirety. For example, if RMAN starts reading the copy of log sequence 123 from `/log1` and discovers corruption in the file, it continues reading from the copy in `/log2`. Because `DELETE ALL INPUT` is specified, RMAN deletes all copies of logs on disk of sequence 123 and higher.

How RMAN Deletes Backup Records from the RMAN Repository

Every RMAN backup produces a corresponding record in the RMAN repository. This backup record is stored in the control file. In addition, if a recovery catalog is used, the record is also stored in the recovery catalog once it has been resynced with the control file.

For example, if you generate a full database backup set, then you can view the record for this backup set in the `V$BACKUP_SET` control file view. If you use a recovery catalog, then you can also access the record in the `RC_BACKUP_SET` catalog view.

The `V$` control file views and recovery catalog tables differ in the way that they store information, and this affects how RMAN handles repository records. The recovery catalog RMAN repository is stored in actual database tables, while the control file version of the repository is stored in an internal structure in the control file.

When you use an RMAN command to delete a backup, RMAN performs the following steps:

- Removes the physical file from the operating system (if the file is still present)
- Updates the backup records in the control file to status `DELETED`

- Removes the backup records from the recovery catalog tables (if RMAN is connected to a recovery catalog)

Because of the way that control file data is stored, RMAN cannot remove the record from the control file, only update it to DELETED status. However, because the catalog tables are ordinary database tables, RMAN removes rows from them in the same way that rows are deleted from any table.

Behavior of DELETE Command When the Repository and Media Do Not Correspond

The RMAN repository record for an object can sometimes fail to reflect the physical status of the object. For example, you backup an archived redo log to disk and then use an operating system utility to delete the object. If you do not run the CROSSCHECK command to update the repository, and if you then run DELETE against the object, then the repository shows that the object is AVAILABLE while the object is in fact missing. The following table indicates the behavior of DELETE in such situations.

Repository Status	Physical Status	Behavior of DELETE Command
AVAILABLE	Not found on media	Does not delete the object and reports the list of mismatched objects at the end of the job. RMAN does not update the repository status.
EXPIRED	Found on media	Does not delete the object and reports the list of mismatched objects at the end of the job. RMAN does not update the repository status.
UNAVAILABLE	Any	Removes repository record and deletes object if it exists. All I/O errors are ignored.

If you use the FORCE option of DELETE, RMAN will remove the repository record and delete the file if it exists. All I/O errors are ignored, and RMAN displays the number of objects deleted at the end of the job.

CHANGE... AVAILABLE and UNAVAILABLE with RMAN Backups

RMAN can update the repository to show backups as AVAILABLE or UNAVAILABLE. An unavailable backup is one that cannot be accessed at a particular moment but that has not been deleted. For example, you may have backups on tapes that are temporarily stored offsite and are inaccessible. You can use the CHANGE . . . UNAVAILABLE command to update the repository status for these backups to UNAVAILABLE so that RMAN will not select them for use in recovery operations.

When the tapes become available again, you can issue the CHANGE . . . AVAILABLE command to update the RMAN repository to show that these backups now can be used. After setting the files back to status AVAILABLE, you can also run a CROSSCHECK to verify that RMAN can access the files.

See Also:

- *Oracle Database Backup and Recovery Basics* to learn how to make backups available or unavailable
- *Oracle Database Backup and Recovery Reference* for CHANGE syntax

Changing Retention Policy Status of RMAN Backups

Use `CHANGE . . . KEEP` or `CHANGE . . . NOKEEP` to specify whether a backup should be subject to the configured retention policy or kept until a different date or even indefinitely.

The `KEEP` option exempts a backup from the current retention policy either indefinitely or until the specified `UNTIL` time. RMAN does not mark the files as obsolete even if they would be considered obsolete under the retention policy. Such backups are called **long-term backups**. `CHANGE . . . NOKEEP` is used to undo the effects of `CHANGE . . . KEEP`, so that the configured retention policy applies to the backup.

For example, the following command blocks RMAN from considering backupsets with the tag `'year_end_2002'` as obsolete under the retention policy:

```
RMAN> CHANGE BACKUPSET TAG year_end_2002 KEEP FOREVER NOLOGS;
```

To allow backupsets with the tag `year_end_2002` to be marked as obsolete based on the retention policy, use this command:

```
RMAN> CHANGE BACKUPSET TAG year_end_2002 NOKEEP;
```

If you want to prevent the use of a backup marked with `KEEP` in restore and recovery operations, then mark these backups as `UNAVAILABLE`. RMAN will not delete the records for these backups from the RMAN repository, but will not try to use them in restore and recovery until they are marked `AVAILABLE` again.

See Also:

- *Oracle Database Backup and Recovery Reference* for `CHANGE . . . KEEP` and `CHANGE . . . NOKEEP` syntax

Monitoring RMAN Through V\$ Views

When `LIST`, `REPORT` and `SHOW` do not provide all the information you need on RMAN activities, there are a number of useful `V$` views that can provide more details.

Sometimes it is useful to identify exactly what a server session performing a backup or recovery task is doing, or view the results of recent RMAN backup activities. The views described in the following table are useful in monitoring the progress of or obtaining information about RMAN jobs.

View	Description
<code>V\$PROCESS</code>	Identifies currently active processes.
<code>V\$SESSION</code>	Identifies currently active sessions. Use this view to determine which database server sessions correspond to which RMAN allocated channels.
<code>V\$SESSION_LONGOPS</code>	Provides progress reports on RMAN backup and restore jobs.
<code>V\$SESSION_WAIT</code>	Lists the events or resources for which sessions are waiting.
<code>V\$BACKUP_SYNC_IO</code>	Displays rows when the I/O is synchronous to the process (or thread on some platforms) performing the backup.

View	Description
V\$BACKUP_ASYNC_IO	Displays rows when the I/O is asynchronous to the process (or thread on some platforms) performing the backup. Note: Where asynchronous I/O is not supported by the host operating system, it may be implemented using slave I/O processes.

The following aspects of RMAN performance can be monitored through these views:

- [Correlating Server Sessions with RMAN Channels](#)
- [Monitoring RMAN Job Progress](#)
- [Monitoring RMAN Interaction with the Media Manager](#)
- [Monitoring RMAN Job Performance](#)

Correlating Server Sessions with RMAN Channels

To identify which server sessions correspond to which RMAN channels, you can query V\$SESSION and V\$PROCESS. The SPID column of V\$PROCESS identifies the operating system ID number for the process or thread. For example, on UNIX the SPID column shows the process ID, whereas on Windows the SPID column shows the thread ID. You have two basic methods for obtaining this information, depending on whether you have multiple RMAN sessions active concurrently.

Matching Server Sessions with Channels When One RMAN Session Is Active

When only one RMAN session is active, the easiest method for determining the server session ID for an RMAN channel is to execute the following query on the target database while the RMAN job is executing:

```
COLUMN CLIENT_INFO FORMAT a30
COLUMN SID FORMAT 999
COLUMN SPID FORMAT 9999

SELECT s.SID, p.SPID, s.CLIENT_INFO
FROM V$PROCESS p, V$SESSION s
WHERE p.ADDR = s.PADDR
AND CLIENT_INFO LIKE 'rman%'
;
```

If you do not run the SET COMMAND ID command in the RMAN job, then the CLIENT_INFO column displays in the following format:

```
rman channel=channel_id
```

For example, the following shows sample output:

```
SID SPID          CLIENT_INFO
-----
14 8374          rman channel=ORA_SBT_TAPE_1
```

Matching Server Sessions with Channels in Multiple RMAN Sessions

If more than one RMAN session is active, it is possible for the V\$SESSION.CLIENT_INFO column to yield the same information for a channel in each session. For example:

```
SID SPID          CLIENT_INFO
-----
14 8374          rman channel=ORA_SBT_TAPE_1
```

```
9 8642          rman channel=ORA_SBT_TAPE_1
```

In this case, you have the following methods for determining which channel corresponds to which SID value.

Obtaining the Channel ID from the RMAN Output In this method, you must first obtain the sid values from the RMAN output and then use these values in your SQL query.

To correlate a process with a channel during a backup:

1. In one of the active sessions, run the RMAN job as normal and examine the output to get the sid for the channel. For example, the output may show:

```
Starting backup at 21-AUG-01
allocated channel: ORA_SBT_TAPE_1
channel ORA_SBT_TAPE_1: sid=14 devtype=SBT_TAPE
```

2. Start a SQL*Plus session and then query the joined V\$SESSION and V\$PROCESS views *while the RMAN job is executing*. For example, enter:

```
COLUMN CLIENT_INFO FORMAT a30
COLUMN SID FORMAT 999
COLUMN SPID FORMAT 9999

SELECT s.SID, p.SPID, s.CLIENT_INFO
FROM V$PROCESS p, V$SESSION s
WHERE p.ADDR = s.PADDR
AND CLIENT_INFO LIKE 'rman%'
/
```

Use the sid value obtained from the first step to determine which channel corresponds to which server session:

SID	SPID	CLIENT_INFO
14	2036	rman channel=ORA_SBT_TAPE_1
12	2066	rman channel=ORA_SBT_TAPE_1

Correlating Server Sessions with Channels by Using SET COMMAND ID In this method, you specify a command ID string in the RMAN backup script. You can then query V\$SESSION.CLIENT_INFO for this string.

To correlate a process with a channel during a backup:

1. In each session, set the COMMAND ID to a different value *after* allocating the channels and then back up the desired object. For example, enter the following in session 1:

```
RUN
{
  ALLOCATE CHANNEL c1 TYPE disk;
  SET COMMAND ID TO 'sess1';
  BACKUP DATABASE;
}
```

Set the command ID to a string such as sess2 in the job running in session 2:

```
RUN
{
  ALLOCATE CHANNEL c1 TYPE sbt;
  SET COMMAND ID TO 'sess2';
  BACKUP DATABASE;
```

}

2. Start a SQL*Plus session and then query the joined V\$SESSION and V\$PROCESS views *while the RMAN job is executing*. For example, enter:

```
SELECT SID, SPID, CLIENT_INFO
   FROM V$PROCESS p, V$SESSION s
  WHERE p.ADDR = s.PADDR
        AND CLIENT_INFO LIKE '%id=sess%';
```

If you run the SET COMMAND ID command in the RMAN job, then the CLIENT_INFO column displays in the following format:

```
id=command_id,rman channel=channel_id
```

For example, the following shows sample output:

```

SID SPID          CLIENT_INFO
-----
11 8358          id=sess1
15 8638          id=sess2
14 8374          id=sess1,rman channel=c1
9  8642          id=sess2,rman channel=c1
```

The rows that contain the string `rman channel` show the channel performing the backup. The remaining rows are for the connections to the target database.

See Also: *Oracle Database Backup and Recovery Reference* for SET COMMAND ID syntax, and *Oracle Database Reference* for more information on V\$SESSION and V\$PROCESS

Monitoring RMAN Job Progress

Monitor the progress of backups and restores by querying the view V\$SESSION_LONGOPS. RMAN uses two types of rows in V\$SESSION_LONGOPS: detail and aggregate rows. Detail rows describe the files being processed by one job step, while aggregate rows describe the files processed by all job steps in an RMAN command. A job step is the creation or restore of one backup set or datafile copy. Detail rows are updated with every buffer that is read or written during the backup step, so their granularity of update is small. Aggregate rows are updated when each job step completes, so their granularity of update is large.

Table 9–2 describes column in V\$SESSION_LONGOPS that are most relevant for RMAN. Typically, you will view the detail rows rather than the aggregate rows to determine the progress of each backup set.

Table 9–2 Columns of V\$SESSION_LONGOPS Relevant for RMAN

Column	Description for Detail Rows
SID	The server session ID corresponding to an RMAN channel.
SERIAL#	The server session serial number. This value changes each time a server session is reused.
OPNAME	A text description of the row. Examples of details rows include RMAN: datafile copy, RMAN: full datafile backup, and RMAN: full datafile restore. Note: RMAN: aggregate input and RMAN: aggregate output are the only aggregate rows.

Table 9–2 (Cont.) Columns of V\$SESSION_LONGOPS Relevant for RMAN

Column	Description for Detail Rows
CONTEXT	For backup output rows, this value is 2. For all other rows except proxy copy (which does not update this column), the value is 1.
SOFAR	The meaning of this column depends on the type of operation described by this row: <ul style="list-style-type: none"> ■ For image copies, the number of blocks that have been read. ■ For backup input rows, the number of blocks that have been read from the files being backed up. ■ For backup output rows, the number of blocks that have been written to the backup piece. ■ For restores, the number of blocks that have been processed to the files that are being restored in this one job step. ■ For proxy copies, the number of files that have been copied.
TOTALWORK	The meaning of this column depends on the type of operation described by this row: <ul style="list-style-type: none"> ■ For image copies, the total number of blocks in the file. ■ For backup input rows, the total number of blocks to be read from all files processed in this job step. ■ For backup output rows, the value is 0 because RMAN does not know how many blocks that it will write into any backup piece. ■ For restores, the total number of blocks in all files restored in this job step. ■ For proxy copies, the total number of files to be copied in this job step.

Each server session performing a backup or restore reports its progress compared to the total amount of work required for a job step. For example, if you perform a database restore that uses two channels, and each channel has two backup sets to restore (a total of four sets), then each server session reports its progress through a single backup set. When that set is completely restored, RMAN begins reporting progress on the next set to restore.

To monitor job progress:

1. Before starting the job, create a script file (called, for this example, `longops`) containing the following SQL statement:

```
SELECT SID, SERIAL#, CONTEXT, SOFAR, TOTALWORK,
       ROUND(SOFAR/TOTALWORK*100,2) "%_COMPLETE"
FROM V$SESSION_LONGOPS
WHERE OPNAME LIKE 'RMAN%'
      AND OPNAME NOT LIKE '%aggregate%'
      AND TOTALWORK != 0
      AND SOFAR <> TOTALWORK
;
```

2. After connecting to the target database and, if desired, the recovery catalog database, start an RMAN job. For example, enter:

```
RESTORE DATABASE;
```

3. While the job is running, start SQL*Plus connected to the target database, and execute the `longops` script to check the progress of the RMAN job. If you repeat the query while the restore progresses, then you see output such as the following:

```
SQL> @longops
```

SID	SERIAL#	CONTEXT	SOFAR	TOTALWORK	%_COMPLETE
8	19	1	10377	36617	28.34

```
SQL> @longops
```

SID	SERIAL#	CONTEXT	SOFAR	TOTALWORK	% COMPLETE
8	19	1	21513	36617	58.75

```
SQL> @longops
```

SID	SERIAL#	CONTEXT	SOFAR	TOTALWORK	% COMPLETE
8	19	1	29641	36617	80.95

```
SQL> @longops
```

SID	SERIAL#	CONTEXT	SOFAR	TOTALWORK	% COMPLETE
8	19	1	35849	36617	97.9

```
SQL> @longops
```

```
no rows selected
```

- If you run the script at intervals of two minutes or more and the `%_COMPLETE` column does not increase, then RMAN is encountering a problem. Refer to ["Monitoring RMAN Interaction with the Media Manager"](#) on page 9-15 to obtain more information.

If you frequently monitor the execution of long-running tasks, you could create a shell script or batch file under your host operating system that runs SQL*Plus to execute this query repeatedly.

Monitoring RMAN Interaction with the Media Manager

You can use the event names in the dynamic performance event views to monitor RMAN calls to the media management API. The event names have one-to-one correspondence with `sbt` functions, as shown in the following examples:

```
Backup: sbtinit
Backup: ssbtopen
Backup: ssbtread
Backup: ssbtwrite
Backup: ssbtbackup
```

```
.
.
.
```

To obtain the complete list of `sbt` events, you can use the following query:

```
select name from v$event_name where name like '%sbt%';
```

Before making a call to any of functions in the media management API, the server adds a row in `V$SESSION_WAIT`, with the `STATE` column including the string `WAITING`. The `V$SESSION_WAIT.SECONDS_IN_WAIT` column shows the number of seconds that the server has been waiting for this call to return. After an `sbt` function is returned from the media manager, this row disappears.

A row in `V$SESSION_WAIT` corresponding to an `sbt` event name does not indicate a problem, because the server updates these rows at runtime. The rows appear and disappear as calls are made and returned. However, if the `SECONDS_IN_WAIT` column is high, then the media manager may be hung.

To monitor the `sbt` events, you can run the following SQL query:

```

COLUMN EVENT FORMAT a10
COLUMN SECONDS_IN_WAIT FORMAT 999
COLUMN STATE FORMAT a20
COLUMN CLIENT_INFO FORMAT a30

SELECT p.SPID, EVENT, SECONDS_IN_WAIT AS SEC_WAIT,
       sw.STATE, CLIENT_INFO
FROM V$SESSION_WAIT sw, V$SESSION s, V$PROCESS p
WHERE sw.EVENT LIKE 's%bt%'
      AND s.SID=sw.SID
      AND s.PADDR=p.ADDR
;

```

Examine the SQL output to determine which `sbt` functions are waiting. For example, the following output indicates that RMAN has been waiting for the `sbtbackup` function to return for ten minutes:

SPID	EVENT	SEC_WAIT	STATE	CLIENT_INFO
8642	Backup: sbtbackup	600	WAITING	rman channel=ORA_SBT_TAPE_1

Note: The `V$SESSION_WAIT` view shows only database events, not media manager events.

See Also: *Oracle Database Reference* for descriptions of `V$SESSION_WAIT`

Monitoring RMAN Job Performance

Monitor backup and restore performance by querying `V$BACKUP_SYNC_IO` and `V$BACKUP_ASYNC_IO`.

See Also: *Oracle Database Reference* for more information on these `V$` views, and "[Step 5: Query V\\$ Views to Identify Bottlenecks](#)" on page 11-8 to learn how to use these views to tune backup performance

Managing the Recovery Catalog

This chapter describes how to manage an RMAN recovery catalog, which holds RMAN repository data for one or more databases in a separate database schema, in addition to using the control files of the databases.

This chapter contains these topics:

- [Creating a Recovery Catalog](#)
- [Managing Target Database Records in the Recovery Catalog](#)
- [Resynchronizing the Recovery Catalog](#)
- [Managing the Control File When You Use a Recovery Catalog](#)
- [Backing Up and Recovering the Recovery Catalog](#)
- [Exporting and Importing the Recovery Catalog](#)
- [Increasing Availability of the Recovery Catalog](#)
- [Querying Recovery Catalog Views](#)
- [Determining the Schema Version of the Recovery Catalog](#)
- [Upgrading the Recovery Catalog](#)
- [Dropping the Recovery Catalog](#)

See Also: *Oracle Database Backup and Recovery Basics* to learn how to manage the RMAN repository as stored in the control file, without a recovery catalog

Creating a Recovery Catalog

Creating a recovery catalog is a three-step process: you must configure the database that will contain the recovery catalog, create the recovery catalog owner, and then create the recovery catalog itself.

Configuring the Recovery Catalog Database

When you use a recovery catalog, RMAN requires that you maintain a recovery catalog schema. The recovery catalog is stored in the default tablespace of the schema. Note that `SYS` cannot be the owner of the recovery catalog.

Decide which database you will use to install the recovery catalog schema, and also how you will back up this database.

Note: Do not use the target database to be backed up as the database for the recovery catalog. The recovery catalog must be protected in the event of the loss of the target database.

Also, decide whether to operate the catalog database in ARCHIVELOG mode, which is recommended.

Planning the Size of the Recovery Catalog Schema

You must allocate space to be used by the catalog schema. The size of the recovery catalog schema depends upon the number of databases monitored by the catalog. The schema also grows as the number of archived redo log files and backups for each database increases. Finally, if you use RMAN stored scripts stored in the catalog, some space must be allocated for those scripts.

For an example, assume that the `trgt` database has 100 files, and you back up the database once a day, producing 50 backup sets containing 1 backup piece each. If you assume that each row in the backup piece table uses the maximum amount of space, then one daily backup will consume less than 170 KB in the recovery catalog. So, if you back up once a day for a year, then the total storage in this period is about 62 MB. Assume approximately the same amount for archived logs. Hence, the worst case is about 120 MB for a year for metadata storage.

For a more typical case in which only a portion of the backup piece row space is used, 15 MB for each year is a realistic estimate.

If you plan to register multiple databases in your recovery catalog, remember to add up the space required for each one based on the previous calculation to arrive at a total size for the default tablespace of the recovery catalog schema.

Allocating Disk Space for the Recovery Catalog Database

If you are creating your recovery catalog in an already-existing database, add enough room to hold the default tablespace to the recovery catalog schema. If you are creating a new database to hold your recovery catalog, then, in addition to the space for the recovery catalog schema itself, you must allow space for other files in the recovery catalog database:

- SYSTEM tablespace
- Temporary tablespaces
- Rollback segment tablespaces
- Online redo log files

Most of the space used in the recovery catalog database is devoted to supporting tablespaces, for example, the `SYSTEM`, temporary, and rollback or undo tablespaces. [Table 10–1, "Typical Recovery Catalog Space Requirements for 1 Year"](#) describes typical space requirements.

Table 10–1 Typical Recovery Catalog Space Requirements for 1 Year

Type of Space	Space Requirement
SYSTEM tablespace	90 MB
Temp tablespace	5 MB
Rollback or undo tablespace	5 MB

Table 10–1 (Cont.) Typical Recovery Catalog Space Requirements for 1 Year

Type of Space	Space Requirement
Recovery catalog tablespace	15 MB for each database registered in the recovery catalog
Online redo logs	1 MB each (3 groups, each with 2 members)

Caution: Ensure that the recovery catalog and target databases do *not* reside on the same disk. If both your recovery catalog and your target database suffer hard disk failure, your recovery process is much more difficult. If possible, take other measures as well to eliminate common points of failure between your recovery catalog database and the databases you are backing up.

Creating the Recovery Catalog Owner

After choosing the recovery catalog database and creating necessary space, you are ready to create the owner of the recovery catalog and grant this user necessary privileges.

Assume the following background information for the instructions in the following sections:

- User `SYS` with password `oracle` has `SYSDBA` privileges on the recovery catalog database `catdb`.
- A tablespace called `tools` in the recovery catalog database `catdb` stores the recovery catalog. Note that to use an RMAN reserved word as a tablespace name, you must enclose it in quotes and put it in uppercase. (Refer to *Oracle Database Backup and Recovery Reference* for a list of RMAN reserved words.)
- A tablespace called `temp` exists in the recovery catalog database.
- The database is configured in the same way as all normal databases, for example, `catalog.sql` and `catproc.sql` have successfully run.

To create the recovery catalog schema in the recovery catalog database:

1. Start SQL*Plus and then connect with administrator privileges to the database containing the recovery catalog. For example, enter:

```
CONNECT SYS/oracle@catdb AS SYSDBA
```

2. Create a user and schema for the recovery catalog. For example, enter:

```
CREATE USER rman IDENTIFIED BY cat
TEMPORARY TABLESPACE temp
DEFAULT TABLESPACE tools
QUOTA UNLIMITED ON tools;
```

3. Grant the `RECOVERY_CATALOG_OWNER` role to the schema owner. This role provides the user with all privileges required to maintain and query the recovery catalog.

```
SQL> GRANT RECOVERY_CATALOG_OWNER TO rman;
```

Creating the Recovery Catalog

After creating the catalog owner, create the catalog tables with the RMAN CREATE CATALOG command. The command creates the catalog in the default tablespace of the catalog owner.

To create the recovery catalog:

1. Connect to the database that will contain the catalog as the catalog owner. For example:

```
% rman
RMAN> CONNECT CATALOG rman/cat@catdb
```

2. Run the CREATE CATALOG command to create the catalog. The creation of the catalog can take several minutes. If the catalog tablespace is this user's default tablespace, then you can run this command:

```
CREATE CATALOG;
```

You can specify the tablespace name for the catalog in the CREATE CATALOG command. For example:

```
CREATE CATALOG TABLESPACE cat_ts;
```

Note: If the tablespace name you wish to use for the recovery catalog is an RMAN reserved word, then it *must* be uppercase and enclosed in quotes. For example:

```
CREATE CATALOG TABLESPACE 'CATALOG';
```

You can check the results by using SQL*Plus to query the recovery catalog to see which tables were created:

```
SQL> SELECT TABLE_NAME FROM USER_TABLES;
```

See Also: *Oracle Database SQL Reference* for the SQL syntax for the GRANT and CREATE USER statements, and *Oracle Database Backup and Recovery Reference* for CREATE CATALOG command syntax

Managing Target Database Records in the Recovery Catalog

This section describes how to maintain target database records in the recovery catalog. It contains the following sections:

- [Registering a Database in the Recovery Catalog](#)
- [Unregistering a Target Database from the Recovery Catalog](#)
- [Resetting the Database Incarnation in the Recovery Catalog](#)
- [Removing DELETED Records From the Recovery Catalog After Upgrade](#)

Registering a Database in the Recovery Catalog

The first step in using a recovery catalog with a target database is registering the database in the recovery catalog. Use the following procedure:

1. After making sure the recovery catalog database is open, connect RMAN to the target database and recovery catalog database. For example, issue the following to connect to the catalog database `catdb` as user `rman` (who owns the catalog schema):

```
% rman TARGET / CATALOG rman/cat@catdb
```

2. If the target database is not mounted, then mount or open it:

```
RMAN> STARTUP MOUNT;
```

3. Register the target database in the connected recovery catalog:

```
RMAN> REGISTER DATABASE;
```

RMAN creates rows in the catalog tables to contain information about the target database, then copies all pertinent data about the target database from the control file into the catalog, synchronizing the catalog with the control file.

Verify that the registration was successful by running `REPORT SCHEMA`:

```
RMAN> REPORT SCHEMA;
```

```
Report of database schema
File Size(MB)  Tablespace      RB segs  Datafile Name
-----
1          307200 SYSTEM          NO   /oracle/oradata/trgt/system01.dbf
2           20480 UNDOTBS         YES  /oracle/oradata/trgt/undotbs01.dbf
3           10240 CWMLITE         NO   /oracle/oradata/trgt/cwmlite01.dbf
4           10240 DRSYS           NO   /oracle/oradata/trgt/drsys01.dbf
5           10240 EXAMPLE         NO   /oracle/oradata/trgt/example01.dbf
6           10240 INDX            NO   /oracle/oradata/trgt/indx01.dbf
7           10240 TOOLS          NO   /oracle/oradata/trgt/tools01.dbf
8           10240 USERS          NO   /oracle/oradata/trgt/users01.dbf
```

Cataloging Older Files in the Recovery Catalog

If you have datafile copies, backup pieces or archive logs on disk, you can catalog them in the recovery catalog using the `CATALOG` command. When using a recovery catalog, cataloging older backups that have aged out of the control file lets RMAN use the older backups during restore operations. For example:

```
RMAN> CATALOG DATAFILECOPY '/disk1/old_datafiles/01_01_2003/users01.dbf';
RMAN> CATALOG ARCHIVELOG '/disk1/arch_logs/archive1_731.dbf',
'/disk1/arch_logs/archive1_732.dbf';
RMAN> CATALOG BACKUPPIECE '/disk1/backups/backup_820.bkp';
```

You can also catalog multiple backup files in a directory at once, using the `CATALOG START WITH` command, as shown in this example:

```
RMAN> CATALOG START WITH '/disk1/backups/';
```

RMAN lists the files to be added to the RMAN repository and prompts for confirmation before adding the backups.

Note: Be careful when creating your prefix for `CATALOG START WITH`. RMAN scans all paths for all files on disk which begin with the specified prefix. The prefix is not just a directory name. Using the wrong prefix can cause the cataloging of the wrong set of files. For example, a group of directories `/disk1/backups`, `/disk1/backups-year2003`, `/disk1/backupsets`, and `/disk1/backupsets/test` and so on, all contain backup files. The command

```
RMAN> CATALOG START WITH '/disk1/backups';
```

catalogs all files in all of these directories, because `/disk1/backups` is a prefix for the paths for all of these directories. In order to catalog only backups in the `/disk1/backups` directory, the correct command would be:

```
RMAN> CATALOG START WITH '/disk1/backups/';
```

Cataloging Oracle7 Datafile Copies in the Recovery Catalog In general, only Oracle8 and higher files can be cataloged. Datafile copies from Oracle7 databases can also be cataloged, if they do not require the application of Oracle7 redo before they can be opened. Datafile copies made in the following circumstances satisfy this requirement:

- Datafile copies made when the database was shut down consistently. The database must not have been opened again before migration to a higher version of Oracle.
- Datafile copies made after a tablespace became offline normal or read-only. The tablespaces must not have been brought online or made read/write again before migration to a higher version of Oracle.

See Also:

- *Oracle Database Backup and Recovery Reference* for `REGISTER` syntax
- *Oracle Database Upgrade Guide* for issues relating to database migration

Registering Multiple Databases in a Recovery Catalog

You can register multiple target databases in a single recovery catalog, if they do not have duplicate DBIDs. RMAN uses the DBID to distinguish one database from another.

In general, if you use the `DUPLICATE RMAN` command or `CREATE DATABASE SQL` statement, the database created is assigned a unique DBID. If you create a database by some other means, such as a user-managed copy, then the new database may have the same DBID as the one from which it was copied. You will not be able to register both databases in the same recovery catalog until you change the DBID of the copied database using the `DBNEWID` utility.

Note that you can also register a single target databases in multiple recovery catalogs.

See Also:

- *Oracle Database Backup and Recovery Reference* for DUPLICATE syntax
- *Oracle Database Utilities* to learn how to use the DBNEWID utility to change the DBID
- *Oracle Database Upgrade Guide* for issues relating to database migration

Unregistering a Target Database from the Recovery Catalog

The UNREGISTER DATABASE command is used to unregister a database from the recover catalog.

Note: When a database is unregistered from the recovery catalog, all RMAN repository records in the recovery catalog are lost. The database can be registered again, but the recovery catalog records for that database are then based on the contents of the control file at the time of re-registration. Records older than the CONTROLFILE_RECORD_KEEP_TIME setting in the target database control file are lost. Stored scripts, which are not stored in the control file, are also lost.

To unregister a database:

1. Start RMAN and connect to the target database. For example, enter:

```
% rman TARGET / CATALOG rman/cat@catdb
```

```
connected to target database: RDBMS (DBID=1237603294)
connected to recovery catalog database
```

Make a note of the DBID as displayed by RMAN at startup. If there is more than one database registered in the recovery catalog, you will need the DBID to uniquely identify the database.

It is not necessary to connect to the target database, but if you do not, then you must specify the name of the target database in the UNREGISTER command. If more than one database has the same name in the recovery catalog, then you must create a RUN block around the command and use SET DBID to set the DBID for the database.

2. As a precaution, it may be useful to list all of the backups recorded in the recovery catalog using LIST BACKUP SUMMARY and LIST COPY SUMMARY. This way, you can re-catalog backups not known to the control file if you later decide to re-register the database.
3. If your intention is to actually delete all backups of the database completely, rather than just removing the database from the recovery catalog and relying on the control file to store the RMAN repository for this database, then run DELETE statements to delete all existing backups. For example:

```
DELETE BACKUP DEVICE TYPE sbt;
DELETE BACKUP DEVICE TYPE DISK;
DELETE COPY;
```

RMAN will list the backups that it intends to delete and prompt for confirmation before deleting them.

4. Run the UNREGISTER DATABASE command. For example:

```
UNREGISTER DATABASE;
```

RMAN displays the database name and DBID, and prompts you for a confirmation:

```
database name is "RDBMS" and DBID is 931696259
```

```
Do you really want to unregister the database (enter YES or NO)? yes
```

When the process is complete, RMAN outputs the following message:

```
database unregistered from the recovery catalog
```

Resetting the Database Incarnation in the Recovery Catalog

When you run either the RMAN command or the SQL statement ALTER DATABASE OPEN RESETLOGS, you create a new incarnation of the database. You can access a record of the new incarnation in the V\$DATABASE_INCARNATION view of the target database.

If you run the RMAN command or the SQL statement ALTER DATABASE OPEN RESETLOGS, then a new database incarnation record is automatically created in the recovery catalog. The database also implicitly and automatically issues a RESET DATABASE command, which specifies that this new incarnation of the database is the current incarnation. All subsequent backups and log archiving done by the target database is associated with the new database incarnation.

You may need to change the current incarnation of the database in the recovery catalog for some recovery tasks. The steps for changing the incarnation are described in the following procedure.

To reset the recovery catalog to an older incarnation:

1. Determine the incarnation key of the desired database incarnation. Obtain the incarnation key value by issuing a LIST command:

```
LIST INCARNATION OF DATABASE trgt;
```

```
List of Database Incarnations
```

DB Key	Inc Key	DB Name	DB ID	STATUS	Reset SCN	Reset Time
1	2	TRGT	1224038686	PARENT	1	02-JUL-02
1	582	TRGT	1224038686	CURRENT	59727	10-JUL-02

The incarnation key is listed in the "Inc Key" column.

2. Reset the database to the old incarnation. For example, enter:

```
RESET DATABASE TO INCARNATION 2;
```

3. If the control file of the previous incarnation is available and mounted, then skip to step 6 of this procedure. Otherwise, shut down the database and start it without mounting. For example:

```
SHUTDOWN IMMEDIATE
STARTUP NOMOUNT
```


- Restore a control file from the old incarnation. If you have a control file tagged, then specify the tag. Otherwise, you can run the `SET UNTIL` command, as in this example:

```
RUN
{
  SET UNTIL 'SYSDATE-45';
  RESTORE CONTROLFILE; # only if current control file is not available
}
```

- Mount the restored control file:

```
ALTER DATABASE MOUNT;
```

- Run `RESTORE` and `RECOVER` commands to restore and recover the database files from the prior incarnation, then open the database with the `RESETLOGS` option. For example, enter:

```
RESTORE DATABASE;
RECOVER DATABASE;
ALTER DATABASE OPEN RESETLOGS;
```

See Also: *Oracle Database Backup and Recovery Reference* for `RESET DATABASE` syntax, *Oracle Database Backup and Recovery Reference* for `LIST` syntax

Removing DELETED Records From the Recovery Catalog After Upgrade

In Oracle9i and later, RMAN always removes catalog records on deleting files, and never updates them to status `DELETED`. However, in releases prior to Oracle9i, RMAN updated recovery catalog records to `DELETED` status after deleting the physical files rather than removing the records. Thus, records with status `DELETED` can appear in the recovery catalog when you upgrade a recovery catalog created prior to Oracle9i to the current release. For this special case, you can run the `prgmanc.sql` script, which is stored in an operating system specific location (`$ORACLE_HOME/rdbms/admin` on UNIX).

To remove all backup records with status DELETED:

- Start a SQL*Plus session and connect to the recovery catalog. This example connects to the database `rcat` as user `rman`:

```
% sqlplus rman/cat@catdb
```

- Run the script `prgmanc.sql` script:

```
SQL> @?/rdbms/admin/prgmanc.sql
```

The script removes all records with status `DELETED` from the recovery catalog.

Resynchronizing the Recovery Catalog

When RMAN performs a **resynchronization**, it compares the recovery catalog to either the current control file of the target database or a backup control file and updates the recovery catalog with information that is missing or changed. When resynchronizing, RMAN does the following:

- Creates a snapshot control file.
- Compares the recovery catalog to the snapshot control file.

- Updates the recovery catalog with missing or changed information.

RMAN performs resynchronizations automatically as needed when you execute certain commands, including `BACKUP`. You can also manually perform a full resynchronization using the `RESYNC CATALOG` command.

Types of Records Updated When Recovery Catalog is Resynchronized

Table 10–2, "Records Updated During a Resynchronization" describes the types of records that RMAN resynchronizes.

Table 10–2 Records Updated During a Resynchronization

Records	Description
Log history	Created when an online redo log switch occurs.
Archived redo logs	Associated with archived logs that were created by archiving an online log, copying an existing archived redo log, or restoring an archived redo log backup set. RMAN tracks this information so that it knows which archived logs it should expect to find.
Backup history	Associated with backup sets, backup pieces, proxy copies, and file copies. The <code>RESYNC CATALOG</code> command updates these records when a <code>BACKUP</code> command is executed.
Incarnation history	Associated with database incarnations.
Physical schema	Associated with datafiles and tablespaces. If the target database is open, then undo segment information is also updated. Physical schema information in the recovery catalog is updated only when the target has the current control file mounted. If the target database has mounted a backup control file, a freshly created control file, or a control file that is less current than a control file that was seen previously, then physical schema information in the recovery catalog is <i>not</i> updated. Physical schema information is also not updated when you use the <code>RESYNC CATALOG FROM CONTROLFILECOPY</code> command.

Full and Partial Resynchronization

Resynchronizations can be full or partial. In a partial resynchronization, RMAN reads the current control file to update changed information about new backups, new archived logs, and so forth. However, RMAN does not resynchronize metadata about the database physical schema: datafiles, tablespaces, redo threads, rollback segments (only if the database is open), and online redo logs. In a full resynchronization, RMAN updates all changed records, including those for the database schema.

Note: Although RMAN performs partial resynchronizations when using a backup control file, it does not perform full resynchronizations. A backup control file may not have correct information about the database physical schema, so a full resynchronization could update the recovery catalog with incorrect information.

Note: You can use Oracle Enterprise Manager to perform catalog resynchronizations.

See Also: *Oracle Database Backup and Recovery Reference* for more information about the `RESYNC` command

When to Resynchronize the Recovery Catalog

RMAN automatically performs full or partial resynchronizations in most situations in which they are needed. Most RMAN commands such as `BACKUP`, `DELETE`, and so forth perform a full or partial resynchronization (depending on whether the schema metadata has changed) automatically when the target database control file is mounted and the recovery catalog database is available. Thus, you should not need to manually run `RESYNC CATALOG` very often.

The following sections describe situations in which a manual resynchronization is required.

Resynchronizing After the Recovery Catalog is Unavailable

If the recovery catalog is unavailable when you issue RMAN commands that cause a partial resynchronization, then open the catalog database later and resynchronize it manually with the `RESYNC CATALOG` command.

For example, the target database may be in New York while the recovery catalog database is in Japan. You may not want to make daily backups of the target database in `CATALOG` mode, to avoid depending on the availability of a geographically distant database. In such a case you could connect to the catalog as often as feasible (for example, once each week) and run the `RESYNC CATALOG` command.

Resynchronizing in ARCHIVELOG Mode When You Back Up Infrequently

Assume that you do the following:

- Run the database in `ARCHIVELOG` mode
- Back up the database infrequently (for example, hundreds of archive logs are archived between database backups)
- Generate a high number of log switches every day (for example, 1000 switches between catalog resynchronizations)

In this case, you may want to manually resynchronize the recovery catalog regularly because the recovery catalog is *not* updated automatically when a redo log switch occurs or when a redo log is archived. The database stores information about log switches and archived redo logs only in the control file. You must periodically resynchronize in order to propagate this information into the recovery catalog.

How frequently you need to resynchronize the recovery catalog depends on the rate at which the database archives redo logs. The cost of the operation is proportional to the number of records in the control file that have been inserted or changed since the previous resynchronization. If no records have been inserted or changed, then the cost of resynchronization is very low; if many records have been inserted or changed, then the resynchronization is more time-consuming.

Resynchronizing After Physical Database Changes

Resynchronize the recovery catalog after making any change to the physical structure of the target database. As with redo log archive operations, the recovery catalog is not updated automatically after physical schema changes, such as adding or dropping tablespaces, adding datafiles to a tablespace, or adding or dropping rollback segments.

Forcing a Full Resynchronization of the Recovery Catalog

Use `RESYNC CATALOG` to force a full resynchronization of the recovery catalog.

1. Connect RMAN to the target and recovery catalog databases, and then mount or open the target database if it is not already mounted or open:

```
STARTUP MOUNT;
```

2. Run the `RESYNC CATALOG` command at the RMAN prompt:

```
RESYNC CATALOG;
```

See Also: *Oracle Database Backup and Recovery Reference* for `RESYNC CATALOG` command syntax

Resynchronizing the Recovery Catalog and `CONTROL_FILE_RECORD_KEEP_TIME`

If you maintain a recovery catalog, then use the RMAN `RESYNC CATALOG` command often enough to ensure that control file records are propagated to the recovery catalog before they are reused.

Make sure that `CONTROL_FILE_RECORD_KEEP_TIME` is longer than the interval between backups or resynchronizations. Otherwise, control file records could be reused before they are propagated to the recovery catalog. An extra week is a safe margin in most circumstances.

Caution: Never set `CONTROL_FILE_RECORD_KEEP_TIME` to 0. If you do, then backup records may be overwritten in the control file before RMAN is able to add them to the catalog.

See Also: *Oracle Database Backup and Recovery Basics* to learn how to monitor the overwriting of control file records

Managing the Control File When You Use a Recovery Catalog

Your goal is to ensure that the metadata in the recovery catalog is current. Because the recovery catalog obtains its metadata from the target control file, the currency of the data in the catalog depends on the currency of the data in the control file. You need to make sure that the backup metadata in the control file is recorded in the catalog before it is overwritten with new records.

The `CONTROL_FILE_RECORD_KEEP_TIME` initialization parameter determines the minimum number of days that records are retained in the control file before they are candidates for being overwritten. Hence, you must ensure that you resynchronize the recovery catalog with the control file records before these records are erased. As described in "[Resynchronizing the Recovery Catalog and `CONTROL_FILE_RECORD_KEEP_TIME`](#)" on page 10-12, you should perform either of the following actions at intervals less than the `CONTROL_FILE_RECORD_KEEP_TIME` setting:

- Make a backup, thereby performing an implicit resynchronization of the recovery catalog
- Manually resynchronize the recovery catalog with the `RESYNC CATALOG` command

So, to ensure the currency of the information in the recovery catalog, the frequency of resynchronizations should be related to the value for the `CONTROL_FILE_RECORD_KEEP_TIME` initialization parameter.

One problem can arise if the control file becomes too large. The size of the target database control file grows depending on the number of:

- Backups that you perform
- Archived redo logs that the database generates
- Days that this information is stored in the control file

As explained in *Oracle Database Backup and Recovery Basics*, if the control file grows so large that it can no longer expand because it has reached either the maximum number of blocks or the maximum number of records, then the database may overwrite the oldest records even if their age is less than the `CONTROL_FILE_RECORD_KEEP_TIME` setting. In this case, the database writes a message to the alert log. If you discover that this situation occurs frequently, then reducing the value of `CONTROL_FILE_RECORD_KEEP_TIME` and increase the frequency of resynchronizations.

Note: The maximum size of the control file is port-specific. Typically, the maximum size is 20,000 Oracle blocks. Refer to your platform-specific Oracle documentation for more information.

See Also:

- *Oracle Database Reference* for more information about the `CONTROL_FILE_RECORD_KEEP_TIME` parameter
- *Oracle Database Administrator's Guide* for more detailed information on other aspects of control file management

Working with RMAN Stored Scripts in the Recovery Catalog

Stored scripts offer an alternative to command files for managing frequently used sequences of RMAN commands.

The chief advantage of a stored script over a command file is that a stored script is always available to any RMAN client that can connect to the target database and recovery catalog, whereas command files are only available if the RMAN client has access to the file system on which they are stored.

Stored scripts can be global or local. A local stored script is associated with the target database to which RMAN is connected when the script is created, and can only be executed when you are connected to that target database. A global stored script can be run against any database registered in the recovery catalog, if the RMAN client is connected to the recovery catalog and a target database.

Note that to work with stored scripts, even global ones, you must be connected to both a recovery catalog and a target instance.

Creating Stored Scripts: CREATE SCRIPT

Make sure RMAN is connected to the right target database and the recovery catalog. Then run the `CREATE SCRIPT` command, as shown in this example:

```
CREATE SCRIPT full_backup
{
  BACKUP DATABASE PLUS ARCHIVELOG;
  DELETE OBSOLETE;
}
```

Examine the output. If no errors are displayed, then the script was successfully created and stored in the recovery catalog.

For a global script, the syntax is similar:

```
CREATE GLOBAL SCRIPT global_full_backup
{
  BACKUP DATABASE PLUS ARCHIVELOG;
  DELETE OBSOLETE;
}
```

You can also provide a COMMENT with descriptive information:

```
CREATE GLOBAL SCRIPT global_full_backup
COMMENT 'use only with ARCHIVELOG mode databases'
{
  BACKUP DATABASE PLUS ARCHIVELOG;
  DELETE OBSOLETE;
}
```

Finally, you can create a local or global script, reading its contents from a text file:

```
CREATE SCRIPT full_backup FROM FILE 'my_script_file.txt';
```

The file must begin with a { character, contain a series of commands valid within a RUN block, and end with a } character. Otherwise, a syntax error is signalled, just as if the commands were entered at the keyboard.

Running Stored Scripts: EXECUTE SCRIPT

To run a stored script, connect to the target database and recovery catalog, and use EXECUTE SCRIPT. EXECUTE SCRIPT requires a RUN block, as shown:

```
RUN { EXECUTE SCRIPT
full_backup
; }
```

This command invokes a local script if one is with the name specified. If no local script is found, but there is a global script with the name specified, RMAN will execute the global script. You can also use EXECUTE GLOBAL SCRIPT to control which script is invoked if a local and a global script have the same name. Assuming there is no local script called global_full_backup, the following two commands have the same effect:

```
RUN { EXECUTE GLOBAL SCRIPT global_
full_backup
; }
RUN { EXECUTE SCRIPT global_
full_backup
; }
```

Executing a global script only affects the connected target database; to run a global script across multiple databases, you must connect the RMAN client to each one separately and execute the script.

Your script will use the automatic channels configured at the time you execute the script. Use ALLOCATE CHANNEL commands in the script if you need to override the configured channels. Note that, because of the RUN block, if an RMAN command in the script fails, subsequent RMAN commands in the script will not execute.

See Also: *Oracle Database Backup and Recovery Reference* for EXECUTE SCRIPT command syntax

Displaying a Stored Script: PRINT SCRIPT

The PRINT SCRIPT command displays a stored script or writes it out to a file. With RMAN connected to the target database and recovery catalog, use the PRINT SCRIPT command as shown here:

```
PRINT SCRIPT full_backup;
```

To send the contents of a script to a file, use this form of the command:

```
PRINT SCRIPT full_backup TO FILE 'my_script_file.txt';
```

For global scripts, the analogous syntax would be:

```
PRINT GLOBAL SCRIPT global_full_backup;
```

and

```
PRINT GLOBAL SCRIPT global_full_backup TO FILE 'my_script_file.txt';
```

See Also: *Oracle Database Backup and Recovery Reference* for PRINT SCRIPT command syntax

Listing Stored Scripts: LIST SCRIPT NAMES

Use the LIST SCRIPT NAMES command to display the names of scripts defined in the recovery catalog. This command displays the names of all stored scripts, both global and local, that can be executed for the currently connected target database:

```
LIST SCRIPT NAMES;
```

If RMAN is not connected to a target database when the LIST SCRIPT NAMES command is run, then RMAN will respond with an error.

To view only global script names, use this form of the command:

```
LIST GLOBAL SCRIPT NAMES;
```

To view the names of all scripts stored in the current recovery catalog, including global scripts and local scripts for all target databases registered in the recovery catalog, use this form of the command:

```
LIST ALL SCRIPT NAMES;
```

The output will indicate for each script listed which target database the script is defined for (or whether a script is global).

Note: LIST GLOBAL SCRIPT NAMES and LIST ALL SCRIPT NAMES are the only commands that work when RMAN is connected to a recovery catalog without connecting to a target instance.

See Also: *Oracle Database Backup and Recovery Reference* for LIST SCRIPT NAMES command syntax and output format.

Updating Stored Scripts: REPLACE SCRIPT

To update stored scripts, connect to the target database and recovery catalog and use the `REPLACE SCRIPT` command. If the script does not already exist, then RMAN creates it.

This command updates stored script `full_backup` with new contents:

```
REPLACE SCRIPT full_backup
{
  BACKUP DATABASE PLUS ARCHIVELOG;
}
```

Global scripts can be updated using the `REPLACE GLOBAL SCRIPT` command when connected to a recovery catalog, as follows:

```
REPLACE GLOBAL SCRIPT global_full_backup
COMMENT 'A script for full backup to be used with any database'
{
  BACKUP AS BACKUPSET DATABASE PLUS ARCHIVELOG;
}
```

As with `CREATE SCRIPT`, you can update a local or global stored script from a text file, with this form of the command:

```
REPLACE GLOBAL SCRIPT global_full_backup FROM FILE 'my_script_file.txt';
```

See Also: *Oracle Database Backup and Recovery Reference* for `REPLACE SCRIPT` command syntax

Deleting Stored Scripts: DELETE SCRIPT

To delete a stored script from the recovery catalog, connect to the catalog and a target database, and use the `DELETE SCRIPT` command:

```
DELETE SCRIPT 'full_backup';
```

To delete a global stored script, use `DELETE GLOBAL SCRIPT`:

```
DELETE GLOBAL SCRIPT 'global_full_backup';
```

If you use `DELETE SCRIPT` without `GLOBAL`, and there is no stored script for the target database with the specified name, RMAN will look for a global stored script by the specified name and delete the global script if it exists. So, if you were to enter the command

```
DELETE SCRIPT 'global_full_backup';
```

RMAN would look for a script `'global_full_backup'` defined for the connected target database, and if it did not find one, it would search the global scripts for a script called `'global_full_backup'` and delete that script.

See Also: *Oracle Database Backup and Recovery Reference* for `DELETE SCRIPT` command syntax

Starting the RMAN Client and Running a Stored Script

To run the RMAN client and start a stored script in the recovery catalog on startup, use the `SCRIPT` argument when starting the RMAN client.

```
% rman TARGET SYS/oracle@trgt CATALOG rman/cat@catdb SCRIPT 'full_backup';
```


You must connect to a recovery catalog (which contains the stored script) and target database (to which the script will apply) when starting the RMAN client.

See Also: *Oracle Database Backup and Recovery Reference* for full RMAN client command line syntax

Restrictions on Stored Script Names

There are some issues to be aware of about how RMAN resolves script names, especially when a local and global script share the same name.

- RMAN permits but generally does not require that you use quotes around the name of a stored script. However, if the name begins with a digit or if the name is an RMAN reserved word, you will have to put quotes around the name to use it as a stored script name. Consider avoiding stored script names that begin with characters other than A-Z or that are the same as RMAN reserved words.

See Also: *Oracle Database Backup and Recovery Reference*

- When starting the RMAN client with a SCRIPT argument on the command line, if local and global scripts are defined with the same name, then RMAN will always execute the local script.
- For the EXECUTE SCRIPT, DELETE SCRIPT and PRINT SCRIPT commands, if the script name passed as an argument is not the name of a script defined for the connected target instance, RMAN will look for a global script by the same name to execute, delete or print. For example, if the a stored script global_full_backup is in the recovery catalog as a global script, but no local stored script global_full_backup is defined for the target database, the following command will delete the global script:

```
DELETE SCRIPT global_full_backup;
```

Consider using some naming convention to avoid mistakes due to confusion between global stored scripts and local stored scripts.

See Also: *Oracle Database Backup and Recovery Reference* for the list of RMAN reserved words.

Backing Up and Recovering the Recovery Catalog

Include the recovery catalog database in your backup and recovery strategy. If you do not back up the recovery catalog and a disk failure occurs that destroys the recovery catalog database, then you may lose the metadata in the catalog. Without the recovery catalog contents, recovery of your other databases is likely to be more difficult.

Backing Up the Recovery Catalog

Here are some general guidelines you should follow when developing a strategy for backing up the recovery catalog.

Back Up the Recovery Catalog Often

The recovery catalog database is a database like any other, and is also a key part of your backup and recovery strategy. Protect the recovery catalog as you would protect any other part of your database, by backing it up. The backup strategy for your recovery catalog database should be part of your overall backup and recovery strategy.

Back up the recovery catalog with the same frequency that you back up the target database. For example, if you make a weekly whole database backup of the target database, then back up the recovery catalog immediately after all target database backups, in order to protect the record of the whole database backup. This backup can help you in a disaster recovery scenario. Even if you have to restore the recovery catalog database using a control file autobackup, you can then use the full record of backups in your restored recovery catalog database to restore the target database without using a control file autobackup for the target database.

Choosing the Appropriate Method for Physical Backups

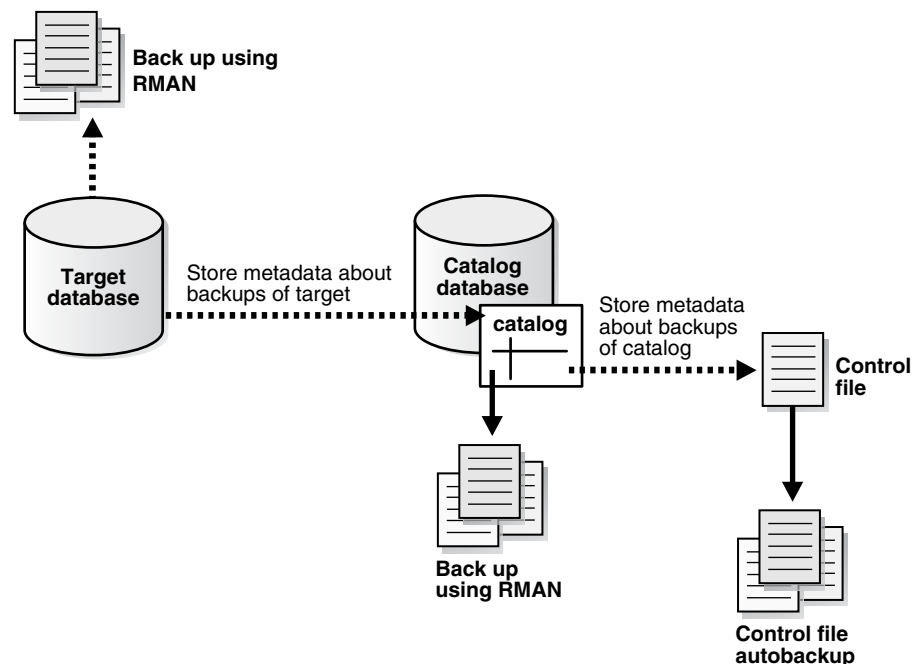
When backing up the recovery catalog database, you can use RMAN to make the backups. As illustrated in [Figure 10-1](#), you should start RMAN with the `NOCATALOG` option so that the repository for the recovery catalog is the control file in the catalog database.

Follow these guidelines when developing an RMAN backup strategy for the recovery catalog database:

- Run the recovery catalog database in `ARCHIVELOG` mode so that you can do point-in-time recovery if needed.
- Set the retention policy to a `REDUNDANCY` value greater than 1.
- Back up the database onto two separate media (for example, disk and tape).
- Run `BACKUP DATABASE PLUS ARCHIVELOG` at regular intervals, to a media manager if available, or just to disk.
- Do not use another recovery catalog as the repository for the backups.
- Configure the control file autobackup feature to `ON`.

With this strategy, the control file autobackup feature ensures that the recovery catalog database can always be recovered, so long as the control file autobackup is available.

Figure 10-1 Using the Control File as the Repository for Backups of the Recovery Catalog



See Also: ["Performing Disaster Recovery"](#) on page 7-10 for more information for recovery with a control file autobackup

Safe Storage of the Recovery Catalog

Never store a recovery catalog containing the RMAN repository for a database in the same database as the target database or on the same disks as the target database. For example, do not store the catalog for database `prod1` in `prod1`. A recovery catalog for `prod1` is only effective if it is separated from the data that it is designed to protect.

If `prod1` suffers a total media failure, and if the recovery catalog for `prod1` is also stored in `prod1`, then if you lose the database you also lose the recovery catalog. At that point the only option is to restore an autobackup of the control file for `prod1` and use it to restore and recover the database without the benefit of any information stored in the recovery catalog.

Separating the target and catalog databases is especially important when you back up a recovery catalog database. The following example shows what you should *not* do. For example, consider the following:

- Target database `prod1` and catalog database `catdb` are on different hosts.
- `catdb` contains the recovery catalog repository for target database `prod1`.

You decide to use a recovery catalog to back up `catdb`, but are not sure where to create it. If you create the catalog containing the repository for `catdb` in database `catdb`, then if you lose `catdb` due to a media failure, you will have difficulty restoring `catdb` *and* will leave `prod1` without a recovery catalog to use in a restore scenario.

Exporting the Recovery Catalog Data for Logical Backups

Logical backups of the RMAN recovery catalog created with one of Oracle's export utilities can be a useful supplement for physical backups. In the event of damage to a recovery catalog database, you can quickly reimport the exported recovery catalog data into another database and rebuild the catalog this way.

Restoring and Recovering the Recovery Catalog from Backup

Restoring and recovering the recovery catalog is much like restoring and recovering any other database, if you backed it up with RMAN.

You can restore the control file and SPFILE for the recovery catalog database from an autobackup, then restore and perform complete recovery on the rest of the database. The processes required are all described in *Oracle Database Backup and Recovery Basics*. You can also use another recovery catalog to record metadata about backups of this recovery catalog database, if you are in a situation where you are using multiple recovery catalogs.

Re-Creating the Recovery Catalog

If the recovery catalog database is lost or damaged, and recovery of the recovery catalog database through the normal Oracle recovery procedures is not possible, then you must re-create the catalog. Examples of this worst-case scenario include:

- A recovery catalog database that has never been backed up
- A recovery catalog database that has been backed up, but cannot be recovered because the datafile backups or archived logs are not available

You have these options for partially re-creating the contents of the missing recovery catalog:

- Use the `RESYNC CATALOG` command to update the recovery catalog with any RMAN repository information from the control file of the target database or a control file copy. Note that any metadata from control file records that aged out of the control file is lost.
- Issue `CATALOG START WITH . . .` commands to recatalog any available backups.

To minimize the likelihood of this worst-case scenario, your backup strategy should at least include backing up the recovery catalog using RMAN as described in ["Backing Up the Recovery Catalog"](#) on page 10-17.

See Also:

- *Oracle Database Backup and Recovery Reference* for information about the `CATALOG` command
- *Oracle Database Backup and Recovery Reference* for information about the `CROSSCHECK` command
- ["Managing the Control File When You Use a Recovery Catalog"](#) on page 10-12 to learn about how records age out of the control file

Exporting and Importing the Recovery Catalog

To move the recovery catalog from one database to another, export the catalog from the old database, and import it into the new one. You can only import the catalog into a supported version of the Oracle database server. In general, you can import the catalog into a database of the same release or later.

Exports can also serve as logical backups of the RMAN recovery catalog. If the recovery catalog database is damaged, you can quickly reimport the exported recovery catalog data into another database and rebuild the catalog.

This section contains the following topics:

- [Considerations When Moving Catalog Data](#)
- [Exporting the Recovery Catalog](#)
- [Importing the Recovery Catalog](#)

Considerations When Moving Catalog Data

You should only import the recover catalog into a schema that does not already contain a recovery catalog schema. In other words, the user who will own the imported recovery catalog schema should not already own a recovery catalog schema. For example, if user `rman` owns the recovery catalog on database `catdb`, and you want to export the recovery catalog on `catdb` and import it into database `catdb2`, then `rman` should not already own a recovery catalog on `catdb2`. You should either create a new recovery catalog owner on `catdb2`, or drop the current user `rman` on `catdb2` and then re-create the user. You cannot merge a recovery catalog into an existing recovery catalog.

The basic steps for exporting a recovery catalog from a primary database and importing the catalog into a secondary database are as follows:

1. Use one of the Oracle Export utilities to export the catalog data from the primary database. See ["Exporting the Recovery Catalog"](#) on page 10-21 for an example.

2. Create a user on the secondary database as described in "[Creating the Recovery Catalog Owner](#)" on page 10-3, and grant the user necessary privileges.
3. Use the import utility corresponding to the export utility in step 1 to import the catalog data into the schema created in the previous step. See "[Importing the Recovery Catalog](#)" on page 10-21 for an example.

You should *not* run the CREATE CATALOG command either before or after the Import of the catalog into the secondary database. By importing the catalog data into the new schema, you effectively create the catalog in the secondary database.

Note: You cannot import data exported from two different recovery catalogs to merge them into one catalog. It is not currently possible to merge two or more recovery catalog schemas into one.

Exporting the Recovery Catalog

This example uses the Original Export utility described in *Oracle Database Utilities* to create a logical export of the recovery catalog. Refer to *Oracle Database Utilities* for concepts and procedures relating to the Data Pump Export utility.

The following procedure creates a logical export of the recovery catalog.

1. Execute the Oracle export utility at the operating system command line, making sure to do the following:
 - a. Connect as the owner of the recovery catalog
 - b. Specify the OWNER option
 - c. Specify an output file

For example, if the owner of the catalog in database `catdb` is `rman`, you can issue the following at the UNIX command line to export the catalog to file `cat.dmp`:

```
% exp rman/cat@catdb FILE=cat.dmp OWNER=rman
```

2. Examine the output to make sure you were successful:

Export terminated successfully without warnings.

Importing the Recovery Catalog

This example uses the Original Import utility described in *Oracle Database Utilities* to create a logical export of the recovery catalog. Refer to *Oracle Database Utilities* for concepts and procedures relating to the Data Pump Import utility.

To make a logical import of the recovery catalog from the command line:

1. Create a new user in another database. For the recommended SQL syntax for creating a new user in a recovery catalog database, see "[Creating the Recovery Catalog Owner](#)" on page 10-3. Be sure to grant the new user the necessary privileges.
2. Import the catalog data from the export file. Execute the import at the command line, making sure to do the following:
 - a. Connect as the new owner of the recovery catalog.
 - b. Specify the old owner with the FROMUSER parameter.
 - c. Specify the new owner with the TOUSER parameter.

- d. Specify the import file.

For example, assume the following:

- The old owner of the catalog in database `prod1` is `rman`.
- The user in the new recovery catalog database `catdb2` is `rman2`.
- The file containing the export of the catalog is `cat.dmp`.

The command is then as follows:

```
% imp USERID=rman2/cat2@catdb2 FILE=cat.dmp FROMUSER=rman TOUSER=rman2
```

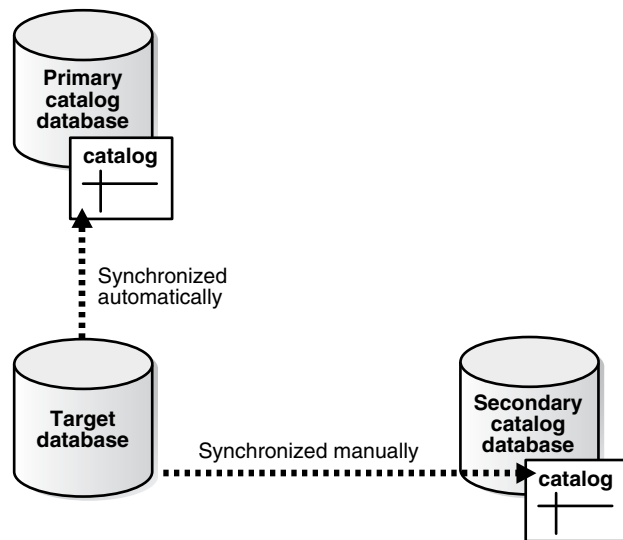
3. Use the imported catalog data for restore and recovery of your target database.

Increasing Availability of the Recovery Catalog

You may have a production system in which you want to maintain high availability for the catalog database. For example, you may have 100 target databases registered in the recovery catalog. In case the primary catalog database goes down, you can create redundancy by storing a secondary recovery catalog in a separate database, as illustrated in [Figure 10-2](#). You must register the target database in the secondary catalog.

In this availability scenario, the main catalog is synchronized as normal during regular backups, while the secondary catalog is synchronized periodically with the `RESYNC CATALOG` command. If the primary catalog database goes down or requires routine maintenance, then you can resynchronize the secondary catalog and use it as the new primary catalog during the interim.

Figure 10-2 Registering One Target Database in Two Recovery Catalogs



See Also: ["Creating a Recovery Catalog"](#) on page 10-1

Querying Recovery Catalog Views

The `LIST`, `REPORT`, and `SHOW` commands provide the easiest means of accessing the data in the control file and the recovery catalog. Nevertheless, you can sometimes also

obtain useful information from the recovery catalog views, which are views in the catalog schema prefixed with RC_.

See Also: *Oracle Database Backup and Recovery Reference* for reference information about the recovery catalog views

RMAN obtains backup and recovery metadata from the target database control file and stores it in the tables of the recovery catalog. The recovery catalog views are derived from these tables. Note that the recovery catalog views are not normalized or optimized for user queries.

In general, the recovery catalog views are not as user-friendly as the RMAN reporting commands. For example, when you start RMAN and connect to a target database, you obtain the information for this target database only when you issue LIST, REPORT, and SHOW commands. If you have 10 different target databases registered in the same recovery catalog, then any query of the recovery catalog views show the information for all incarnations of all databases registered in the catalog. You often have to perform complex selects and joins among the views to extract usable information about a specific database and incarnation.

Most of the catalog views have a corresponding dynamic performance view (or V\$ view) in the database server. For example, RC_BACKUP_PIECE corresponds to V\$BACKUP_PIECE. The primary difference between the recovery catalog view and corresponding server view is that each catalog view contains information about all the databases registered in the catalog, whereas the database server view contains information only about itself. The RC_ views and corresponding V\$ views use different primary keys to uniquely identify rows.

Identifying Rows for a Database in the Catalog Views

Most of the catalog views contain the columns DB_KEY and DBINC_KEY. Each target database can be uniquely identified by either the primary key, which is the DB_KEY column value, or the DBID, which is the 32-bit unique database identifier. Each incarnation of a target database is uniquely identified by the DBINC_KEY primary key. When querying data about a specific incarnation of a target database, you should use these columns to specify the database. Then, you can perform joins with most of the other catalog views to obtain the desired information.

Identifying Rows for a Database Object in the Catalog Views

An important difference between catalog and V\$ views is that a different system of unique identifiers is used for backup and recovery objects. For example, many V\$ views such as V\$ARCHIVED_LOG use the RECID and STAMP columns to form a concatenated primary key. The corresponding catalog view uses a derived value as its primary keys and stores this value in a single column. For example, the primary key in RC_ARCHIVED_LOG is the AL_KEY column. The AL_KEY column value is the primary key that RMAN displays in the LIST command output.

Querying Catalog Views for the Target DB_KEY or DBID Values

The DB_KEY value, which is the primary key for a target database, is used only in the recovery catalog. The easiest way to obtain the DB_KEY is to use the DBID of the target database, which is displayed whenever you connect RMAN to the target database. The DBID, which is a unique system-defined number given to every Oracle database, is used to distinguish among databases registered in the RMAN recovery catalog.

Assume that you want to obtain information about one of the target databases registered in the recovery catalog. You can easily determine the DBID from this database either by looking at the output displayed when RMAN connects to the database, querying `V$RMAN_OUTPUT`, or querying a `V$DATABASE` view as in the following:

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

```
DBID
-----
598368217
```

You can then obtain the `DB_KEY` for a target database by running the following query, where `dbid_of_target` is the DBID that you previously obtained:

```
SELECT DB_KEY
FROM RC_DATABASE
WHERE DBID = dbid_of_target;
```

To obtain information about the current incarnation of a target database, specify the target database `DB_KEY` value and perform a join with `RC_DATABASE_INCARNATION` by using a `WHERE` condition to specify that the `CURRENT_INCARNATION` column value is set to `YES`. For example, to obtain information about backup sets in the current incarnation of a target database with the `DB_KEY` value of 1, you can execute the following script:

```
SELECT BS_KEY, BACKUP_TYPE, COMPLETION_TIME
FROM RC_DATABASE_INCARNATION i, RC_BACKUP_SET b
WHERE i.DB_KEY = 1
AND i.DB_KEY = b.DB_KEY
AND i.CURRENT_INCARNATION = 'YES';
```

You should use the `DB_NAME` column to specify a database *only if* you do not have more than one database registered in the recovery catalog with the same `DB_NAME`. RMAN permits you to register more than one database with the same database name, but requires that the DBID values be different. For example, you can have ten databases with the `DB_NAME` value of `prod1`, each with a different DBID. Because the DBID is the unique identifier for every database in the metadata, use this value to obtain the `DB_KEY` and then use `DB_KEY` to uniquely identify the database.

Using `RC_BACKUP_FILES` and `DBMS_RCVMAN.SETDATABASE`

The view `RC_BACKUP_FILES` can be queried for information about all backups of any database registered in the recovery catalog. However, before querying `RC_BACKUP_FILES` you must call `DBMS_RCVMAN.SETDATABASE`, specifying the DBID of one of the databases registered in the recovery catalog, as shown in the following example:

```
SQL> CALL DBMS_RCVMAN.SETDATABASE(null,null,null,2283997583,null);
```

The fourth parameter must be the DBID of a database registered in the recovery catalog. The other parameters must all be `NULL`.

- See also:**
- *Oracle Database Backup and Recovery Reference* for details about the `RC_BACKUP_FILES` view
 - *Oracle Database Backup and Recovery Basics* for methods of determining the DBID of a database

Determining the Schema Version of the Recovery Catalog

The schema version of the recovery catalog is stored in the recovery catalog itself. The information is important in case you maintain multiple databases of different versions in your production system, and need to determine whether the catalog schema version is usable with a specific target database version.

See Also: *Oracle Database Backup and Recovery Reference* for the complete set of compatibility rules governing the RMAN environment

To determine the schema version of the recovery catalog:

1. Start SQL*Plus and connect to the recovery catalog database as the catalog owner. For example:

```
% sqlplus rman/cat@catdb
```

2. Query the RCVER table to obtain the schema version, as in the following example (sample output included):

```
SELECT *
FROM rcver;
```

```
VERSION
-----
09.02.00
```

If the table displays multiple rows, then the highest version in the RCVER table is the current catalog schema version. The table stores only the major version numbers and not the patch numbers. For example, assume that the rcver table displays the following rows:

```
VERSION
-----
08.01.07
09.00.01
09.02.00
```

These rows indicate that the catalog was created with a release 8.1.7 executable, then upgraded to release 9.0.1, and finally upgraded to release 9.2.0. The current version of the catalog schema is 9.2.0.

Upgrading the Recovery Catalog

If you use a version of the recovery catalog that is older than that required by the RMAN client, then you must upgrade it. For example, you must upgrade the catalog if you use a release 8.1 version of the RMAN client with a release 8.0 version of the recovery catalog.

You receive an error when issuing UPGRADE CATALOG if the recovery catalog is already at a version greater than that required by the RMAN client. RMAN permits the UPGRADE CATALOG command to be run if the recovery catalog is current and does not require upgrading, however, so that you can re-create packages at any time if necessary. Check the message log for error messages generated during the upgrade.

To upgrade the recovery catalog:

1. To install the new recovery catalog schema, the recovery catalog user must have TYPE privilege:

```
sqlplus> connect sys/oracle@catdb as sysdba;
sqlplus> grant TYPE to rman;
```

2. Use RMAN to connect to the target and recovery catalog databases. For example, enter:

```
% rman TARGET / CATALOG rman/cat@catdb

connected to recovery catalog database
PL/SQL package rcat.DBMS_RCVCAT version 08.00.04 in RCVCAT database
is too old
```

3. Issue the UPGRADE CATALOG command:

```
UPGRADE CATALOG;

recovery catalog owner is rman
enter UPGRADE CATALOG command again to confirm catalog upgrade
```

4. Enter the UPDATE CATALOG command again to confirm:

```
UPGRADE CATALOG;

recovery catalog upgraded to version 09.02.00
DBMS_RCVMAN package upgraded to version 09.02.00
DBMS_RCVCAT package upgraded to version 09.02.00
```

See Also:

- *Oracle Database Backup and Recovery Reference* for UPGRADE CATALOG command syntax
- *Oracle Database Backup and Recovery Reference* for information about recovery catalog compatibility
- *Oracle Database Upgrade Guide* for complete compatibility and migration information

Dropping the Recovery Catalog

If you do not want to maintain a recovery catalog, then you can drop the recovery catalog schema from the tablespace. The `DROP CATALOG` command deletes all information from the recovery catalog. Hence, if you have no backups of the recovery catalog schema, then backups of all target databases managed by this catalog may become unusable. (The control file of the target database will still retain a record of recent backups.)

The `DROP CATALOG` command is not appropriate for unregistering a single database from a recovery catalog that has multiple target databases registered. Dropping the catalog deletes the recovery catalog record of backups for all target databases registered in the catalog.

To drop the recovery catalog schema:

1. Use RMAN to connect to the target and recovery catalog databases.

```
% rman TARGET / CATALOG rman/cat@catdb
```

2. Issue the `DROP CATALOG` command twice to confirm:

```
DROP CATALOG;

recovery catalog owner is rman
```

enter DROP CATALOG command again to confirm catalog removal

DROP CATALOG;

Note: Even after you drop the recovery catalog, the control file still contains records about the backups. To purge RMAN repository records from the control file, re-create the control file.

See Also: *Oracle Database Backup and Recovery Reference* for DROP CATALOG command syntax, and "[Unregistering a Target Database from the Recovery Catalog](#)" on page 10-7 to learn how to unregister a database from the catalog

Tuning Backup and Recovery

Tuning RMAN performance is mostly a matter of maximizing the speed with which RMAN creates your backups and restores from backups, on disk and especially on tape. A secondary concern is limiting the effect of backup activities on database throughput.

You may also need to tune performance of the database during instance recovery.

This chapter covers the concepts needed for performance tuning, and the features in RMAN that can help you. The discussion is divided into the following sections:

- [Tuning Recovery Manager: Overview](#)
- [Features and Options Used to Tune RMAN Performance](#)
- [Tuning RMAN Backup Performance: Procedure](#)
- [Instance Recovery Performance Tuning: Fast-Start Fault Recovery](#)

Tuning Recovery Manager: Overview

RMAN backup and restore operations have the following distinct components:

- Reading or writing input data
- Processing data by validating blocks and copying them from the input to the output buffers

The slowest of these operations in any RMAN task is called the **bottleneck**. RMAN tuning involves identifying the bottlenecks for a given task and using RMAN commands, initialization parameter settings, or adjustments to physical media to improve performance on the backup.

The key to tuning RMAN is understanding how it performs I/O. RMAN's backup and restore jobs use two types of I/O buffers: `DISK` and tertiary storage (usually tape). When performing a backup, RMAN reads input files using disk buffers and writes the output backup file by using either disk or tape buffers. Restore operations use disk or tape buffers for input, depending on where the backup is stored, and disk buffers for output.

To tune RMAN effectively, you must thoroughly understand concepts such as synchronous and asynchronous I/O, disk and tape buffers, and channel architecture. When you understand these concepts, then you can learn how to use fixed views to monitor bottlenecks, and use the techniques described in "[Tuning RMAN Backup Performance: Procedure](#)" on page 11-6 to solve problems.

There are a number of concepts that affect RMAN performance and that can therefore influence your strategy for backup performance tuning:

- [I/O Buffer Allocation](#)
- [Allocation of Tape Buffers](#)
- [Synchronous and Asynchronous I/O](#)
- [Factors Affecting Backup Speed to Tape](#)
- [Using the RATE Parameter to Control Disk Bandwidth Usage](#)

I/O Buffer Allocation

RMAN I/O uses two different types of buffers: disk and tape. These buffers are typically different sizes. They are allocated differently, depending upon the device type and the role the buffer plays in an RMAN operation.

Allocation for Disk Buffers

To understand how RMAN allocates buffers to read datafiles during backups, you must understand how RMAN multiplexing works.

RMAN **multiplexing** is RMAN's ability to read a number of files in a backup simultaneously from different sources to improve reading performance, and then write them to a single backup piece. The level of multiplexing is the number of files read simultaneously.

Multiplexing is described at greater length in "[Multiplexed Backup Sets](#)" on page 2-12. The level of multiplexing is determined by the algorithm described in "[Algorithm for Multiplexed Backups](#)" on page 2-12. Review this section before proceeding.

When RMAN backs up from disk, it uses the rules described in the following table to determine how large to make the buffers.

Table 11–1 Datafile Read Buffer Sizing Algorithm

Level of Multiplexing	Resulting Buffer Size
Less than or equal to 4	RMAN allocates buffers of size 1 MB so that the total buffer size for all the input files is 16 MB.
Greater than 4 but less than or equal to 8	RMAN allocates disk buffers of size 512 KB so that the total buffer size for all the files is less than 16 MB.
Greater than 8	RMAN allocates a fixed 4 disk buffers of 128 KB for each file, so that the total size is 512 KB for each file.

The number of buffers allocated depends on the following rules:

- When the output of the backup resides on disk, 4 buffers are allocated, their size being operating system dependent.
- If the operation is a restore, and the backup resides on disk, 4 buffers are allocated, their size being operating system dependent.
- When restoring a backup, for each active datafile 4 buffers of 128K are allocated.
- When image copies are produced, only 4 buffers in total are allocated, each of an operating system dependent size.

Allocation of Tape Buffers

If you backup to or restore from an sbt device, then by default the database allocates four buffers for each channel for the tape writers (or reads if doing a restore). The size of these buffers is platform dependent, but is typically 256K. This value can be

changed using the `ALLOCATE` or `SEND` command using the `PARMS` and the `BLKSIZE` option.

RMAN allocates the tape buffers in the SGA if I/O slaves are being used, or the PGA otherwise.

If you use I/O slaves, then set the `LARGE_POOL_SIZE` initialization parameter to set aside SGA memory dedicated to holding these large memory allocations. This prevents RMAN I/O buffers from competing with the library cache for SGA memory. If I/O slaves for tape I/O were requested but there is not enough space in the SGA for them, slaves are not used, and a message appears in the alert log.

Synchronous and Asynchronous I/O

When RMAN reads or writes data, the I/O is either **synchronous** or **asynchronous**. When the I/O is synchronous, a server process can perform only one task at a time. When it is asynchronous, a server process can begin an I/O and then perform other work while waiting for the I/O to complete. It can also begin multiple I/O operations before waiting for the first to complete.

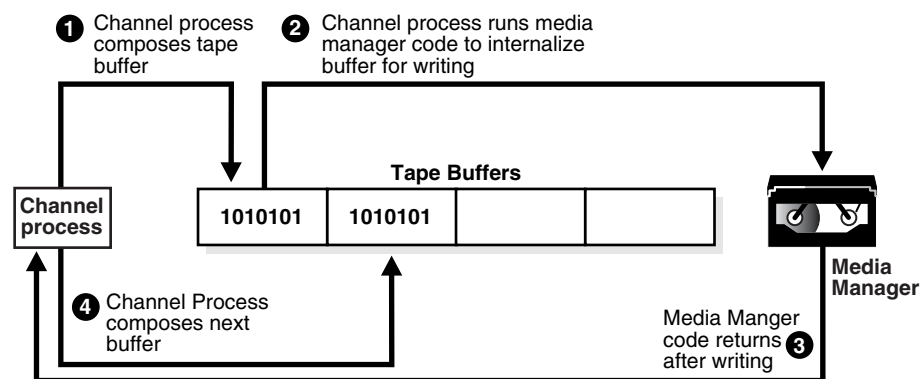
Some operating systems support native asynchronous disk I/O. The database takes advantage of this feature if it is available. On operating systems that do not support native asynchronous I/O, the database can simulate it with special I/O slave processes that are dedicated to performing I/O on behalf of another process. You can control disk I/O slaves by setting the `DBWR_IO_SLAVES` parameter to a nonzero value. The database allocates four backup disk I/O slaves for any nonzero value of `DBWR_IO_SLAVES`.

By contrast, tape I/O is always synchronous. For tape I/O, each channel allocated (whether manually or based on a `CONFIGURE` command) corresponds to a server process, called here a channel process.

Synchronous I/O: Example

Figure 11-1 shows synchronous I/O in a backup to tape.

Figure 11-1 Synchronous I/O



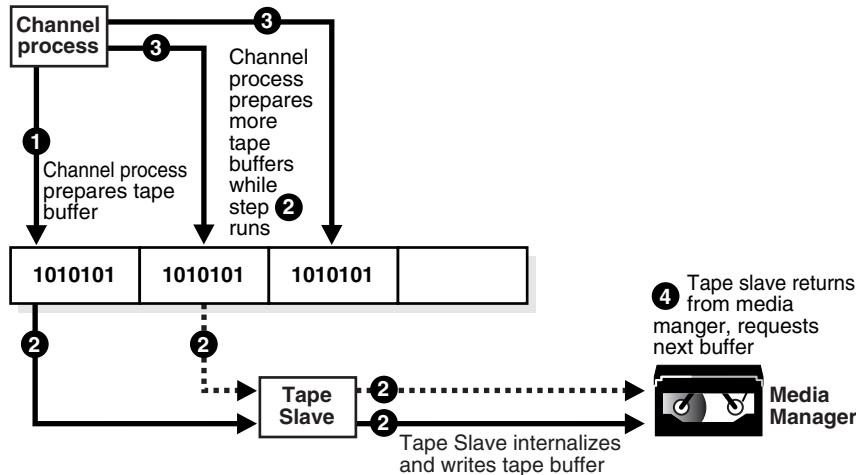
The following steps occur:

1. The channel process composes a tape buffer.
2. The channel process executes media manager code that processes the tape buffer and internalizes it for further processing and storage by the media manager.

3. The media manager code returns a message to the server process stating that it has completed writing.
4. The channel process can initiate a new task.

Figure 11–2 shows **asynchronous I/O** in a tape backup. Asynchronous I/O to tape is simulated by using tape slaves. In this case, each allocated channel corresponds to a server process, which in the explanation which follows is identified as a channel process. For each channel process, one tape slave is started (or more than one, in the case of multiple copies).

Figure 11–2 Asynchronous I/O



The following steps occur:

1. A channel process writes blocks to a tape buffer.
2. The channel process sends a message to the tape slave process to process the tape buffer. The tape slave process executes media manager code that processes the tape buffer and internalizes it so that the media manager can process it.
3. While the tape slave process is writing, the channel process is free to read data from the datafiles and prepare more output buffers.
4. Once the tape slave channel returns from the media manager code, it requests a new tape buffer, which usually is ready. Thus waiting time for the channel process is reduced, and the backup is completed faster.

Factors Affecting Backup Speed to Tape

The following factors affect the speed of the backup to tape:

- [Native Transfer Rate](#)
- [Tape Compression](#)
- [Tape Streaming](#)
- [Physical Tape Block Size](#)

Native Transfer Rate

The tape native transfer rate is the speed of writing to a tape without compression. This speed represents the upper limit of the backup rate. The upper limit of your backup performance should be the aggregate transfer rate of all of your tape drives. If

your backup is already performing at that rate, and if it is not using an excessive amount of CPU, then RMAN performance tuning will not help.

Tape Compression

The level of tape compression is very important for backup performance. If the tape has good compression, then the sustained backup rate is faster. For example, if the compression ratio is 2:1 and native transfer rate of the tape drive is 6 MB/s, then the resulting backup speed is 12 MB/s. In this case, RMAN must be able to read disks with a throughput of more than 12 MB/s or the disk becomes the bottleneck for the backup.

Note: You should not use both tape compression provided by the media manager and binary backupset compression as provided by RMAN. If the media manager compression is efficient, then it is usually the better choice. Using RMAN compressed backupsets can be an effective alternative if you need to reduce bandwidth used to move uncompressed backupsets over a network to the media manager, and if the CPU overhead required to compress the data in RMAN is acceptable.

See *Oracle Database Backup and Recovery Basics* for more on using compressed backupsets.

Tape Streaming

Tape **streaming** during write operations has a major impact on tape backup performance. Almost all tape drives currently on the market are fixed-speed, streaming tape drives. Because such drives can only write data at one speed, when they run out of data to write to tape, the tape must slow down and stop. Generally, when the drive's buffer empties, the tape is moving so quickly that it actually overshoots; to continue writing, the drive must rewind the tape to locate the point where it stopped writing.

Physical Tape Block Size

The physical tape block size can affect backup performance. The block size is the amount of data written by media management software to a tape in one write operation. In general, the larger the tape block size, the faster the backup. Note that physical tape block size is not controlled by RMAN or the Oracle database server, but by media management software. See your media management software's documentation for details.

Features and Options Used to Tune RMAN Performance

There are a number of features you can use to tune your backup performance, once you have sufficient knowledge of your database and its workload and bottlenecks.

Using the RATE Parameter to Control Disk Bandwidth Usage

The RATE parameter specifies the bytes/second that RMAN reads on this channel. Use this parameter to set an upper limit for bytes read so that RMAN does not consume excessive disk bandwidth and degrade online performance.

For example, set RATE=1500K. If each disk drive delivers 3 MB/second, then RMAN leaves some disk bandwidth available to the online system.

Tuning RMAN Backup Performance: Procedure

Many factors can affect backup performance. Often, finding the solution to a slow backup is a process of trial and error. To get the best performance for a backup, follow the suggested steps in this section:

- [Step 1: Remove RATE Parameters from Configured and Allocated Channels](#)
- [Step 2: If You Use Synchronous Disk I/O, Set DBWR_IO_SLAVES](#)
- [Step 3: If You Fail to Allocate Shared Memory, Set LARGE_POOL_SIZE](#)
- [Step 4: Tune RMAN Tape Streaming Performance Bottlenecks](#)
- [Step 5: Query V\\$ Views to Identify Bottlenecks](#)

Step 1: Remove RATE Parameters from Configured and Allocated Channels

The `RATE` parameter on a channel is intended to reduce, rather than increase, backup throughput, so that more disk bandwidth is available for other database operations.

If your backup is not streaming to tape, then make sure that the `RATE` parameter is not set on the `ALLOCATE CHANNEL` or `CONFIGURE CHANNEL` commands.

Step 2: If You Use Synchronous Disk I/O, Set DBWR_IO_SLAVES

If and only if your disk does *not* support asynchronous I/O, then try setting the `DBWR_IO_SLAVES` initialization parameter to a nonzero value. Any nonzero value for `DBWR_IO_SLAVES` causes a fixed number (four) of disk I/O slaves to be used for backup and restore, which simulates asynchronous I/O. If I/O slaves are used, I/O buffers are obtained from the SGA. The large pool is used, if configured. Otherwise, the shared pool is used.

Note: By setting `DBWR_IO_SLAVES`, the database writer processes will use slaves as well. You may need to increase the value of the `PROCESSES` initialization parameter.

Step 3: If You Fail to Allocate Shared Memory, Set LARGE_POOL_SIZE

Set this initialization parameter if the database reports an error in the `alert.log` stating that it does not have enough memory and that it will not start I/O slaves. The message should resemble the following:

```
ksfqxcre: failure to allocate shared memory means sync I/O will be used whenever
async I/O to file not supported natively
```

When attempting to get shared buffers for I/O slaves, the database does the following:

- If `LARGE_POOL_SIZE` is set, then the database attempts to get memory from the large pool. If this value is not large enough, then an error is recorded in the alert log, the database does not try to get buffers from the shared pool, and asynchronous I/O is not used.
- If `LARGE_POOL_SIZE` is not set, then the database attempts to get memory from the shared pool.
- If the database cannot get enough memory, then it obtains I/O buffer memory from the PGA and writes a message to the `alert.log` file indicating that synchronous I/O is used for this backup.

The memory from the large pool is used for many features, including the shared server (formerly called multi-threaded server), parallel query, and RMAN I/O slave buffers.

Configuring the large pool prevents RMAN from competing with other subsystems for the same memory.

Requests for contiguous memory allocations from the shared pool are usually small (under 5 KB) in size. However, it is possible that a request for a large contiguous memory allocation can either fail or require significant memory housekeeping to release the required amount of contiguous memory. Although the shared pool may be unable to satisfy this memory request, the large pool is able to do so. The large pool does not have a least recently used (LRU) list; the database does not attempt to age memory out of the large pool.

Use the `LARGE_POOL_SIZE` initialization parameter to configure the large pool. To see in which pool (shared pool or large pool) the memory for an object resides, query `V$SGASTAT.POOL`.

The formula for setting `LARGE_POOL_SIZE` is as follows:

$$\text{LARGE_POOL_SIZE} = \text{number_of_allocated_channels} * (\text{16 MB} + (4 * \text{size_of_tape_buffer}))$$

See Also: *Oracle Database Concepts* for more information about the large pool, and *Oracle Database Reference* for complete information about initialization parameters

Step 4: Tune RMAN Tape Streaming Performance Bottlenecks

There are several tasks you can perform to identify and remedy bottlenecks that affect RMAN's performance on tape backups:

Using BACKUP... VALIDATE To Distinguish Between Tape and Disk Bottlenecks

One reliable way to determine whether the tape streaming or disk I/O is the bottleneck in a given backup job is to compare the time required to run backup tasks with the time required to run `BACKUP VALIDATE` of the same tasks. `BACKUP VALIDATE` of a backup to tape performs the same disk reads as a real backup but performs no tape I/O. If the time required for the `BACKUP VALIDATE` to tape is significantly less than the time required for a real backup to tape, then writing to tape is the likely bottleneck.

Using Multiplexing to Improve Tape Streaming with Disk Bottlenecks

In some situations when performing a backup to tape, RMAN may not be able to send data blocks to the tape drive fast enough to support streaming. For example, during an incremental backup, RMAN only backs up blocks changed since a previous datafile backup as part of the same strategy. If you do not turn on change tracking, RMAN must scan entire datafiles for changed blocks, and fill output buffers as it finds such blocks. If there are not many changed blocks, RMAN may not fill output buffers fast enough to keep the tape drive streaming.

You can improve performance by increasing the degree of multiplexing used for backing up. This increases the rate at which RMAN fills tape buffers, which makes it more likely that buffers are sent to the media manager fast enough to maintain streaming.

Using Incremental Backups to Improve Backup Performance With Tape Bottlenecks

If writing to tape is the source of a bottleneck for your backups, consider using incremental backups as part of your backup strategy. Incremental level 1 backups write only the changed blocks from datafiles to tape, so that any bottleneck on writing

to tape has less impact on your overall backup strategy. In particular, if tape drives are not locally attached to the node running the database being backed up, then incremental backups can be faster.

Step 5: Query V\$ Views to Identify Bottlenecks

If none of the previous steps improves backup performance, then try to determine the exact source of the bottleneck. Use the `V$BACKUP_SYNC_IO` and `V$BACKUP_ASYNC_IO` views to determine the source of backup or restore bottlenecks and to see detailed progress of backup jobs.

`V$BACKUP_SYNC_IO` contains rows when the I/O is synchronous to the process (or thread on some platforms) performing the backup. `V$BACKUP_ASYNC_IO` contains rows when the I/O is asynchronous. Asynchronous I/O is obtained either with I/O processes or because it is supported by the underlying operating system.

See Also: *Oracle Database Reference* for more information about these views

To determine whether your tape is streaming when the I/O is synchronous, query the `EFFECTIVE_BYTES_PER_SECOND` column in the `V$BACKUP_SYNC_IO` or `V$BACKUP_ASYNC_IO` view. If `EFFECTIVE_BYTES_PER_SECOND` is less than the raw capacity of the hardware, then the tape is not streaming. If `EFFECTIVE_BYTES_PER_SECOND` is greater than the raw capacity of the hardware, the tape may or may not be streaming. Compression may cause the `EFFECTIVE_BYTES_PER_SECOND` to be greater than the speed of real I/O.

Identifying Bottlenecks with Synchronous I/O

With synchronous I/O, it is difficult to identify specific bottlenecks because all synchronous I/O is a bottleneck to the process. The only way to tune synchronous I/O is to compare the rate (in bytes/second) with the device's maximum throughput rate. If the rate is lower than the rate that the device specifies, then consider tuning this aspect of the backup and restore process. The `DISCRETE_BYTES_PER_SECOND` column in the `V$BACKUP_SYNC_IO` view displays the I/O rate. If you see data in `V$BACKUP_SYNC_IO`, then the problem is that you have not enabled asynchronous I/O or you are not using disk I/O slaves.

Identifying Bottlenecks with Asynchronous I/O

Long waits are the number of times the backup or restore process told the operating system to wait until an I/O was complete. **Short waits** are the number of times the backup or restore process made an operating system call to poll for I/O completion in a nonblocking mode. **Ready** indicates the number of time when I/O was already ready for use and so there was no need to made an operating system call to poll for I/O completion.

The simplest way to identify the bottleneck is to query `V$BACKUP_ASYNC_IO` for the datafile that has the largest ratio for `LONG_WAITS` divided by `IO_COUNT`.

Note: If you have synchronous I/O but you have set `BACKUP_DISK_IO_SLAVES`, then the I/O will be displayed in `V$BACKUP_ASYNC_IO`.

See Also: *Oracle Database Reference* for descriptions of the `V$BACKUP_SYNC_IO` and `V$BACKUP_ASYNC_IO` views

Instance Recovery Performance Tuning: Fast-Start Fault Recovery

This section describes instance recovery, and how Oracle's Fast-Start Fault Recovery improves availability in the event of a crash or instance failure. It also offers guidelines for tuning the time required to perform crash and instance recovery.

This section contains the following topics:

- [Understanding Instance Recovery](#)
- [Checkpointing and Cache Recovery](#)
- [Configuring the Duration of Cache Recovery: FAST_START_MTTR_TARGET](#)
- [Tuning FAST_START_MTTR_TARGET and Using MTTR Advisor](#)

Understanding Instance Recovery

Instance and crash recovery are the automatic application of redo log records to Oracle data blocks after a crash or system failure. During normal operation, if an instance is shut down cleanly (as when using a `SHUTDOWN IMMEDIATE` statement), rather than terminated abnormally, then the in-memory changes that have not already been written to the datafiles on disk are written to disk as part of the checkpoint performed during shutdown.

However, if a single instance database crashes or if all instances of an Oracle Real Application Cluster configuration crash, then Oracle performs crash recovery at the next startup. If one or more instances of an Oracle Real Application Cluster configuration crash, then a surviving instance performs instance recovery automatically. Instance and crash recovery occur in two steps: cache recovery followed by transaction recovery.

The database can be opened as soon as cache recovery completes, so improving the performance of cache recovery is important for increasing availability.

Cache Recovery (Rolling Forward)

During the cache recovery step, Oracle applies all committed and uncommitted changes in the redo log files to the affected data blocks. The work required for cache recovery processing is proportional to the rate of change to the database (update transactions each second) and the time between checkpoints.

Transaction Recovery (Rolling Back)

To make the database consistent, the changes that were not committed at the time of the crash must be undone (in other words, rolled back). During the transaction recovery step, Oracle applies the rollback segments to undo the uncommitted changes.

Checkpointing and Cache Recovery

Periodically, Oracle records a checkpoint. A checkpoint is the highest system change number (SCN) such that all data blocks less than or equal to that SCN are known to be written out to the data files. If a failure occurs, then only the redo records containing changes at SCNs higher than the checkpoint need to be applied during recovery. The duration of cache recovery processing is determined by two factors: the number of data blocks that have changes at SCNs higher than the SCN of the checkpoint, and the number of log blocks that need to be read to find those changes.

How Checkpoints Affect Performance

Frequent checkpointing writes dirty buffers to the datafiles more often than otherwise, and so reduces cache recovery time in the event of an instance failure. If checkpointing is frequent, then applying the redo records in the redo log between the current checkpoint position and the end of the log involves processing relatively few data blocks. This means that the cache recovery phase of recovery is fairly short.

However, in a high-update system, frequent checkpointing can reduce runtime performance, because checkpointing causes DBWn processes to perform writes.

Fast Cache Recovery Trade-offs To minimize the duration of cache recovery, you must force Oracle to checkpoint often, thus keeping the number of redo log records to be applied during recovery to a minimum. However, in a high-update system, frequent checkpointing increases the overhead for normal database operations.

If daily operational efficiency is more important than minimizing recovery time, then decrease the frequency of writes to data files due to checkpoints. This should improve operational efficiency, but also increase cache recovery time.

Configuring the Duration of Cache Recovery: FAST_START_MTTR_TARGET

The Fast-Start Fault Recovery feature reduces the time required for cache recovery, and makes the recovery bounded and predictable by limiting the number of dirty buffers and the number of redo records generated between the most recent redo record and the last checkpoint.

The foundation of Fast-Start Fault Recovery is the Fast-Start checkpointing architecture. Instead of conventional event-driven (that is, log switching) checkpointing, which does bulk writes, fast-start checkpointing occurs incrementally. Each DBWn process periodically writes buffers to disk to advance the checkpoint position. The oldest modified blocks are written first to ensure that every write lets the checkpoint advance. Fast-Start checkpointing eliminates bulk writes and the resultant I/O spikes that occur with conventional checkpointing.

With the Fast-Start Fault Recovery feature, the `FAST_START_MTTR_TARGET` initialization parameter simplifies the configuration of recovery time from instance or system failure. `FAST_START_MTTR_TARGET` specifies a target for the expected mean time to recover (MTTR), that is, the time (in seconds) that it should take to start up the instance and perform cache recovery. Once `FAST_START_MTTR_TARGET` is set, the database manages incremental checkpoint writes in an attempt to meet that target. If you have chosen a practical value for `FAST_START_MTTR_TARGET`, you can expect your database to recover, on average, in approximately the number of seconds you have chosen.

Note: You must disable or remove the `FAST_START_IO_TARGET`, `LOG_CHECKPOINT_INTERVAL`, and `LOG_CHECKPOINT_TIMEOUT` initialization parameters when using `FAST_START_MTTR_TARGET`. Setting these parameters interferes with the mechanisms used to manage cache recovery time to meet `FAST_START_MTTR_TARGET`.

Practical Values for FAST_START_MTTR_TARGET

The maximum value for `FAST_START_MTTR_TARGET` is 3600 seconds (one hour). If you set the value to more than 3600, then Oracle rounds it to 3600.

The following example shows how to set the value of `FAST_START_MTTR_TARGET`:

```
SQL> ALTER SYSTEM SET FAST_START_MTTR_TARGET=30;
```

In principle, the minimum value for `FAST_START_MTTR_TARGET` is one second. However, the fact that you can set `FAST_START_MTTR_TARGET` this low does not mean that target can be achieved. There are practical limits to the minimum achievable MTTR target, due to such factors as database startup time.

The MTTR target that your database can achieve given the current value of `FAST_START_MTTR_TARGET` is called the **effective MTTR target**. You can view your current effective MTTR by viewing the `TARGET_MTTR` column of the `V$INSTANCE_RECOVERY` view.

The practical range of MTTR target values for your database is defined to be the range between the lowest effective MTTR target for your database and the longest that startup and cache recovery will take in the worst-case scenario (that is, when the whole buffer cache is dirty). A procedure for determining the range of achievable MTTR target values, one step in the process of tuning your `FAST_START_MTTR_TARGET` value, is described in "[Determine the Practical Range for FAST_START_MTTR_TARGET](#)" on page 11-13.

Note: It is usually not useful to set your `FAST_START_MTTR_TARGET` to a value outside the practical range. If your `FAST_START_MTTR_TARGET` value is shorter than the lower limit of the practical range, the effect is as if you set it to the lower limit of the practical range. In such a case, the effective MTTR target will be the best MTTR target the system can achieve, but checkpointing will be at a maximum, which can affect normal database performance. If you set `FAST_START_MTTR_TARGET` to a time longer than the practical range, the MTTR target will be no better than the worst-case situation.

Reducing Checkpoint Frequency to Optimize Runtime Performance

To reduce the checkpoint frequency and optimize runtime performance, you can do the following:

- Set the value of `FAST_START_MTTR_TARGET` to 3600. This enables Fast-Start checkpointing and the Fast-Start Fault Recovery feature, but minimizes its effect on runtime performance while avoiding the need for performance tuning of `FAST_START_MTTR_TARGET`.
- Size your online redo log files according to the amount of redo your system generates. A good rule of thumb is to switch logs at most every twenty minutes. Having your log files too small can increase checkpoint activity and reduce performance. Also note that all redo log files should be the same size.

See Also: *Oracle Database Concepts* for a complete discussion of checkpoints

Monitoring Cache Recovery with `V$INSTANCE_RECOVERY`

The `V$INSTANCE_RECOVERY` view displays the current recovery parameter settings. You can also use statistics from this view to determine which factor has the greatest influence on checkpointing.

The following table lists those columns most useful in monitoring predicted cache recovery performance:

Table 11–2 V\$INSTANCE_RECOVERY Columns

Column	Description
TARGET_MTTR	Effective mean time to recover (MTTR) target in seconds. This field is 0 if FAST_START_MTTR_TARGET is not specified.
ESTIMATED_MTTR	The current estimated mean time to recover (MTTR) in seconds, based on the current number of dirty buffers and log blocks. This field is always calculated, whether or not FAST_START_MTTR_TARGET is specified.

For more details on the columns in V\$INSTANCE_RECOVERY, see *Oracle Database Reference*.

As part of the ongoing monitoring of your database, you can periodically compare V\$INSTANCE_RECOVERY.TARGET_MTTR to your FAST_START_MTTR_TARGET. The two values should generally be the same if the FAST_START_MTTR_TARGET value is in the practical range. If TARGET_MTTR is consistently longer than FAST_START_MTTR_TARGET, then set FAST_START_MTTR_TARGET to a value no less than TARGET_MTTR. If TARGET_MTTR is consistently shorter, then set FAST_START_MTTR_TARGET to a value no greater than TARGET_MTTR.

Tuning FAST_START_MTTR_TARGET and Using MTTR Advisor

To determine the appropriate value for FAST_START_MTTR_TARGET for your database, use the following four step process:

- [Calibrate the FAST_START_MTTR_TARGET](#)
- [Determine the Practical Range for FAST_START_MTTR_TARGET](#)
- [Evaluate Different Target Values with MTTR Advisor](#)
- [Determine Optimal Size for Redo Logs](#)

Calibrate the FAST_START_MTTR_TARGET

The FAST_START_MTTR_TARGET initialization parameter causes the database to calculate internal system trigger values, in order to limit the length of the redo log and the number of dirty data buffers in the data cache. This calculation uses estimated time to read a redo block, estimates of the time to read and write a data block, as well as characteristics of typical workload of the system, such as how many dirty buffers corresponds to how many change vectors, and so on.

Initially, internal defaults are used in the calculation. These defaults are replaced over time by data gathered on I/O performance during system operation and actual cache recoveries.

You will have to perform several instance recoveries in order to calibrate your FAST_START_MTTR_TARGET value properly. Before starting calibration, you must decide whether FAST_START_MTTR_TARGET is being calibrated for a database crash or a hardware crash. This is a consideration if your database files are stored in a file system or if your I/O subsystem has a memory cache, because there is a considerable difference in the read and write time to disk depending on whether or not the files are cached. The appropriate value for FAST_START_MTTR_TARGET will depend upon which type of crash is more important to recover from quickly.

To effectively calibrate FAST_START_MTTR_TARGET, make sure that you run the typical workload of the system for long enough, and perform several instance

recoveries to ensure that the time to read a redo block and the time to read or write a data block during recovery are recorded accurately.

Determine the Practical Range for FAST_START_MTTR_TARGET

After calibration, you can perform tests to determine the practical range for FAST_START_MTTR_TARGET for your database.

Determining Lower Bound for FAST_START_MTTR_TARGET: Scenario To determine the lower bound of the practical range, set FAST_START_MTTR_TARGET to 1, and start up your database. Then check the value of V\$INSTANCE_RECOVERY.TARGET_MTTR, and use this value as a good lower bound for FAST_START_MTTR_TARGET. Database startup time, rather than cache recovery time, is usually the dominant factor in determining this limit.

For example, set the FAST_START_MTTR_TARGET to 1:

```
SQL> ALTER SYSTEM SET FAST_START_MTTR_TARGET=1;
```

Then, execute the following query immediately after opening the database:

```
SQL> SELECT TARGET_MTTR, ESTIMATED_MTTR
       FROM V$INSTANCE_RECOVERY;
```

Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR
18          15
```

The TARGET_MTTR value of 18 seconds is the minimum MTTR target that the system can achieve, that is, the lowest practical value for FAST_START_MTTR_TARGET. This minimum is calculated based on the average database startup time.

The ESTIMATED_MTTR field contains the estimated mean time to recovery based on the current state of the running database. Because the database has just opened, the system contains few dirty buffers, so not much cache recovery would be required if the instance failed at this moment. That is why ESTIMATED_MTTR can, for the moment, be lower than the minimum possible TARGET_MTTR.

ESTIMATED_MTTR can be affected in the short term by recent database activity. Assume that you query V\$INSTANCE_RECOVERY immediately after a period of heavy update activity in the database. Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR
18          30
```

Now the effective MTTR target is still 18 seconds, and the estimated MTTR (if a crash happened at that moment) is 30 seconds. This is an acceptable result. This means that some checkpoints writes might not have finished yet, so the buffer cache contains more dirty buffers than targeted.

Now wait for sixty seconds and reissue the query to V\$INSTANCE_RECOVERY. Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR
18          25
```

The estimated MTTR at this time has dropped to 25 seconds, because some of the dirty buffers have been written out during this period

Determining Upper Bound for FAST_START_MTTR_TARGET To determine the upper bound of the practical range, set FAST_START_MTTR_TARGET to 3600, and operate your database under a typical workload for a while. Then check the value of V\$INSTANCE_RECOVERY.TARGET_MTTR. This value is a good upper bound for FAST_START_MTTR_TARGET.

The procedure is substantially similar to that in "[Determining Lower Bound for FAST_START_MTTR_TARGET: Scenario](#)" on page 11-13.

Selecting Preliminary Value for FAST_START_MTTR_TARGET Once you have determined the practical bounds for the FAST_START_MTTR_TARGET parameter, select a preliminary value for the parameter. Choose a higher value within the practical range if your concern is with database performance, and a lower value within the practical range if your priority is shorter recovery times. The narrower the practical range, of course, the easier the choice becomes.

For example, if you discovered that the practical range was between 17 and 19 seconds, it would be quite simple to choose 19, because it makes relatively little difference in recovery time and at the same time minimizes the effect of checkpointing on system performance. However, if you found that the practical range was between 18 and 40 seconds, you might choose a compromise value of 30, and set the parameter accordingly:

```
SQL> ALTER SYSTEM SET FAST_START_MTTR_TARGET=30;
```

You might then go on to use the MTTR Advisor to determine an optimal value.

Evaluate Different Target Values with MTTR Advisor

Once you have selected a preliminary value for FAST_START_MTTR_TARGET, you can use MTTR Advisor to evaluate the effect of different FAST_START_MTTR_TARGET settings on system performance, compared to your chosen setting.

Enabling MTTR Advisor To enable MTTR Advisor, set the two initialization parameters STATISTICS_LEVEL and FAST_START_MTTR_TARGET.

STATISTICS_LEVEL governs whether all advisors are enabled and is not specific to MTTR Advisor. Make sure that it is set to TYPICAL or ALL. Then, when FAST_START_MTTR_TARGET is set to a non-zero value, the MTTR Advisor is enabled.

Using MTTR Advisor After enabling MTTR Advisor, run a typical database workload for a while. When MTTR Advisor is ON, the database simulates checkpoint queue behavior under the current value of FAST_START_MTTR_TARGET, and up to four other different MTTR settings within the range of valid FAST_START_MTTR_TARGET values. (The database will in this case determine the valid range for FAST_START_MTTR_TARGET itself before testing different values in the range.)

Viewing MTTR Advisor Results: V\$MTTR_TARGET_ADVICE The dynamic performance view V\$MTTR_TARGET_ADVICE lets you view statistics or advisories collected by MTTR Advisor.

The database populates V\$MTTR_TARGET_ADVICE with advice about the effects of each of the FAST_START_MTTR_TARGET settings for your database. For each possible value of FAST_START_MTTR_TARGET, the row contains details about how many cache writes would be performed under the workload tested for that value of FAST_START_MTTR_TARGET.

Specifically, each row contains information about cache writes, total physical writes (including direct writes), and total I/O (including reads) for that value of FAST_

`START_MTTR_TARGET`, expressed both as a total number of operations and a ratio compared to the operations under your chosen `FAST_START_MTTR_TARGET` value. For instance, a ratio of 1.2 indicates 20% more cache writes.

Knowing the effect of different `FAST_START_MTTR_TARGET` settings on cache write activity and other I/O enables you to decide better which `FAST_START_MTTR_TARGET` value best fits your recovery and performance needs.

If MTTR Advisor is currently on, then `V$MTTR_TARGET_ADVICE` shows the Advisor information collected. If MTTR Advisor is currently OFF, the view shows information collected the last time MTTR Advisor was ON since database startup, if any. If the database has been restarted since the last time the MTTR Advisor was used, or if it has never been used, the view will not show any rows.

See Also: *Oracle Database Reference* for column details of the `V$MTTR_TARGET_ADVICE` view

Determine Optimal Size for Redo Logs

You can use the `V$INSTANCE_RECOVERY` view column `OPTIMAL_LOGFILE_SIZE` to determine the size of your online redo logs. This field shows the redo log file size in megabytes that is considered optimal based on the current setting of `FAST_START_MTTR_TARGET`. If this field consistently shows a value greater than the size of your smallest online log, then you should configure all your online logs to be at least this size.

Note, however, that the redo log file size affects the MTTR. In some cases, you may be able to refine your choice of the optimal `FAST_START_MTTR_TARGET` value by re-running the MTTR Advisor with your suggested optimal log file size.

Recovery Manager Troubleshooting

This chapter describes how to troubleshoot Recovery Manager. This chapter contains these topics:

- [Interpreting RMAN Message Output](#)
- [Testing the Media Management API](#)
- [Terminating an RMAN Command](#)
- [RMAN Troubleshooting Scenarios](#)

Interpreting RMAN Message Output

Recovery Manager provides detailed error messages that can aid in troubleshooting problems. Also, the Oracle database server and third-party media vendors generate useful debugging output of their own. The discussion which follows explains how to identify and interpret the different errors you may encounter.

Identifying Types of Message Output

Output that is useful for troubleshooting failed or hung RMAN jobs is located in several different places, as explained in the following table.

Type of Output	Produced By	Location	Description
RMAN messages	RMAN	<p>Completed job information is in <code>V\$RMAN_STATUS</code> and <code>RC_RMAN_STATUS</code>. Current job information is in <code>V\$RMAN_OUTPUT</code>.</p> <p>When running RMAN from the command line, you can direct output to the following places:</p> <ul style="list-style-type: none"> ■ Standard output ■ A log file specified by <code>LOG</code> on the command line or the <code>SPOOL LOG</code> command ■ A file created by redirecting RMAN output (for example, <code>UNIX > operator</code>) 	<p>Contains actions relevant to the RMAN job as well as error messages generated by RMAN, the database server, and the media vendor. RMAN error messages have an <code>RMAN-xxxxx</code> prefix. Normal action descriptions do not have a prefix.</p>

Type of Output	Produced By	Location	Description
alert_SID.log	Oracle database server	The directory named in the BACKGROUND_DUMP_DEST initialization parameter.	Contains a chronological log of errors, initialization parameter settings, and administration operations. Records values for overwritten control file records (refer to <i>Oracle Data Guard Concepts and Administration</i>).
Oracle trace file	Oracle database server	The directory specified in the USER_DUMP_DEST initialization parameter.	Contains detailed output generated by Oracle server processes. This file is created when an ORA-600 or ORA-3113 error message occurs, whenever RMAN cannot allocate a channel, and when the database fails to load the media management library.
sbtio.log	Third-party media management software	The directory specified in the USER_DUMP_DEST initialization parameter.	Contains vendor-specific information written by the media management software. This log does not contain Oracle server or RMAN errors.
Media manager log file	Third-party media management software	The filenames for any media manager logs other than sbtio.log are determined by the media management software.	Contains information on the functioning of the media management device.

Recognizing RMAN Error Message Stacks

RMAN reports errors as they occur. If an error is not retrievable, that is, RMAN cannot perform failover to another channel to complete a particular job step, then RMAN also reports a summary of the errors after all job sets complete. This feature is known as **deferred error reporting**.

One way to determine whether RMAN encountered an error is to examine its return code, as described in ["Identifying RMAN Return Codes"](#) on page 12-7. A second way is to search the RMAN output for the string RMAN-00569, which is the message number for the error stack banner. All RMAN errors are preceded by this error message. If you do not see an RMAN-00569 message in the output, then there are no errors. Following is sample output for a syntax error:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-00558: error encountered while parsing input commands
RMAN-01005: syntax error: found "): expecting one of: "archivelog, backup,
backupset, controlfilecopy, current, database, datafile, datafilecopy, (, plus, ;,
tablespace"
RMAN-01007: at line 1 column 18 file: standard input
```

Identifying Error Codes

Typically, you find the following types of error codes in RMAN message stacks:

- Errors prefixed with RMAN-
- Errors prefixed with ORA-
- Errors preceded by the line Additional information:

See Also: *Oracle Database Error Messages* for explanations of RMAN and ORA error codes

RMAN Error Message Numbers

Table 12–1 indicates the error ranges for common RMAN error messages, all of which are described in *Oracle Database Error Messages*.

Table 12–1 RMAN Error Message Ranges

Error Range	Cause
0550-0999	Command-line interpreter
1000-1999	Keyword analyzer
2000-2999	Syntax analyzer
3000-3999	Main layer
4000-4999	Services layer
5000-5499	Compilation of RESTORE or RECOVER command
5500-5999	Compilation of DUPLICATE command
6000-6999	General compilation
7000-7999	General execution
8000-8999	PL/SQL programs
9000-9999	Low-level keyword analyzer
10000-10999	Server-side execution
11000-11999	Interphase errors between PL/SQL and RMAN
12000-12999	Recovery catalog packages

ORA-19511: Media Manager Errors

In the event of a media manager error, ORA-19511 is signalled, and the media manager is expected to provide RMAN a descriptive error. RMAN will display the error passed back to it by the media manager. For example, you might see this:

```
ORA-19511: Error received from media manager layer, error text:
  sbtpvt_open_input: file .* does not exist or cannot be accessed, errno = 2
```

The message from the media manager should provide you with enough information to let you fix the root problem. If it does not, you should refer to the documentation for your media manager or contact your media management vendor support representative for further information. ORA-19511 errors originate with the media manager, not the Oracle database. The database merely passes the message on from the media manager. The cause can only be addressed by the media management vendor.

Note that if you are still using an SBT 1.1-compliant media management layer, you may see some additional error message text. Output from an SBT 1.1-compliant media management layer is similar to the following:

```
ORA-19507: failed to retrieve sequential file, handle="c-140148591-20031014-06",
parms=""
ORA-27007: failed to open file
Additional information: 7000
Additional information: 2
ORA-19511: Error received from media manager layer, error text:
```

```
SBT error = 7000, errno = 0, sbtopen: backup file not found
```

The "Additional information" provided uses error codes specific to SBT 1.1. The values displayed correspond to the media manager message numbers and error text listed in [Table 12-2](#). RMAN re-signals the error, as an ORA-19511 Error received from media manager layer error, and a general error message related to the error code returned from the media manager and including the SBT 1.1 error number is then displayed.

The SBT 1.1 error messages are listed here for your reference. [Table 12-2](#) lists media manager message numbers and their corresponding error text. In the error codes, *O/S* stands for *operating system*. The errors prefixed with an asterisk are internal and should not typically be seen during normal operation.

Table 12-2 Media Manager Error Message Ranges

Cause	No.	Message
sbtopen	7000	Backup file not found (only returned for read)
	7001	File exists (only returned for write)
	7002*	Bad mode specified
	7003	Invalid block size specified
	7004	No tape device found
	7005	Device found, but busy; try again later
	7006	Tape volume not found
	7007	Tape volume is in-use
	7008	I/O Error
	7009	Can't connect with Media Manager
	7010	Permission denied
	7011	O/S error for example malloc, fork error
	7012*	Invalid argument(s) to sbtopen
sbtclose	7020*	Invalid file handle or file not open
	7021*	Invalid flags to sbtclose
	7022	I/O error
	7023	O/S error
	7024*	Invalid argument(s) to sbtclose
	7025	Can't connect with Media Manager
sbtwrite	7040*	Invalid file handle or file not open
	7041	End of volume reached
	7042	I/O error
	7043	O/S error
	7044*	Invalid argument(s) to sbtwrite
sbtread	7060*	Invalid file handle or file not open
	7061	EOF encountered
	7062	End of volume reached
	7063	I/O error
	7064	O/S error
	7065*	Invalid argument(s) to sbtread

Table 12–2 (Cont.) Media Manager Error Message Ranges

Cause	No.	Message
sbtremove	7080	Backup file not found
	7081	Backup file in use
	7082	I/O Error
	7083	Can't connect with Media Manager
	7084	Permission denied
	7085	O/S error
	7086*	Invalid argument(s) to sbtremove
sbtinfo	7090	Backup file not found
	7091	I/O Error
	7092	Can't connect with Media Manager
	7093	Permission denied
	7094	O/S error
	7095*	Invalid argument(s) to sbtinfo
sbtinit	7110*	Invalid argument(s) to sbtinit
	7111	O/S error

Interpreting RMAN Error Stacks

Sometimes you may find it difficult to identify the useful messages in the RMAN error stack. Note the following tips and suggestions:

- Read the messages from the bottom up, because this is the order in which RMAN issues the messages. The last one or two errors displayed in the stack are often the most informative.
- When using an SBT 1.1 media management layer and presented with SBT 1.1 style error messages containing the "Additional information:" numeric error codes, look for the ORA-19511 message that follows for the text of error messages passed back to RMAN by the media manager. These should identify the real failure in the media management layer.
- Look for the RMAN-03002 or RMAN-03009 message (RMAN-03009 is the same as RMAN-03002 but includes the channel ID), immediately following the error banner. These messages indicate which command failed. Syntax errors generate RMAN-00558.
- Identify the basic type of error according to the error range chart in [Table 12–1](#) and then refer to *Oracle Database Error Messages* for information on the most important messages.

Interpreting RMAN Errors: Example

You attempt a backup of tablespace `users` and receive the following message:

```
Starting backup at 29-AUG-02
using channel ORA_DISK_1
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of backup command at 08/29/2002 15:14:03
RMAN-20202: tablespace not found in the recovery catalog
RMAN-06019: could not translate tablespace name "USESR"
```

The RMAN-03002 error indicates that the BACKUP command failed. You read the last two messages in the stack first and immediately see the problem: no tablespace users appears in the recovery catalog because you mistyped the name.

Interpreting Server Errors: Example

Assume that you attempt to recover a tablespace and receive the following errors:

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

```
Starting recover at 29-AUG-01
using channel ORA_DISK_1
```

```
starting media recovery
media recovery failed
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of recover command at 08/29/2001 15:18:43
RMAN-11003: failure during parse/execution of SQL statement: alter database
recover if needed tablespace USERS
ORA-00283: recovery session canceled due to errors
ORA-01124: cannot recover data file 8 - file is in use or recovery
ORA-01110: data file 8: '/oracle/oradata/trgt/users01.dbf'
```

As suggested, you start reading from the bottom up. The ORA-01110 message explains there was a problem with the recovery of datafile users01.dbf. The second error indicates that the database cannot recover the datafile because it is in use or already being recovered. The remaining RMAN errors indicate that the recovery session was cancelled due to the server errors. Hence, you conclude that because you were not already recovering this datafile, the problem must be that the datafile is online and you need to take it offline and restore a backup.

Interpreting SBT 2.0 Media Management Errors: Example

Assume that you use a tape drive and receive the following output during a backup job:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
ORA-19624: operation failed, retry possible
ORA-19507: failed to retrieve sequential file, handle="/tmp/foo", parms=""
ORA-27029: skgfrtrv: sbtrestore returned error
ORA-19511: Error received from media manager layer, error text:
  sbtpvt_open_input:file /tmp/foo does not exist or cannot be accessed, errno=2
```

The error text displayed following the ORA-19511 error is generated by the media manager and describes the real source of the failure. Refer to the media manager documentation to interpret this error.

Interpreting SBT 1.1 Media Management Errors: Example

Assume that you use a tape drive and receive the following output during a backup job:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03009: failure of backup command on c1 channel at 09/04/2001 13:18:19
```

```
ORA-19506: failed to create sequential file, name="07d36ecp_1_1", parms=""
ORA-27007: failed to open file
SVR4 Error: 2: No such file or directory
Additional information: 7005
Additional information: 1
ORA-19511: Error received from media manager layer, error text:
  SBT error = 7005, errno = 2, sbtopen: system error
```

The main information of interest returned by SBT 1.1 media managers is the error code in the "Additional information" line:

```
Additional information: 7005
```

Referring to [Table 12–2, "Media Manager Error Message Ranges"](#), you discover that error 7005 means that the media management device is busy. So, the media management software is not able to write to the device because it is in use or there is a problem with it.

Note: The `sbtio.log` contains information written by the media management software, *not* the Oracle database server. Hence, you must consult your media vendor documentation to interpret the error codes and messages. If no information is written to the `sbtio.log`, contact your media manager support to ask whether they are writing error messages in some other location, or whether there are steps you need to take to have the media manager errors appear in `sbtio.log`.

Identifying RMAN Return Codes

One way to determine whether RMAN encountered an error is to examine its return code or exit status. The RMAN client returns 0 to the shell from which it was invoked if no errors occurred, and a nonzero error value otherwise.

How you access this return code depends upon the environment from which you invoked the RMAN client. For example, if you are running UNIX with the C shell, then, when RMAN completes, the return code is placed in a shell variable called `$status`. The method of returning exit status is a detail specific to the host operating system rather than the RMAN client.

Testing the Media Management API

On some platforms, Oracle provides a diagnostic tool called `sbttest`. This utility performs a simple test of the media management software by attempting to communicate with the media manager as the Oracle database server would.

Obtaining the `sbttest` Utility

On UNIX, the `sbttest` utility is typically located in `$ORACLE_HOME/bin`. If for some reason the utility is not included with your platform, then contact Oracle Support to obtain the C version of the program. You can compile this version of the program on all UNIX platforms.

Note that on platforms such as Solaris, you do not have to relink when using `sbttest`. On other platforms, relinking may be necessary.

Obtaining Online Documentation for the sbttest Utility

For online documentation of `sbttest`, issue the following on the command line:

```
% sbttest
```

The program displays the list of possible arguments for the program:

```
Error: backup file name must be specified
Usage: sbttest backup_file_name # this is the only required parameter
    <-dbname database_name>
    <-trace trace_file_name>
    <-remove_before>
    <-no_remove_after>
    <-read_only>
    <-no_regular_backup_restore>
    <-no_proxy_backup>
    <-no_proxy_restore>
    <-file_type n>
    <-copy_number n>
    <-media_pool n>
    <-os_res_size n>
    <-pl_res_size n>
    <-block_size block_size>
    <-block_count block_count>
    <-proxy_file os_file_name bk_file_name
                [os_res_size pl_res_size block_size block_count]>
    <-libname sbt_library_name>
```

The display also indicates the meaning of each argument. For example, following is the description for two optional parameters:

Optional parameters:

```
-dbname specifies the database name which will be used by SBT
        to identify the backup file. The default is "sbt_db"
-trace  specifies the name of a file where the Media Management
        software will write diagnostic messages.
```

Using the sbttest Utility

Use `sbttest` to perform a quick test of the media manager.

If `sbttest` returns 0, then the test ran without error, which means that the media manager is correctly installed and can accept a data stream and return the same data when requested. If `sbttest` returns a non-zero value, then either the media manager is not installed or it is not configured correctly.

To use sbttest:

1. Make sure the program is installed and included in the system path by typing `sbttest` at the command line:

```
% sbttest
```

If the program is operational, then you should see a display of the online documentation.

2. Execute the program, specifying any of the arguments described in the online documentation. For example, enter the following to create test file `some_file.f` and write the output to `sbtio.log`:

```
% sbttest some_file.f -trace sbtio.log
```

You can also test a backup of an existing datafile. For example, this command tests datafile `tbs_33.f` of database `prod`:

```
% sbttest tbs_33.f -dbname prod
```

3. Examine the output. If the program encounters an error, then it provides messages describing the failure. For example, if the database cannot find the library, you see:

```
libobk.so could not be loaded. Check that it is installed properly, and that
LD_LIBRARY_PATH environment variable (or its equivalent on your platform)
includes the directory where this file can be found. Here is some additional
information on the cause of this error:
ld.so.1: sbttest: fatal: libobk.so: open failed: No such file or directory
```

Note that in some cases `sbttest` can work but an RMAN backup does not. The reasons can be the following:

- The user who starts `sbttest` is not the owner of the Oracle processes.
- If the database server is not linked with the media management library or cannot load it dynamically when needed, then RMAN backups to the media manager fail, but `sbttest` may still work.
- The `sbttest` program passes all environment parameters from the shell but RMAN does not.

Terminating an RMAN Command

There are several ways to terminate an RMAN command in the middle of execution:

- The preferred method is to press `CTRL+C` (or the equivalent "attention" key combination for your system) in the RMAN interface. This will also terminate allocated channels, unless they are hung in the media management code, as happens when, for example, when they are waiting for a tape to be mounted.
- You can kill the server session corresponding to the RMAN channel by running the SQL `ALTER SYSTEM KILL SESSION` statement.
- You can terminate the server session corresponding to the RMAN channel on the operating system.

Terminating the Session with ALTER SYSTEM KILL SESSION

You can identify the Oracle session ID for an RMAN channel by looking in the RMAN log for messages with the format shown in the following example:

```
channel ch1: sid=15 devtype=SBT_TAPE
```

The `sid` and `devtype` are displayed for each allocated channel. Note that the Oracle `sid` is different from the operating system process ID. You can kill the session using a SQL `ALTER SYSTEM KILL SESSION` statement.

`ALTER SYSTEM KILL SESSION` takes two arguments, the `sid` printed in the RMAN message and a serial number, both of which can be obtained by querying `V$SESSION`. For example, run the following statement, where `sid_in_rman_output` is the number from the RMAN message:

```
SELECT SERIAL# FROM V$SESSION WHERE SID=sid_in_rman_output;
```

Then, run the following statement, substituting the `sid_in_rman_output` and serial number obtained from the query:

```
ALTER SYSTEM KILL SESSION 'sid_in_rman_output,serial#';
```

Note that this will not unhang the session if the session is hung in media manager code..

Terminating the Session at the Operating System Level

Finding and killing the processes that are associated with the server sessions is operating system specific. On some platforms the server sessions are not associated with any processes at all. Refer to your operating system specific documentation for more information.

Terminating an RMAN Session That Is Hung in the Media Manager

You may sometimes need to kill an RMAN job that is hung in the media manager. The best way to terminate RMAN when the channel connections are hung in the media manager is to kill the session in the media manager. If this action does not solve the problem, then on some platforms, such as Unix, you may be able to kill the Oracle processes of the connections. (Note that killing the Oracle processes may cause problems from the media manager. See your media manager documentation for details.)

Components of an RMAN Session

The nature of an RMAN session depends on the operating system. In UNIX, an RMAN session has the following processes associated with it:

- The **RMAN client process** itself
- The **default channel**, the initial connection to the target database
- One **target connection** to the target database corresponding to each allocated channel
- The **catalog connection** to the recovery catalog database, if you use a recovery catalog
- An **auxiliary connection** to an auxiliary instance, during **DUPLICATE** or **TSPITR** operations
- A **polling connection** to the target database, used for monitoring RMAN command execution on the various allocated channels. By default, RMAN makes one polling connection. RMAN makes additional polling connections if you use different connect strings in the **ALLOCATE CHANNEL** or **CONFIGURE CHANNEL** commands. One polling connection exists for each distinct connect string used in the **ALLOCATE CHANNEL** or **CONFIGURE CHANNEL** command.

Process Behavior During a Hung Job

RMAN usually hangs because one of the channel connections is waiting in the media manager code for a tape resource. The catalog connection and the default channel appear to hang, because they are waiting for RMAN to tell them what to do. Polling connections seem to be in an infinite loop while polling the RPC under the control of the RMAN process.

If you kill the RMAN process itself, then you also kill the catalog connection, the auxiliary connection, the default channel, and the polling connections. If target and auxiliary connections are not hung in the media manager code, they also terminate. If either the target connection or any of the auxiliary connections are executing in the

media management layer, they will not terminate until the processes are manually killed at the operating system level.

Not all media managers can detect the termination of the Oracle process. Those which cannot may keep resources busy or continue processing. Consult your media manager documentation for details.

Terminating the catalog connection does not cause the RMAN process to terminate because RMAN is not performing catalog operations while the backup or restore is in progress. Removing default channel and polling connections causes the RMAN process to detect that one of the channels has died and then proceed to exit. In this case, the connections to the hung channels remain active as described previously.

Terminating an RMAN Session: Basic Steps

Once the hung channels in the media manager code are killed, the RMAN process detects this termination and proceed to exit, removing all connections except target connections that are still operative in the media management layer. The warning about the media manager resources still applies in this case.

To terminate an Oracle process that is hung in the media manager:

1. Query V\$SESSION and V\$SESSION_WAIT as described in "[Monitoring RMAN Through V\\$ Views](#)" on page 9-10. For example, execute the following query:

```
COLUMN EVENT FORMAT a10
COLUMN SECONDS_IN_WAIT FORMAT 999
COLUMN STATE FORMAT a20
COLUMN CLIENT_INFO FORMAT a30

SELECT p.SPID, EVENT, SECONDS_IN_WAIT AS SEC_WAIT,
       sw.STATE, CLIENT_INFO
FROM V$SESSION_WAIT sw, V$SESSION s, V$PROCESS p
WHERE sw.EVENT LIKE 'sbt%'
      AND s.SID=sw.SID
      AND s.PADDR=p.ADDR
;
```

Examine the SQL output to determine which sbt functions are waiting. For example, the output may be as follows:

SPID	EVENT	SEC_WAIT	STATE	CLIENT_INFO
8642	sbtwrite2	600	WAITING	rman channel=ORA_SBT_TAPE_1
8374	sbtwrite2	600	WAITING	rman channel=ORA_SBT_TAPE_2

2. Using operating system-level tools appropriate to your platform, kill the hung sessions. For example, on Solaris execute a `kill -9` command:

```
% kill -9 8642 8374
```

On Windows, there is a command-line utility called ORAKILL which lets you kill a specific thread in this situation. From a command prompt, run the following command:

```
orakill sid thread_id
```

where *sid* identifies the database instance to target, and the *thread_id* is the SPID value from the query in step 1.

3. Check that the media manager also clears its processes. If any remain, the next backup or restore operation may hang again, due to the previous hang. In some

media managers, the only solution is to shut down and restart the media manager. If the documentation from the media manager does not provide the needed information, contact technical support for the media manager.

See Also: Your operating system specific documentation for the relevant commands

RMAN Troubleshooting Scenarios

This section contains these topics:

- [After Installation of Media Manager, RMAN Channel Allocation Fails: Scenario](#)
- [Backup Job Is Hanging: Scenario](#)
- [RMAN Fails to Start RPC Call: Scenario](#)
- [Backup Fails with Invalid RECID Error: Scenario](#)
- [Backup Fails Because of Control File Enqueue: Scenario](#)
- [RMAN Fails to Delete All Archived Logs: Scenario](#)
- [Backup Fails Because RMAN Cannot Locate an Archived Log: Scenario](#)
- [RMAN Does Not Recognize Character Set Name: Scenario](#)
- [RMAN Denies Logon to Target Database: Scenario](#)
- [Database Duplication Fails Because of Missing Log: Scenario](#)
- [Duplication Fails with Multiple RMAN-06023 Errors: Scenario](#)
- [UNKNOWN Database Name Appears in Recovery Catalog: Scenario](#)

After Installation of Media Manager, RMAN Channel Allocation Fails: Scenario

In this scenario, you install and test the media manager as explained in "[Configuring RMAN to Make Backups to a Media Manager](#)" on page 5-4, but you still cannot make RMAN back up to tape. For example, after allocating the `sbt` channel, you receive an error stack similar to the following:

```

RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03009: failure of allocate command on c1 channel at 08/29/2001 17:16:54
ORA-19554: error allocating device, device type: SBT_TAPE, device name:
ORA-27211: Failed to load Media Management Library
Additional information: 25
    
```

The most important line of the error output is the ORA-27211 error. It indicates the basic problem, that the media management library could not be loaded. Typically, there is no need to refer to the trace file or `sbtio.log` in such a case.

After Installation of Media Manager, RMAN Channel Allocation Fails: Diagnosis

The ORA-27211 error indicates that the channel allocation is failing because the database is not loading the media management library. If the channel allocation fails, then the database generates a trace file in the `USER_DUMP_DEST` location that contains the error that caused the channel allocation to fail. The trace file should have the complete path name of the media management library loaded by the database as well as any other media manager errors or operating system errors. For example, the trace

file on UNIX may be called something like `/oracle/rdbms/log/prod1_oracle_16226.trc`, and may contain information such as the following:

```
*** 2001-08-29 17:16:54.385
SKGFQ OSD: Error in function sbtinit on line 2396
SKGFQ OSD: Look for SBT Trace messages in file /oracle/rdbms/log/sbtio.log
SBT Initialize failed for oracle.static
```

The last line of this output indicates that Oracle is loading the default static library instead of the media management library that you installed.

You may find more detailed information in the file `sbtio.log`, as described in the error message. Note, however, that writing SBT trace messages is the responsibility of the media management software, not the Oracle database or RMAN. The media management vendor may not have implemented the writing of trace messages in a particular situation. Contact the media management vendor for details about the trace messages written to `sbtio.log`.

To test the loading of the media management library, try allocating a channel by using the `PARMS` parameter `SBT_LIBRARY` to force the loading of the media management library. For example, if your library is called `/vendor/lib/some_mmm_lib.so`, then run a command such as the following, making sure to specify whatever `PARMS` settings are required by your media manager:

```
RUN
{
  ALLOCATE CHANNEL c1 DEVICE TYPE sbt
  PARMS='SBT_LIBRARY=/vendor/lib/some_mmm_lib.so',
  'ENV=(NSR_SERVER=tape_svr,NSR_CLIENT=oracleclnt,NSR_GROUP=oracle_tapes)';
}
```

If the channel allocation fails, then check the trace file again to see whether you can learn anything new. If the channel allocation with `SBT_LIBRARY` succeeds, but an ordinary `sbt` channel allocation fails, then the database is probably trying to load a library other than the one you installed. By default, the database expects to find the media management library at `$ORACLE_HOME/lib/libobk.so` on UNIX, or `%ORACLE_HOME%/bin/orasbt.dll` on NT. You may have more than one library in the operating system path, and the database is loading the wrong one.

After Installation of Media Manager, RMAN Channel Allocation Fails: Solution

If the problem is that the database is not loading the correct library, then make sure that the library is named correctly in the `SBT_LIBRARY` parameter.

See Also: *Oracle Database Backup and Recovery Reference* for descriptions of the legal `PARMS` parameters

Backup Job Is Hanging: Scenario

In this scenario, an RMAN backup job starts as normal and then pauses inexplicably:

```
Recovery Manager: Release 10.1.0.2.0 - Production
```

```
Copyright (c) 1995, 2003, Oracle. All rights reserved.
```

```
connected to target database: TRGT
connected to recovery catalog database
```

```
RMAN> BACKUP TABLESPACE SYSTEM, tools;
```

```

allocated channel: t1
channel t1: sid=16 devtype=SBT_TAPE

channel t1: starting datafile backupset
set_count=15 set_stamp=338309600
channel t1: including datafile 2 in backupset
channel t1: including datafile 1 in backupset
channel t1: including current control file in backupset
# Hanging here for 30 minutes now

```

Backup Job Is Hanging: Diagnosis

If a backup job is hanging, that is, not proceeding, then several scenarios are possible:

- A server-side or media management error occurred.
- RMAN is waiting for an event such as the insertion of a new cassette into the tape device.

Query sbt wait events to gain more information. For example, run the following query on the target instance:

```

COLUMN EVENT FORMAT a10
COLUMN SECONDS_IN_WAIT FORMAT 999
COLUMN STATE FORMAT a20
COLUMN CLIENT_INFO FORMAT a30

SELECT p.SPID, EVENT, SECONDS_IN_WAIT AS SEC_WAIT,
       sw.STATE, CLIENT_INFO
FROM V$SESSION_WAIT sw, V$SESSION s, V$PROCESS p
WHERE sw.EVENT LIKE 'sbt%'
      AND s.SID=sw.SID
      AND s.PADDR=p.ADDR
;

```

Examine the SQL output to determine which sbt functions are waiting. For example, the output may be as follows:

SPID	EVENT	SEC_WAIT	STATE	CLIENT_INFO
8642	sbtbackup	1500	WAITING	rman channel=ORA_SBT_TAPE_1

Backup Job Is Hanging: Solution

Because the causes of a hung backup job can be varied, so are the solutions. For example, backup jobs often hang simply because the tape device has completely filled the current cassette and is waiting for a new tape to be inserted. Ideally, the query of the sbt wait events should indicate the problem.

In this example, a single sbtbackup has taken 1500 seconds, so RMAN is waiting on the media manager to finish its write operation. Check that the media manager is functioning normally, and contact the media management vendor's technical support for assistance.

If the sbt wait event query is unhelpful, then examine media manager process, log, and trace files for signs of abnormal termination or other errors (refer to the description of message files in ["Identifying Types of Message Output"](#) on page 12-1).

See Also: ["Terminating an RMAN Session: Basic Steps"](#) on page 12-11 to learn how to kill an RMAN session that is hanging

RMAN Fails to Start RPC Call: Scenario

In this scenario, you run a backup job and receive message output similar to the following:

```
channel c8: including datafile number 47 in backupset
RPC call appears to have failed to start on channel c9
RPC call ok on channel c9
channel c3: including datafile number 18 in backupset
```

RMAN Fails to Start RPC Call: Diagnosis

The RPC call appears to have failed message does not usually indicate a problem. The message indicates one of the following:

- The target database instance is slow.
- A timing problem occurred.

Timing problems occur in this way. When RMAN begins an RPC, it checks the `V$SESSION` performance view. The RPC updates the information in the view to indicate when it starts and finishes. Sometimes RMAN checks `V$SESSION` before the RPC has indicated it has started, which in turn generates the following message:

```
RPC call appears to have failed
```

If a message stating "RPC call ok" does not appear in the output immediately following the message stating "RPC call appears to have failed", then the backup job encountered an internal problem. Contact Oracle Support for further assistance.

Backup Fails with Invalid RECID Error: Scenario

In this scenario, you attempt a backup and receive the following error messages:

```
RMAN-3014: Implicit resync of recovery catalog failed
RMAN-6038: Recovery catalog package detected an error
RMAN-20035: Invalid high RECID error
```

Backup Fails with Invalid RECID Error: Diagnosis

In one common scenario, you restore a backup control file created through a non-Oracle mechanism, and then open the database without the `RESETLOGS` option. If you had created the backup control file through the `RMAN BACKUP` command or the `SQL ALTER DATABASE BACKUP CONTROLFILE` statement, then the database would have required you to reset the online logs.

The control file and the recovery catalog are now not synchronized. The database control file is older than the recovery catalog, because at one time the recovery catalog resynchronized with the old current control file, and now the database is using a backup control file. RMAN detects that the control file currently in use is older than the control file previously used to resynchronize.

Another common scenario occurs when you attempt to copy the target database to a new machine as follows:

1. On machine 1, you shut down the database and make a copy of the control file with an operating system utility. You do not use `CATALOG` to add this control file copy to the repository.
2. You transfer the control file copy to machine 2.

3. On machine 2, you create a new initialization parameter file and new database instance.
4. You mount the control file copy on machine 2. The database does not recognize the control file as a backup control file: to the database it looks like the current control file.
5. You start RMAN and connect to the new target database and the recovery catalog on machine 2. Because the control file was not created with RMAN and was not cataloged as a control file copy, RMAN sees the database on machine 2 as the database on machine 1.
6. You restore and recover database the new database on machine 2 and then open it. As a consequence, various records are added to the recovery catalog during the restore and recovery. For example, the highest RECID in the recovery catalog moves from 90 to 100.
7. On machine 1, you start RMAN and connect to the original target database and recovery catalog. The recovery catalog indicates that the highest RECID is 100, but the control file indicates that the highest RECID is 90. The control file RECID should always be greater than or equal to the recovery catalog RECID, so RMAN issues RMAN-20035.

Backup Fails with Invalid RECID Error: Solution 1

This solution is safest and is strongly recommended. It preserves the control file, so that the historical information about the database stored in the control file continues to be available after the procedure.

To reset the database with RMAN:

1. Connect to the target database with SQL*Plus. For example, enter:

```
% sqlplus '/ AS SYSDBA'
```
2. Mount the database if it is not already mounted. For example, enter:

```
ALTER DATABASE MOUNT;
```
3. Start cancel-based recovery by using the backup control file, then cancel it. The reason for canceling is that the USING BACKUP CONTROLFILE clause stamps the control file as a backup, which then permits OPEN RESETLOGS. For example, enter:

```
ALTER DATABASE RECOVER DATABASE UNTIL CANCEL USING BACKUP CONTROLFILE;  
ALTER DATABASE RECOVER CANCEL;
```
4. Use RMAN to connect to the target database and recovery catalog. For example, enter:

```
% rman TARGET SYS/oracle@trgt CATALOG rman/cat@catdb
```
5. Open the database with the RESETLOGS option. For example, enter:

```
RMAN> ALTER DATABASE OPEN RESETLOGS;
```
6. Take new backups so that you can recover the database if necessary. For example, enter:

```
BACKUP DATABASE PLUS ARCHIVELOG;
```

Backup Fails with Invalid RECID Error: Solution 2

This solution is similar to the previous one, but does require that you re-create your control file. It is better-suited for the case in which you are copying your database to a second system, where you may not want to keep the history from the control file for the copy of the database on the second system, or where you might drop a few datafiles or change the online logs by editing your control file.

To create the control file with SQL*Plus:

1. Connect to the target database with SQL*Plus. For example, enter:

```
% sqlplus 'SYS/oracle@trgt AS SYSDBA'
```

2. Mount the database if it is not already mounted:

```
SQL> ALTER DATABASE MOUNT;
```

3. Back up the control file to a trace file:

```
SQL> ALTER DATABASE BACKUP CONTROLFILE TO TRACE;
```

4. Edit the trace file as necessary. The relevant section of the trace file looks something like the following:

```
# The following commands will create a new control file and use it
# to open the database.
# Data used by the recovery manager will be lost. Additional logs may
# be required for media recovery of offline data files. Use this
# only if the current version of all online logs are available.
STARTUP NOMOUNT
CREATE CONTROLFILE REUSE DATABASE "TRGT" NORESETLOGS ARCHIVELOG
-- STANDBY DATABASE CLUSTER CONSISTENT AND UNPROTECTED
MAXLOGFILES 32
MAXLOGMEMBERS 2
MAXDATAFILES 32
MAXINSTANCES 1
MAXLOGHISTORY 226
LOGFILE
GROUP 1 '/oracle/oradata/trgt/redo01.log' SIZE 25M,
GROUP 2 '/oracle/oradata/trgt/redo02.log' SIZE 25M,
GROUP 3 '/oracle/oradata/trgt/redo03.log' SIZE 500K
-- STANDBY LOGFILE
DATAFILE
'/oracle/oradata/trgt/system01.dbf',
'/oracle/oradata/trgt/undotbs01.dbf',
'/oracle/oradata/trgt/cwmlite01.dbf',
'/oracle/oradata/trgt/drsys01.dbf',
'/oracle/oradata/trgt/example01.dbf',
'/oracle/oradata/trgt/indx01.dbf',
'/oracle/oradata/trgt/tools01.dbf',
'/oracle/oradata/trgt/users01.dbf'
CHARACTER SET WE8DEC
;
# Take files offline to match current control file.
ALTER DATABASE DATAFILE '/oracle/oradata/trgt/tools01.dbf' OFFLINE;
ALTER DATABASE DATAFILE '/oracle/oradata/trgt/users01.dbf' OFFLINE;
# Configure RMAN configuration record 1
VARIABLE RECNO NUMBER;
EXECUTE :RECNO := SYS.DBMS_BACKUP_RESTORE.SETCONFIG('CHANNEL', 'DEVICE TYPE DISK
DEBUG 255');
# Recovery is required if any of the datafiles are restored backups,
# or if the last shutdown was not normal or immediate.
```

```

RECOVER DATABASE
# All logs need archiving and a log switch is needed.
ALTER SYSTEM ARCHIVE LOG ALL;
# Database can now be opened normally.
ALTER DATABASE OPEN;
# Commands to add tempfiles to temporary tablespaces.
# Online tempfiles have complete space information.
# Other tempfiles may require adjustment.
ALTER TABLESPACE TEMP ADD TEMPFILE '/oracle/oradata/trgt/temp01.dbf' REUSE;
# End of tempfile additions.

```

5. Shut down the database:

```
SHUTDOWN IMMEDIATE
```

6. Execute the script to create the control file, recover (if necessary), archive the logs, and open the database:

```

STARTUP NOMOUNT
CREATE CONTROLFILE ...;
EXECUTE ...;
RECOVER DATABASE
ALTER SYSTEM ARCHIVE LOG CURRENT;
ALTER DATABASE OPEN ...;

```

7. If you intend to keep and continue using this copy of the database, use the DBNEWID utility to change the name and DBID of the new database as needed.

Caution: If you do *not* open with the RESETLOGS option, then two copies of an archived redo log for a given log sequence number may exist—even though these two copies have completely different contents. For example, one log may have been created on the original host and the other on the new host. If you accidentally confuse the logs during a media recovery, then the database will be corrupted but Oracle and RMAN cannot detect the problem.

Backup Fails Because of Control File Enqueue: Scenario

In this scenario, a backup job fails because RMAN cannot make a snapshot control file. The message stack is as follows:

```

RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of backup command at 08/30/2001 22:48:44
ORA-00230: operation disallowed: snapshot control file enqueue unavailable

```

Backup Fails Because of Control File Enqueue: Diagnosis

When RMAN needs to back up or resynchronize from the control file, it first creates a **snapshot** or consistent image of the control file. If one RMAN job is already backing up the control file while another needs to create a new snapshot control file, then you may see the following message:

```
waiting for snapshot control file enqueue
```

Under normal circumstances, a job that must wait for the control file enqueue waits for a brief interval and then successfully obtains the enqueue. RMAN makes up to five attempts to get the enqueue and then fails the job. The conflict is usually caused when

two jobs are both backing up the control file, and the job that first starts backing up the control file waits for service from the media manager.

To determine which job is holding the conflicting enqueue:

1. After you see the first message stating "RMAN-08512: waiting for snapshot control file enqueue", start a new SQL*Plus session on the target database:

```
% sqlplus 'SYS/oracle@trgt AS SYSDBA'
```

2. Execute the following query to determine which job is causing the wait:

```
SELECT s.SID, USERNAME AS "User", PROGRAM, MODULE,
       ACTION, LOGON_TIME "Logon", l.*
FROM V$SESSION s, V$ENQUEUE_LOCK l
WHERE l.SID = s.SID
AND l.TYPE = 'CF'
AND l.ID1 = 0
AND l.ID2 = 2;
```

You should see output similar to the following (the output in this example has been truncated):

SID	User	Program	Module	Action	Logon
9	SYS	rman@h13 (TNS V1-V3)	backup full datafile: c10000210	STARTED	21-JUN-01

Backup Fails Because of Control File Enqueue: Solution

Commonly, enqueue situations occur when a job is writing to a tape drive, but the tape drive is waiting for new tape to be inserted. If you start a new job in this situation, then you will probably receive the enqueue message because the first job cannot complete until the new tape is loaded.

After you have determined which job is creating the enqueue, you can do one of the following:

- Wait until the job holding the enqueue completes
- Cancel the current job and restart it after the job holding the enqueue completes
- Cancel the job creating the enqueue

RMAN Fails to Delete All Archived Logs: Scenario

In this scenario, the database archives automatically to two directories: *ORACLE_HOME/oradata/trgt/arch* and *ORACLE_HOME/oradata/trgt/arch2*. You tell RMAN to perform a backup and delete the input archived redo logs afterward in the following script:

```
BACKUP ARCHIVELOG ALL DELETE INPUT;
```

You then run a crosscheck to make sure the logs are gone and find the following:

```
CROSSCHECK ARCHIVELOG ALL;
```

```
validation succeeded for archived log
archivelog filename=/oracle/oradata/trgt/arch2/archive1_964.arc recid=19
stamp=368726072
```

RMAN deleted one set of logs but not the other.

RMAN Fails to Delete All Archived Logs: Diagnosis

This problem is not an error. When you specify `DELETE INPUT` without the `ALL` keyword, RMAN deletes only one copy of each input log. Even if you archive to five destinations, RMAN deletes logs from only one directory.

RMAN Fails to Delete All Archived Logs: Solution

To force RMAN to delete all existing archived redo logs, use the `DELETE ALL INPUT` clause of the `BACKUP` command. For example, enter:

```
BACKUP ARCHIVELOG ALL DELETE ALL INPUT;
```

Backup Fails Because RMAN Cannot Locate an Archived Log: Scenario

In this scenario, you schedule regular backups of the archived redo logs. The next time you make a backup, you receive this error:

```
RMAN-6089: archive log NAME not found or out of sync with catalog
```

Backup Fails Because RMAN Cannot Locate an Archived Log: Diagnosis

This problem occurs when the archived log that RMAN is looking for cannot be accessed by RMAN, or the recovery catalog needs to be resynchronized. Often, this error occurs when you delete archived logs with an operating system command, which means that RMAN is unaware of the deletion. The `RMAN-6089` error occurs because RMAN attempts to back up a log that the repository indicates still exists.

Backup Fails Because RMAN Cannot Locate an Archived Log: Solution

Make sure that the archived logs exist in the specified directory and that the RMAN catalog is synchronized. Check the following:

1. Make sure the archived log file that is specified by the `RMAN-6089` error exists in the correct directory.
2. Check that the operating system permissions are correct for the archived log (`owner = oracle, group = DBA`) to make sure that RMAN can access the file.
3. If the file appears to be correct, then try synchronizing the catalog by running the following command from the RMAN prompt:

```
RESYNC CATALOG;
```

If you know that the logs are unavailable because you deleted them by using an operating system utility, then run the following command at the RMAN prompt to update RMAN metadata:

```
CROSSCHECK ARCHIVELOG ALL;
```

It is always better to use RMAN to delete logs than to use an operating system utility. The easiest method to remove unwanted logs is to specify the `DELETE INPUT` option when backing up archived logs. For example, enter:

```
BACKUP DEVICE TYPE sbt
ARCHIVELOG ALL
DELETE ALL INPUT;
```

RMAN Does Not Recognize Character Set Name: Scenario

In this scenario, you are connected to the target database while it is not open and attempting to perform an RMAN operation. You receive the following error:

PLS-00553: character set name is not recognized

RMAN Does Not Recognize Character Set Name: Diagnosis

Typically, this message means that the character set in the client environment, that is, the environment in which you are running the RMAN client, is different from the character set in the target database environment.

RMAN Does Not Recognize Character Set Name: Solution

1. Query the target database to determine the value of the NLS_CHARACTERSET parameter. For example, run this query:

```
SQL> SELECT VALUE FROM V$NLS_PARAMETERS WHERE PARAMETER='NLS_CHARACTERSET';
```

2. Set the character set environment variable in the client to the same value as the variable in the server. For example, you can set the NLS_LANG environment variable on a UNIX system as follows:

```
% setenv NLS_LANG american_america.we8dec
% setenv NLS_DATE_FORMAT "MON DD YYYY HH24:MI:SS"
```

If the connection is made through a listener, then the listener must be started with the correct Globalization Support settings. Otherwise, the spawned connections inherit the incorrect Globalization Support settings from the listener.

RMAN Denies Logon to Target Database: Scenario

RMAN fails with ORA-01031 (insufficient privileges) or ORA-01017 (invalid username/password) errors when trying to connect to the target database:

```
% rman
Recovery Manager: Release 10.1.0.2.0 - Production

Copyright (c) 1995, 2003, Oracle. All rights reserved.

RMAN> CONNECT TARGET sys/mypass@inst1

RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
ORA-01031: insufficient privileges
```

RMAN Denies Logon to Target Database: Diagnosis

RMAN automatically requests a connection to the target database as SYSDBA. In order to connect to the target as SYSDBA, you must do one of the following:

- Be part of the operating system DBA group with respect to the target database (that is, have the ability to connect with SYSDBA privileges to the target database without a password).
- Create a password file with the `orapwd` command and the initialization parameter `REMOTE_LOGIN_PASSWORDFILE`.
- Make sure you are connecting with the correct username and password.

If the target database does not have a password file, then the user you are logged in as must be validated with operating system authentication.

RMAN Denies Logon to Target Database: Solution

Either create a password file for the target database or add yourself to the administrator list in the operating system.

See Also: *Oracle Database Administrator's Guide* to learn how to create a password file

Database Duplication Fails Because of Missing Log: Scenario

In this scenario, you attempt to duplicate a database with the `DUPLICATE` command, but receive the following error stack:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of Duplicate Db command at 09/04/2001 12:11:29
RMAN-03015: error occurred in stored script Memory Script
RMAN-06053: unable to perform media recovery because of missing log
RMAN-06025: no backup of log thread 1 seq 16 scn 145858 found to restore
```

Database Duplication Fails Because of Missing Log: Diagnosis

The problem is that RMAN is not able to apply all the archived logs needed for complete recovery. For example, if you only backed up logs through sequence 15, but the most recent archived log is sequence 16, then `DUPLICATE` fails.

Database Duplication Fails Because of Missing Log: Solution

When creating the duplication script, use the `SET UNTIL` command to specify a log sequence number for incomplete recovery. For example, to terminate recovery after applying log sequence 15, enter:

```
RUN
{
  SET UNTIL SEQUENCE 16 THREAD 1; # recovers up to but not including log 16
  DUPLICATE TARGET DATABASE TO 'dupdb';
}
```

See Also: ["RMAN DUPLICATE DATABASE at a Past Point in Time: Example"](#) on page 13-22 for more information about performing incomplete recovery during the duplication operation

Duplication Fails with Multiple RMAN-06023 Errors: Scenario

In this scenario, you back up the database, then run the `DUPLICATE` command. You receive the following error stack:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of Duplicate Db command at 09/04/2001 13:55:11
RMAN-03015: error occurred in stored script Memory Script
RMAN-06026: some targets not found - aborting restore
RMAN-06023: no backup or copy of datafile 8 found to restore
RMAN-06023: no backup or copy of datafile 7 found to restore
RMAN-06023: no backup or copy of datafile 6 found to restore
RMAN-06023: no backup or copy of datafile 5 found to restore
RMAN-06023: no backup or copy of datafile 4 found to restore
RMAN-06023: no backup or copy of datafile 3 found to restore
RMAN-06023: no backup or copy of datafile 2 found to restore
```

RMAN-06023: no backup or copy of datafile 1 found to restore

Duplication Fails with Multiple RMAN-06023 Errors: Diagnosis

The DUPLICATE command recovers to archived redo logs, but cannot recover into online redo logs. Thus, if the restored backup cannot be made consistent without applying the online redo logs, then duplication fails with RMAN-06023 errors because RMAN is looking for backups created *before* the most recent archived log.

Duplication Fails with Multiple RMAN-06023 Errors: Solution

After backing up the source database, archive and back up the current redo log:

```
RMAN> SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';
RMAN> BACKUP ARCHIVELOG ALL;
```

This archives all records in the online redo logs so that RMAN can now recover the backup by applying the most recent archived redo log.

UNKNOWN Database Name Appears in Recovery Catalog: Scenario

In this scenario, you list the database incarnations registered in the recovery catalog and see a database with the name UNKNOWN:

```
LIST INCARNATION OF DATABASE;
```

```
RMAN-03022: compiling command: list
```

```
List of Database Incarnations
```

DB Key	Inc Key	DB Name	DB ID	STATUS	Reset SCN	Reset Time
56	57	TRGT	4052472287	CURRENT	1	Sep 03 2001 06:45:51
1	19	UNKNOWN	4141147584	PARENT	1	Jan 08 2001 14:47:28
.

UNKNOWN Database Name Appears in Recovery Catalog: Diagnosis

One way you get the DB_NAME of UNKNOWN is when you register a database that was once opened with the RESETLOGS option. The DB_NAME can be changed during a RESETLOGS operation, so RMAN does not know what the DB_NAME was for those old incarnations of the database because it was not registered in the recovery catalog at the time. Consequently, RMAN sets the DB_NAME column to UNKNOWN when creating the DBINC record.

UNKNOWN Database Name Appears in Recovery Catalog: Solution

The UNKNOWN name entry is expected behavior after a RESETLOGS operation. You should *not* attempt to remove UNKNOWN entries from the recovery catalog.

Part III

Using RMAN for Database Transport, Duplication and Migration

The following chapters describe how to use RMAN as a utility for database and tablespace transport and migration tasks. This part of the book contains these chapters:

- [Chapter 13, "Creating and Updating Duplicate Databases with RMAN"](#)
- [Chapter 14, "Creating Transportable Tablespace Sets from Backup with RMAN"](#)
- [Chapter 15, "RMAN Cross-Platform Transportable Databases and Tablespaces"](#)
- [Chapter 16, "Migrating Databases To and From ASM with Recovery Manager"](#)

Creating and Updating Duplicate Databases with RMAN

This chapter describes how to use the `DUPLICATE` command to create a duplicate database for testing purposes. This chapter contains these topics:

- [Creating a Duplicate Database with RMAN: Overview](#)
- [Renaming Database Files in RMAN Duplicate Database](#)
- [Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps](#)
- [Creating a Duplicate Database on a Local or Remote Host](#)
- [Using RMAN DUPLICATE DATABASE: Examples](#)
- [Using DUPLICATE DATABASE to Resynchronize a Duplicate Database: Example](#)
- [Using RMAN Incremental Backups to Refresh a Standby Database](#)

See Also: *Oracle Data Guard Concepts and Administration* to learn how to create a standby database with the `DUPLICATE` command

Creating a Duplicate Database with RMAN: Overview

You can use the RMAN `DUPLICATE` command to create a **duplicate database** from backups of the **target database** (primary database) while retaining the original target database. The duplicate database can be identical to the target database or contain only a subset of the tablespaces in the target database. The **target site** and the **duplicate site** can be on separate hosts or on the same host.

A duplicate database is a copy of a target database that you can run independently for a variety of purposes. For example, you can use it to:

- Test backup and recovery procedures
- Export data such as a table that was inadvertently dropped from the production database, and then import it back into the production database

For example, you can duplicate the production database on `host1` to `host2`, and then use the duplicate database on `host2` to practice restore and recovery scenarios while the production database on `host1` continues as usual.

A duplicate database is distinct from a standby database, although both types of databases are created with the `DUPLICATE` command. A standby database is a copy of the primary database that you can update continually or periodically with archived logs from the primary database. If the primary database is damaged or destroyed, then you can perform failover to the standby database and transform it into the new primary database. A duplicate database, on the other hand, cannot be used in this

way: it is not intended for failover scenarios and does *not* support the various standby recovery and failover options.

See Also: *Oracle Data Guard Concepts and Administration* to learn how to create a standby database with the `DUPLICATE` command

How Recovery Manager Duplicates a Database

To prepare for database duplication, first create an **auxiliary instance** as described in "[Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps](#)" on page 13-8. For the duplication to work, you must connect RMAN to both the target (primary) database and an auxiliary instance started in `NOMOUNT` mode.

Allocate at least one auxiliary channel on the auxiliary instance. The principal work of the duplication is performed by the auxiliary channel, which starts a server session on the duplicate host. This channel then restores the necessary backups of the primary database, uses them to create the duplicate database, and initiates recovery.

So long as RMAN is able to connect to the primary and auxiliary instances, the RMAN client can run on any host. All backups and archived redo logs used for creating and recovering the duplicate database, however, must be accessible by the server session on the duplicate host. If the duplicate host is not the same as the target host, then you must make backups on disk on the target host available to the duplicate host with the same full path name as in the primary database.

When using disk backups, you can accomplish this goal in any of the following ways:

- Manually transfer the backups from the primary host to the remote host to an identical path. For example, if the backups are in `/disk1/bkp` on the target host, then transfer them to `/disk1/bkp` on the duplicate host.
- Manually transfer the backups from the primary host to the duplicate host at a new location. For example, if the backups are in `/disk1/bkp` on the target host, then you might transfer them to `/disk2/dup` on the duplicate host. The new path—in this example, `/disk2/dup`—must be accessible from both the target and duplicate hosts. Run the `CATALOG` command to add these copies to the RMAN repository at the duplicate host.
- Use NFS or shared disks and make sure that the same path is accessible in the remote host. For example, the NFS mount point for both hosts could be `/home/file_server`.

When using tape backups, you must make the tapes containing the backups accessible to the remote node. You can achieve this goal by physically moving the tape to a drive attached to the remote host or by means of a network-accessible tape server.

As part of the duplicating operation, RMAN automates the following steps:

- Creates a control file for the duplicate database
- Restores the target datafiles to the duplicate database and performs incomplete recovery by using all available incremental backups and archived redo logs
- Shuts down and starts the auxiliary instance (refer to "[Task 4: Start the Auxiliary Instance](#)" on page 13-10 for issues relating to client-side versus server-side initialization parameter files)
- Opens the duplicate database with the `RESETLOGS` option after incomplete recovery to create the online redo logs (except when running `DUPLICATE . . . FOR STANDBY`, in which case RMAN does not open the database)

- Generates a new, unique DBID for the duplicate database (except when you create a standby database with `DUPLICATE . . . FOR STANDBY`, in which case RMAN does *not* create a unique DBID)

During duplication, RMAN *must* perform incomplete recovery because the online redo logs in the target are not backed up and cannot be applied to the duplicate database. The farthest that RMAN can go in recovery of the duplicate database is the most recent redo log archived by the target database.

See Also: *Oracle Data Guard Concepts and Administration* to learn how to create a standby database with RMAN

RMAN DUPLICATE DATABASE: Options

When duplicating a database, you have the following options:

- You can run the `DUPLICATE` command with or without a recovery catalog.
- You can skip read-only tablespaces with the `SKIP READONLY` clause. Read-only tablespaces are included by default. If you omit them, then you can add them later.
- You can exclude tablespaces from the duplicate database with the `SKIP TABLESPACE` clause. You can exclude any tablespace except the `SYSTEM` tablespace or tablespaces containing rollback or undo segments.
- You can create the duplicate database in a new host. If the directory structure is the same on the new host, then you can specify the `NOFILENAMECHECK` option and reuse the target datafile filenames for the duplicate datafiles.
- You can duplicate a target database stored on a traditional file system to an ASM or Oracle Managed Files location.
- By default, the `DUPLICATE` command creates the duplicate database from the most recent backups of the target database and then performs recovery to the most recent consistent point contained in the archived redo logs. You can duplicate a database as it stood at a past point in time in the current incarnation, by using a `RUN` block with a `SET UNTIL` command, or by including an `UNTIL` clause with the `DUPLICATE` command to cause RMAN to recover the duplicate database to a past point in time within the current incarnation. (You cannot, however, use `DUPLICATE` with a point in time in an earlier incarnation.)
- You can register the duplicate database in the same recovery catalog as the target database. This option is possible because RMAN gives the duplicate database a new, unique DBID during duplication.

Note: If you copy the target database by means of operating system utilities, then the DBID of the copied database remains the same as the original database. To register the copy database in the same recovery catalog with the original, you must change the DBID with the `DBNEWID` utility (refer to *Oracle Database Utilities*).

- In some cases, you can set the duplicate database `DB_NAME` differently from the target database `DB_NAME`. More specifically, if the duplicate database exists in the same Oracle home as the target, then the `DB_NAME` initialization parameter *must* be different. If the duplicate database is in a different Oracle home from the target database, then the `DB_NAME` setting for the duplicate database must be unique among databases in its Oracle home. This is true whether or not the duplicate database is on the same host as the target.

RMAN DUPLICATE DATABASE: Prerequisites and Restrictions

Duplicating a database with RMAN involves a number of prerequisites and restrictions. Review the restrictions section of the `DUPLICATE` command in the *Oracle Database Backup and Recovery Reference* for a complete list.

Renaming Database Files in RMAN Duplicate Database

When duplicating a database, RMAN generates names for the required database files. This section describes the different stages of file creation:

- [Renaming Control Files in RMAN DUPLICATE DATABASE](#)
- [Renaming Online Logs in RMAN DUPLICATE DATABASE](#)
- [Renaming Datafiles in RMAN DUPLICATE DATABASE](#)
- [Renaming Tempfiles in RMAN DUPLICATE DATABASE](#)

Renaming Control Files in RMAN DUPLICATE DATABASE

The rules of precedence for determining the names for the duplicate database control files are the same used by the SQL statement `CREATE CONTROLFILE`. When choosing names for the duplicate control files, make sure you set the parameters in the initialization parameter file of the auxiliary database correctly; otherwise, you could overwrite the control files of the target database.

Renaming Online Logs in RMAN DUPLICATE DATABASE

RMAN needs new names for the online redo log files of the duplicate database. Either you can specify the names explicitly in the `DUPLICATE` command, or you can let RMAN generate them according to the rules listed in [Table 13-1](#).

Table 13-1 Order of Precedence for Online Redo Log Filename Generation

Order	Method	Result
1	Specify the <code>LOGFILE</code> clause of <code>DUPLICATE</code> command.	Creates online redo logs as specified.
2	Set <code>LOG_FILE_NAME_CONVERT</code> initialization parameter.	Transforms target filenames, for example, from <code>log_*</code> to <code>duplog_*</code> . Note that you can specify multiple conversion pairs. For details on the use of <code>LOG_FILE_NAME_CONVERT</code> with Oracle Managed Files (OMF), see "Initialization Parameters for RMAN DUPLICATE to OMF Storage" on page 13-17. RMAN uses the <code>REUSE</code> parameter when creating the logs. If an online log file already exists at the named location and is of the correct size, it is reused for the duplicate.
3	Set one of the Oracle Managed Files initialization parameters <code>DB_CREATE_FILE_DEST</code> , <code>DB_CREATE_ONLINE_DEST_n</code> , or <code>DB_RECOVERY_FILE_DEST</code> .	Transforms target filenames based on the parameters set. The rules of precedence among these parameters are the same used by the SQL statement <code>ALTER DATABASE ADD LOGFILE</code> .

Table 13–1 (Cont.) Order of Precedence for Online Redo Log Filename Generation

Order	Method	Result
4	Do none of the preceding steps.	Makes the duplicate filenames the same as the filenames from the target. You must specify the <code>NOFILENAMECHECK</code> option when using this method, and the duplicate database should be in a different host so that the online logs of the duplicate do not conflict with the originals.

Rules higher in the order of precedence override rules lower in the list. For example, if you specify both the `LOGFILE` clause and the `LOG_FILE_NAME_CONVERT` parameter, then RMAN uses the `LOGFILE` clause.

Caution:

- If the target and duplicate databases are in the same host, then do not use the name of an online redo log currently in use by the target database.
 - Do not use the name of an online log currently in use by the target database if the duplicate database is in a different host and `NOFILENAMECHECK` is not used.
-
-

Renaming Datafiles in RMAN DUPLICATE DATABASE

There are several means of specifying new names to be used for the datafiles of your duplicate database. Listed in order of precedence, they are:

1. Use the RMAN command `SET NEWNAME FOR DATAFILE` within a `RUN` block that encloses both the `SET NEWNAME` commands and the `DUPLICATE` command.
2. Use the RMAN command `CONFIGURE AUXNAME` to specify new names for existing datafiles. Run the `CONFIGURE AUXNAME` command before the `DUPLICATE` command.
3. Specify the `DB_FILE_NAME_CONVERT` parameter on the `DUPLICATE` command to specify a rule for converting filenames for any datafiles not renamed with `SET NEWNAME` or `CONFIGURE AUXNAME`.

Note: The `DB_FILE_NAME_CONVERT` clause of the `DUPLICATE` command cannot be used to control generation of new names for files at the duplicate instance which are Oracle Managed Files (OMF) at the target instance. See *Oracle Database Backup and Recovery Reference* for details on this restriction.

4. Set the `DB_FILE_NAME_CONVERT` initialization parameter.

Note: The `DB_FILE_NAME_CONVERT` initialization parameter cannot be used to control generation of new names for files at the duplicate instance which are Oracle Managed Files (OMF) at the target instance. It is subject to the same semantics and limitations as the `DB_FILE_NAME_CONVERT` parameter to the `DUPLICATE` command. It See *Oracle Database Backup and Recovery Reference* for details.

5. Set the `DB_CREATE_FILE_DEST` initialization parameter to create Oracle Managed Files datafiles at the specified location.

If you do not use any of the preceding options, then the duplicate database reuses the original datafile locations from the target database.

Preventing Filename Checking when Duplicating a Database

It is possible for `CONFIGURE AUXNAME`, `SET NEWNAME`, or `DB_FILE_NAME_CONVERT` to generate a name that is already in use in the target database. In this case, specify `NOFILENAMECHECK` on the `DUPLICATE` command to avoid an error message.

For example, assume that the host A database has two files: datafile 1 is named `/oracle/data/file1.f` and datafile 2 is named `/oracle/data/file2.f`. When duplicating to host B, you use a configured channel to duplicate as follows:

```
RUN
{
  SET NEWNAME FOR DATAFILE 1 TO /oracle/data/file2.f; # rename df 1 as file2.f
  SET NEWNAME FOR DATAFILE 2 TO /oracle/data/file1.f; # rename df 2 as file1.f
  DUPLICATE TARGET DATABASE TO newdb;
}
```

Even though you issued `SET NEWNAME` commands for all the datafiles, the `DUPLICATE` command fails because the duplicate filenames are still in use in the target database. Although datafile 1 in the target is not using `/oracle/data/file2.f`, and datafile 2 in the target is not using `/oracle/data/file1.f`, the target filename is used by one of the duplicate datafiles. Thus, you must specify `DUPLICATE . . . NOFILENAMECHECK` to avoid an error.

Renaming Tempfiles in RMAN DUPLICATE DATABASE

RMAN re-creates datafiles for temporary tablespaces as part of the process of duplicating a database. There are several means of specifying locations for duplicate database tempfiles. Listed in order of precedence, they are:

1. Use the `SET NEWNAME FOR TEMPFILE` command within a `RUN` block that encloses both the `SET NEWNAME` commands and the `DUPLICATE` command.
2. Specify the `DB_FILE_NAME_CONVERT` clause to the `DUPLICATE` command to specify a rule for converting tempfiles not renamed with `SET NEWNAME` or `CONFIGURE AUXNAME`.

Note: The `DB_FILE_NAME_CONVERT` clause cannot be used to control generation of new names for files at the duplicate instance which are Oracle Managed Files (OMF) at the target instance. See *Oracle Database Backup and Recovery Reference* for details on this restriction.

3. Set the `DB_FILE_NAME_CONVERT` initialization parameter.

Note: The `DB_FILE_NAME_CONVERT` initialization parameter is subject to the same semantics and limitations as the `DB_FILE_NAME_CONVERT` parameter to the `DUPLICATE` command. See *Oracle Database Backup and Recovery Reference* for details.

4. Set the `DB_CREATE_FILE_DEST` initialization parameter to create Oracle Managed Files tempfiles.

Skipping Tablespaces in RMAN DUPLICATE DATABASE

It is not always necessary to duplicate all tablespaces of a database. For example, you may plan to generate reports at the duplicate that require only a subset of tablespaces from your target database. `RMAN DUPLICATE DATABASE` can skip read-only tablespaces or tablespaces that are currently `OFFLINE NORMAL`. This section describes how to skip such tablespaces during `DUPLICATE`. It contains the following topics:

- [Skipping Read-Only Tablespaces When Duplicating a Database](#)
- [Skipping OFFLINE NORMAL Tablespaces When Duplicating a Database](#)

Skipping Read-Only Tablespaces When Duplicating a Database

When you specify `SKIP READONLY`, RMAN does not duplicate the datafiles of read-only tablespaces. After duplication is complete, you can query the views in the duplicate database described in [Table 13-2](#) and [Table 13-3](#) to determine which datafiles were skipped. The `STATUS` and `ENABLED` columns describe the current status of the duplicate datafile.

Table 13-2 Read-Only Tablespaces in V\$DATAFILE View on Duplicate Database

In the column ...	The value is ...
<code>STATUS</code>	<code>OFFLINE</code>
<code>ENABLED</code>	<code>READ ONLY</code>
<code>NAME</code>	<code>MISSINGxxx</code>

Table 13-3 Read-Only Tablespaces in Data Dictionary Views on Duplicate Database

View	In the column ...	The value is ...
<code>DBA_DATA_FILES</code>	<code>STATUS</code>	<code>AVAILABLE</code>
<code>DBA_TABLESPACES</code>	<code>STATUS</code>	<code>READ ONLY</code>

Skipping OFFLINE NORMAL Tablespaces When Duplicating a Database

When tablespaces are taken offline with the `OFFLINE NORMAL` option before a `DUPLICATE` operation, RMAN does not duplicate the datafiles of these tablespaces. After duplication, you can manually add or drop these tablespaces. Query the views in the duplicate database described in [Table 13-4](#) and [Table 13-5](#) to determine which datafiles are offline, based on the `STATUS` and `ENABLED` columns.

Table 13–4 Offline Tablespaces in V\$ Views on Duplicate Database

In the column ...	The value is ...
STATUS	OFFLINE
ENABLED	DISABLED
NAME	MISSINGxxx

Table 13–5 Offline Tablespaces in Data Dictionary Views on Duplicate Database

View	In the column ...	The value is ...
DBA_DATA_FILES	STATUS	AVAILABLE
DBA_TABLESPACES	STATUS	OFFLINE

Note: RMAN duplicates tablespaces that are taken offline with the IMMEDIATE option, because such a tablespace requires recovery. As with online tablespaces, RMAN requires a valid backup for duplication.

Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps

To prepare the auxiliary instance used in RMAN DUPLICATE DATABASE, carry out the following tasks:

- [Task 1: Create an Oracle Password File for the Auxiliary Instance](#)
- [Task 2: Establish Oracle Net Connectivity to the Auxiliary Instance](#)
- [Task 3: Create an Initialization Parameter File for the Auxiliary Instance](#)
- [Task 4: Start the Auxiliary Instance](#)
- [Task 5: Mount or Open the Target Database](#)
- [Task 6: Make Sure You Have the Necessary Backups and Archived Redo Logs](#)
- [Task 7: Allocate Auxiliary Channels if Automatic Channels Are Not Configured](#)

Task 1: Create an Oracle Password File for the Auxiliary Instance

For instructions on how to create and maintain Oracle password files, refer to *Oracle Database Administrator's Guide*.

Task 2: Establish Oracle Net Connectivity to the Auxiliary Instance

The auxiliary instance must be accessible through Oracle Net. Before proceeding, start SQL*Plus to ensure that you can establish a connection to the auxiliary instance. Note that you must connect to the auxiliary instance with SYSDBA privileges, so a password file must exist.

Task 3: Create an Initialization Parameter File for the Auxiliary Instance

Create a client-side initialization parameter file for the auxiliary instance. The parameters listed in [Table 13–6](#) must be set to the values specified.

Table 13–6 Database Name and Control File Name Initialization Parameters

Parameter	Required Value
DB_NAME	The same name used in the DUPLICATE command. You must set the DB_NAME parameter in the duplicate parameter file to the same database name specified in the DUPLICATE command. You cannot use the same database name for the target and duplicate when the duplicate is in the same Oracle home as the target. If the duplicate is in a different Oracle home from the target, then its DB_NAME just has to differ from other database names in the same Oracle home.
CONTROL_FILES	Refer to "Renaming Control Files in RMAN DUPLICATE DATABASE" on page 13-4.

The block size for the auxiliary database must match that of the target database. If the target database parameter file contains a value for the DB_BLOCK_SIZE initialization parameter, then you must specify the same value in the auxiliary instance parameter file. If no DB_BLOCK_SIZE is specified in the target database parameter file, then do not specify DB_BLOCK_SIZE in the auxiliary instance parameter file.

You can also set the initialization parameters described in [Table 13–7](#).

Table 13–7 Filename Conversion Initialization Parameters

Initialization Parameter	You must specify:
DB_FILE_NAME_CONVERT	Refer to "Renaming Datafiles in RMAN DUPLICATE DATABASE" on page 13-5. You can also specify DB_FILE_NAME_CONVERT as a clause for the DUPLICATE command itself.
LOG_FILE_NAME_CONVERT	Refer to "Renaming Online Logs in RMAN DUPLICATE DATABASE" on page 13-4.

If you do not set the conversion parameters shown in [Table 13–7](#), then you can set the DB_CREATE_FILE_DEST, DB_CREATE_ONLINE_LOG_DEST_1, or DB_RECOVERY_FILE_DEST initialization parameter. In this case, the DUPLICATE command creates the datafiles, online redo logs, and tempfiles as Oracle Managed Files. If the CONTROL_FILES initialization parameter is not set, then the DUPLICATE command also creates the control file as an Oracle Managed File.

Set other initialization parameters, including the parameters that allow you to connect as SYSDBA through Oracle Net, as needed.

When duplicating to the same host or to a new host with a different file system, pay attention to all initialization parameters specifying path names. Verify that all paths are accessible on the host where the database is being duplicated.

[Example 13–1](#) shows sample initialization parameter settings for the duplicate database.

Example 13–1 Sample Initialization Parameters for Duplicate Database

```
DB_NAME=newdb
CONTROL_FILES=(/dup/oracle/oradata/trgt/control01.ct1,
               /dup/oracle/oradata/trgt/control02.ct1)
# note that the following two initialization parameters have equivalents
# on the DUPLICATE command itself
DB_FILE_NAME_CONVERT=(/oracle/oradata/trgt/, /dup/oracle/oradata/trgt/)
LOG_FILE_NAME_CONVERT=(/oracle/oradata/trgt/redo, /dup/oracle/oradata/trgt/redo)
```

After you create the client-side initialization parameter file, you can run the `CREATE SPFILE` command from SQL*Plus to create a server-side initialization parameter file. You can run this command before or after instance startup. For example, you can create a server-side parameter file in the default location as follows, specifying the filename of the client-side initialization parameter file in the `FROM` clause:

```
CREATE SPFILE FROM PFILE='/tmp/initDUPDB.ora';
```

A server-side parameter file in the default location is an advantage when duplicating a database because you do not need to specify the `PFILE` parameter on the `DUPLICATE` command. Because RMAN shuts down and restarts the auxiliary instance as part of the duplication process, you must tell RMAN which client-side file to use if you use a client-side parameter file. It is highly recommended that you create a server-side parameter file for use in database duplication.

Task 4: Start the Auxiliary Instance

Before beginning RMAN duplication, use SQL*Plus to connect to the auxiliary instance and start it in `NOMOUNT` mode (specifying a client-side parameter file if necessary). In this example, `oracle` is the password for the user with `SYSDBA` authority and `aux` is the net service name for the auxiliary instance:

```
CONNECT SYS/oracle@aux AS SYSDBA
-- start instance with the server parameter file
STARTUP FORCE NOMOUNT
```

Because the auxiliary instance does not yet have a control file, you can only start the instance in `NOMOUNT` mode. Do not create a control file or try to mount or open the auxiliary instance.

RMAN shuts down and restarts the auxiliary instance as part of the duplication. Hence, it is a good idea to create a server-side initialization parameter file for the auxiliary instance in the default location.

If you do not have a server-side initialization parameter file for the auxiliary instance in the default location, then you must specify the client-side initialization parameter file with the `PFILE` parameter on the `DUPLICATE` command. The client-side parameter file for the auxiliary instance must reside on the same host as the RMAN client used to perform the duplication.

Task 5: Mount or Open the Target Database

Before beginning RMAN duplication, connect SQL*Plus to the target database and mount or open it if it is not already mounted or open. For example, enter:

```
-- connect to target database
SQL> CONNECT SYS/oracle@trgt AS SYSDBA;
-- mount or open target database
STARTUP MOUNT;
```

Task 6: Make Sure You Have the Necessary Backups and Archived Redo Logs

Make sure backups all target datafiles are accessible on the duplicate host. If you do not have backups of everything, then the duplicate operation fails. The database backup does not have to be a whole database backup: you can use a mix of full and incremental backups of individual datafiles.

Archived redo logs required to recover the duplicate database to the desired point in time must be accessible at the same path by the node where the duplicate database is

to be created. These can be available either as backups (for instance, on a media manager) or as image copies (or the actual archived redo logs). The backups or copies can be transferred to the local disk of the node that contains the duplicate database or possibly mounted across a network by some means such as NFS.

Task 7: Allocate Auxiliary Channels if Automatic Channels Are Not Configured

Start RMAN with a connection to the target database, the auxiliary instance, and, if applicable, the recovery catalog database. You can start the RMAN client on any host so long as it can connect to all the instances. If the auxiliary instance requires a client-side initialization parameter file, then this file must exist on the same host that runs the RMAN client.

In this example, a connection is established to three instances, all through the use of net service names:

```
% rman TARGET SYS/oracle@trgt CATALOG rman/cat@catdb AUXILIARY SYS/oracle@aux
```

If you do not have automatic channels configured, then before issuing the `DUPLICATE` command, manually allocate at least one auxiliary channel within the same `RUN` command. The channel type (`DISK` or `sbt`) must match the media where the backups of the target database are located. If the backups reside on disk, then the more channels you allocate, the faster the duplication will be. For tape backups, limit the number of channels to the number of devices available.

```
RUN
{
  # to manually allocate a channel of type sbt issue:
  ALLOCATE AUXILIARY CHANNEL ch1 DEVICE TYPE sbt;

  # to manually allocate three auxiliary channels for disk issue (specifying
  # whatever channel id that you want):
  ALLOCATE AUXILIARY CHANNEL aux1 DEVICE TYPE DISK;
  ALLOCATE AUXILIARY CHANNEL aux2 DEVICE TYPE DISK;
  ALLOCATE AUXILIARY CHANNEL aux3 DEVICE TYPE DISK;
  .
  .
  .
  DUPLICATE TARGET DATABASE TO dupdb;
}
```

Note: If you configure automatic channels, then RMAN can use configured channels for duplication even if they do not specify the `AUXILIARY` option. Nevertheless, if the auxiliary channels need some special parameters (for example, to point to a different media management subsystem), then you can configure an automatic channel with the `AUXILIARY` option.

Creating a Duplicate Database on a Local or Remote Host

The procedure to create a duplicate database depends on your configuration. This section contains the following topics:

- [Duplicating a Database on a Remote Host with the Same Directory Structure](#)
- [Duplicating a Database on a Remote Host with a Different Directory Structure](#)
- [Creating a Duplicate Database on the Local Host](#)

- [Using RMAN DUPLICATE with OMF and ASM](#)
- [Duplicating a Database to ASM Storage](#)

Duplicating a Database on a Remote Host with the Same Directory Structure

The simplest case is to duplicate the database to a different host and to use the same directory structure. In this case, you do not need to change the initialization parameter file or set new filenames for the duplicate datafiles.

1. Follow the steps in "[Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps](#)" on page 13-8.
2. Run the `DUPLICATE` command, making sure to do the following:
 - If automatic channels are not configured, then allocate at least one auxiliary channel.
 - Specify the `NOFILENAMECHECK` parameter on the `DUPLICATE` command.
 - Specify the `PFILE` parameter if starting the auxiliary instance with a client-side parameter file. The client-side parameter file must exist on the same host as the RMAN client used to perform the duplication.

The following example assumes that the RMAN client is running on the duplicate host. It duplicates the database with an automatic channel, specifies a client-side initialization parameter file, and specifies the `NOFILENAMECHECK` option:

```
DUPLICATE TARGET DATABASE TO dupdb
# specify client-side parameter file (on same host as RMAN client) for
# auxiliary instance if necessary
PFILE = /dup/oracle/dbs/initDUPDB.ora
NOFILENAMECHECK;
```

RMAN automatically allocates the configured channels, then uses all incremental backups, archived redo log backups, and archived redo logs to perform incomplete recovery. Finally, RMAN opens the database with the `RESETLOGS` option to create the online redo logs.

Duplicating a Database on a Remote Host with a Different Directory Structure

If you create the duplicate database on a host with a different directory structure, then you must change several initialization parameters in order to generate new filenames for the duplicate database datafiles on the new directory structure. The different techniques available are described in "[Renaming Datafiles in RMAN DUPLICATE DATABASE](#)" on page 13-5.

Converting Filenames with Only Initialization Parameters

This procedure assumes that you use only initialization parameters to rename the duplicate datafiles and log files.

1. Follow the steps in "[Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps](#)" on page 13-8. You can either create a client-side parameter file, or copy the parameter file from its location in the target host directory structure to the same location in the duplicate host directory structure using operating system utilities.

Perform the following tasks:

- Review all initialization parameters that end in `_DEST` and specify a path name. Some may need to be changed.

- Set `DB_FILE_NAME_CONVERT` so that it captures *all* the target datafiles and converts them appropriately, for example, from `/oracle/oradata/` to `/dup/oracle/oradata/`.
- Set `LOG_FILE_NAME_CONVERT` so that it captures *all* the online redo logs and converts them appropriately, for example, `/oracle/oradata/redo` to `/dup/oracle/oradata/redo`.

You can set multiple conversion pairs in `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT`. For example, you can specify that `DB_FILE_NAME_CONVERT` changes `/disk1/dbs` to `/dup1/dbs` and `/disk2/dbs` to `/dub2/dbs`.

You can also use the `CONFIGURE AUXNAME` or `SET NEWNAME` commands to rename individual datafiles if you cannot easily generate all of your desired filenames using `DB_FILE_NAME_CONVERT` patterns. As described in ["Using RMAN DUPLICATE with OMF and ASM"](#) on page 13-16, you can also duplicate your database to an Oracle Managed Files location and let the database generate names for your files.

2. Perform the following operations when running the duplication:

- If automatic channels are not configured, then allocate at least one auxiliary channel.
- If using a client-side parameter file to start the auxiliary instance, specify the `PFILE` parameter.

Converting Filenames with Only Initialization Parameters: Example The following example assumes that the duplicate host can access the same media manager as the primary database host. The example duplicates the database with an automatic `sbt` channel and uses a server-side parameter file located on the duplicate host to restart the auxiliary instance:

```
DUPLICATE
  TARGET DATABASE TO dupdb
  DEVICE TYPE sbt # restores from tape backups;
# DUPLICATE DEVICE TYPE sbt works only if the sbt device is configured
# by CONFIGURE CHANNEL, CONFIGURE DEVICE TYPE, or CONFIGURE DEFAULT DEVICE.
```

RMAN uses all incremental backups, archived redo log backups, and archived redo logs to perform incomplete recovery. RMAN then shuts down, starts, and opens the database with the `RESETLOGS` option to create the online redo logs.

Converting Filenames with Only DUPLICATE Parameters

This procedure assumes that you use the `DB_FILE_NAME_CONVERT` parameter of the `DUPLICATE` command to rename the duplicate datafiles, and the `LOGFILE` clause to specify names and sizes for the online redo logs.

Perform the following operations when running the duplication:

- If automatic auxiliary channels are not configured, then allocate at least one auxiliary channel.
- Specify the names and sizes for the duplicate database redo logs in the `LOGFILE` clause.
- Specify new filenames for the duplicate database datafiles with the `DB_FILE_NAME_CONVERT` parameter.

- If using a client-side parameter file to start the auxiliary instance, specify the `PFILE` parameter.

The following example duplicates the database using configured channels and specifies an initialization parameter file:

```

DUPLICATE TARGET DATABASE TO dupdb
  # specify client-side parameter file for auxiliary instance
  PFILE = /dup/oracle/dbs/initDUPDB.ora
  DB_FILE_NAME_CONVERT=(/oracle/oradata/trgt/, /dup/oracle/oradata/trgt/)
  LOGFILE
    '/dup/oracle/oradata/trgt/redo01.log' SIZE 200K,
    '/dup/oracle/oradata/trgt/redo02.log' SIZE 200K,
    '/dup/oracle/oradata/trgt/redo03.log' SIZE 200K;

```

Using RMAN DUPLICATE With SET NEWNAME: Example

This procedure assumes that you use the `SET NEWNAME` command to rename the duplicate datafiles.

1. Follow the steps in ["Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps"](#) on page 13-8, making sure to use an operating system utility to copy the parameter file from its location in the target host directory structure to the same location in the duplicate host. Set all initialization parameters that end in `_DEST` and specify a path name.
2. Perform the following operations when running the duplication:
 - If channels are not configured for the auxiliary database, then allocate at least one auxiliary channel.
 - If desired, specify the same number of redo log members and groups that are used in the target database.
 - Specify new filenames for the duplicate database datafiles.
 - If you use a client-side parameter file to start the auxiliary instance, then specify the `PFILE` parameter.

The following example uses configured channels and a default server-side initialization parameter file for the database duplication, and uses the `LOGFILE` clause to specify names and sizes for the online redo logs:

```

RUN
{
  # set new filenames for the datafiles
  SET NEWNAME FOR DATAFILE 1 TO '/dup/oracle/oradata/trgt/system01.dbf';
  SET NEWNAME FOR DATAFILE 2 TO '/dup/oracle/oradata/trgt/undotbs01.dbf';
  . . .
  # issue the duplicate command
  DUPLICATE TARGET DATABASE TO dupdb
  # create at least two online redo log groups
  LOGFILE
    GROUP1
      (
        '/dup/oracle/oradata/trgt/redo01a.log',
        '/dup/oracle/oradata/trgt/redo01b.log',
        '/dup/oracle/oradata/trgt/redo01c.log';
      ) SIZE 200K,
    GROUP2
      (
        '/dup/oracle/oradata/trgt/redo02a.log',
        '/dup/oracle/oradata/trgt/redo02b.log',

```

```

        '/dup/oracle/oradata/trgt/redo02c.log';
    ) SIZE 200K,
GROUP3
(
    '/dup/oracle/oradata/trgt/redo03a.log',
    '/dup/oracle/oradata/trgt/redo03b.log',
    '/dup/oracle/oradata/trgt/redo03c.log';
) SIZE 200K;

```

RMAN uses all incremental backups, archived redo log backups, and archived redo logs to perform incomplete recovery. RMAN shuts down, starts up, and then opens the database with the `RESETLOGS` option to create the online logs.

Using RMAN DUPLICATE With CONFIGURE AUXNAME: Example

This procedure assumes that you use the `CONFIGURE AUXNAME` command to rename the duplicate datafiles.

1. Follow the steps in ["Preparing the RMAN DUPLICATE Auxiliary Instance: Basic Steps"](#) on page 13-8, making sure to use an operating system utility to copy the parameter file from its location in the target host directory structure to the same location in the duplicate host directory structure. Set all initialization parameters that end in `_DEST` and specify a path name.
2. Add the following features when creating the RMAN commands to perform the duplication:
 - Prepare `CONFIGURE AUXNAME` commands for all datafiles, to be executed before database duplication.
 - If automatic auxiliary channels are not allocated, then allocate at least one auxiliary channel.
 - Use a `LOGFILE` clause to specify redo log groups and members for the duplicate database. (You do not have to use the same number of redo log groups or redo log group members in the duplicate database as you did in the target database.)
 - If you start the auxiliary instance with a client-side parameter file, then specify the `PFILE` parameter. The client-side parameter file must reside on the same host as the RMAN client used to perform the duplication.

[Example 13–2](#) uses `CONFIGURE AUXNAME` to set new datafile names, uses automatic channels and a client-side initialization parameter file for the database duplication, and uses the `LOGFILE` clause to specify names and sizes for the online redo logs.

Example 13–2 Using CONFIGURE AUXNAME to Generate Database Filenames

```

# configure the new desired filenames
CONFIGURE AUXNAME FOR DATAFILE 1
    TO '/dup/oracle/oradata/trgt/system01.dbf';
CONFIGURE AUXNAME FOR DATAFILE 2
    TO '/dup/oracle/oradata/trgt/undotbs01.dbf';
# ... add more CONFIGURE AUXNAME commands as needed

# run the DUPLICATE command
DUPLICATE TARGET DATABASE TO dupdb
# specify client-side parameter file for auxiliary instance if necessary
PFILE = /dup/oracle/dbs/initDUPDB.ora
.
.

```

```
.
# create at least two online redo log groups
LOGFILE
  GROUP1
  (
    '/dup/oracle/oradata/trgt/redo01a.log',
    '/dup/oracle/oradata/trgt/redo01b.log',
    '/dup/oracle/oradata/trgt/redo01c.log';
  ) SIZE 200K,
  GROUP2
  (
    '/dup/oracle/oradata/trgt/redo02a.log',
    '/dup/oracle/oradata/trgt/redo02b.log',
    '/dup/oracle/oradata/trgt/redo02c.log';
  ) SIZE 200K,
  GROUP3
  (
    '/dup/oracle/oradata/trgt/redo03a.log',
    '/dup/oracle/oradata/trgt/redo03b.log',
    '/dup/oracle/oradata/trgt/redo03c.log';
  ) SIZE 200K;
```

RMAN uses all incremental backups, archived redo log backups, and archived redo logs to perform incomplete recovery and then opens the database with the RESETLOGS option to create the online redo logs.

After the duplication is complete, clear the configured auxiliary names for the datafiles in the duplicate database, so that they are not overwritten by future operations. For example, enter the following:

```
# clear specified auxiliary names for the datafiles
CONFIGURE AUXNAME FOR DATAFILE 1 CLEAR;
CONFIGURE AUXNAME FOR DATAFILE 2 CLEAR;
.
.
.
```

Creating a Duplicate Database on the Local Host

When creating a duplicate database on the same host as the target database, follow the same procedure as for duplicating to a remote host with a different directory structure as described in "[Duplicating a Database on a Remote Host with a Different Directory Structure](#)" on page 13-12.

You can duplicate the database to the same Oracle home as the target database, but you must use a different database name from the target database, and convert the filenames by means of the same methods used for conversion on a separate host.

Caution: Do not use the NOFILENAMECHECK option when duplicating to the same Oracle home as the primary database. If you do, then the DUPLICATE command may overwrite the datafiles of the target database.

Using RMAN DUPLICATE with OMF and ASM

The following sections discuss requirements when creating some or all files of the duplicate database in OMF or ASM storage:

- [Initialization Parameters for RMAN DUPLICATE to OMF Storage](#)

- [Duplicating a Database to ASM Storage](#)
- [Using SET NEWNAME with RMAN DUPLICATE to OMF](#)

Initialization Parameters for RMAN DUPLICATE to OMF Storage

When duplicating a database and storing the duplicate in OMF storage, you must make the following changes on the initialization parameter file for the duplicate database:

- Set the `DB_CREATE_FILE_DEST` initialization parameter to the desired OMF storage location. Any database files for which no other location is specified are created in `DB_CREATE_FILE_DEST` during `DUPLICATE`.
- For a duplicate database that is not a standby database, change the value of the `DB_NAME` initialization parameter. For a standby database, keep `DB_NAME` the same, but change the `DB_UNIQUE_NAME` initialization parameter.
- If duplicating for standby, set the `STANDBY_FILE_MANAGEMENT` initialization parameter to `AUTO`.
- Do not set the `CONTROL_FILES` initialization parameter if you want to create the new control files as Oracle Managed Files. Oracle recommends that you use an `SPFILE` at the duplicate database when using OMF for the control file.
- Do not set the `DB_FILE_NAME_CONVERT` initialization parameter. This will allow the database to generate valid OMF filenames for the duplicate datafiles.

The default location for OMF files is `DB_CREATE_FILE_DEST`. You can override the default for individual files using `SET NEWNAME`, as described in "[Using SET NEWNAME with RMAN DUPLICATE to OMF](#)" on page 13-18.

- Do not set the `LOG_FILE_NAME_CONVERT` initialization parameter. This will allow the database to generate valid OMF redo log filenames.

To direct duplicate database online redo log files to OMF storage, you can use the `DB_CREATE_FILE_DEST`, `DB_RECOVERY_FILE_DEST` or `DB_CREATE_ONLINE_LOG_DEST_n` initialization parameters to identify an OMF location for the online logs. [Table 13–8, "Oracle Managed Files Initialization Parameter Settings"](#) on page 13-17 explains how to choose among these options.

You must also decide:

- How many copies of the control file are needed
- How many online redo logs members are needed
- Whether the online redo log and control file locations should be different from the datafile locations

These decisions determine whether you use the Oracle Managed Files initialization parameters `DB_RECOVERY_FILE_DEST`, `DB_CREATE_FILE_DEST`, or `DB_CREATE_ONLINE_DEST_n`.

[Table 13–8](#) explains the required initialization parameter values, depending upon your requirements.

Table 13–8 Oracle Managed Files Initialization Parameter Settings

Copies of control files and online redo logs	Control file, online log, and datafile locations	Required Initialization parameters
<i>n</i> , where <i>n</i> is 3 or more	Same or different	<code>DB_CREATE_ONLINE_LOG_DEST_n</code>

Table 13–8 (Cont.) Oracle Managed Files Initialization Parameter Settings

Copies of control files and online redo logs	Control file, online log, and datafile locations	Required Initialization parameters
n , where n is 1 or more	Different	DB_CREATE_ONLINE_LOG_DEST_1
2	Same	DB_RECOVERY_FILE_DEST
1	Same	DB_CREATE_FILE_DEST

Duplicating a Database to ASM Storage

The procedure for creating a duplicate database to an Automatic Storage Management (ASM) location is similar to the procedure described in ["Initialization Parameters for RMAN DUPLICATE to OMF Storage"](#) on page 13-17. The difference is that you must edit the initialization parameter file in the auxiliary instance to set all parameters that control the location where files are created at the duplicate, such as DB_CREATE_FILE_DEST, DB_CREATE_ONLINE_DEST_1, and CONTROL_FILES, to the ASM disk group.

For example, edit the auxiliary instance parameter file as follows to create the database files in the disk group disk1:

```
*.DB_CREATE_FILE_DEST = '+disk1'
```

Using SET NEWNAME with RMAN DUPLICATE to OMF

If you want to store specific datafiles or tempfiles in an OMF destination, independent of the locations of the rest of the database files, then you can use the following steps:

- Set the DB_CREATE_FILE_DEST initialization parameter at the auxiliary instance to the desired location
- Enclose the DUPLICATE command in a RUN block


```
Use SET NEWNAME FOR DATAFILE ... TO NEW and SET NEWNAME FOR
TEMPFILE ... TO NEW commands in the RUN block with the DUPLICATE
command
```

The specified files are created with OMF names in the location specified by DB_CREATE_FILE_DEST.

This technique is similar to the scenario described in ["Using RMAN DUPLICATE With SET NEWNAME: Example"](#) on page 13-14

You can also use SET NEWNAME to direct individual datafiles or tempfiles to a specific ASM disk group. For example:

```
RUN {
  SET NEWNAME FOR DATAFILE 1 TO "+dgroup1";
  SET NEWNAME FOR DATAFILE 2 TO "+dgroup2";
  ...
  DUPLICATE TARGET DATABASE FOR STANDBY;
}
```

See also: *Oracle Database Backup and Recovery Reference* for details on using SET NEWNAME

Using RMAN DUPLICATE DATABASE: Examples

This section presents several representative examples for using RMAN DUPLICATE DATABASE.

- [Duplicating When the Datafiles Use Inconsistent Paths: Example](#)
- [RMAN DUPLICATE DATABASE at a Past Point in Time: Example](#)
- [Duplicating with a Client-Side Parameter File: Example](#)

Duplicating When the Datafiles Use Inconsistent Paths: Example

This example assumes the following:

- You are using recovery catalog database `catdb`.
- The target database `trgt` is on `host1` and contains eight datafiles, which are spread out over multiple directories.
- You want to exclude tablespace `tools` from the duplicate database, but keep all of the other tablespaces.
- You want to duplicate the target to database `dupdb` on remote host `host2`.
- `host1` and `host2` cannot mount each other's file systems by any means such as NFS.
- You want to store the datafiles in `host2` in the `/oradata1, /oradata2 ...` through `/oradata7` subdirectories.
- You have used an operating system utility to copy the initialization parameter file from `host1` to an appropriate location in `host2`.
- You have reset all initialization parameters that end in `_DEST` and specify a path name.
- You want two online redo logs groups, each with two members of size 200 KB, which on the duplicate system will be stored in the directory `"/duplogs"`.
- You have disk copies or backup sets stored on disk for all the datafiles and archived redo logs in the target database, and you have manually moved them to `host2` by means of an operating system utility.
- The auxiliary instance uses a server-side initialization parameter file in the default location (so the `PFILE` parameter is not necessary on the `DUPLICATE` command).

[Example 13-3](#) illustrates a script that creates the duplicate database.

Example 13-3 Duplicating When Datafiles Use Inconsistent Paths

```
CONNECT TARGET /;
CONNECT CATALOG rman/cat@catdb;
CONNECT AUXILIARY SYS/oracle@dupdb;

# note that a RUN command is necessary because you can only execute SET NEWNAME
# within a RUN command
RUN
{
  # The DUPLICATE command uses an automatic sbt channel.
  # Because the target datafiles are spread across multiple directories,
  # run SET NEWNAME rather than DB_FILE_NAME_CONVERT
  SET NEWNAME FOR DATAFILE 1 TO '/oradata1/system01.dbf';
  SET NEWNAME FOR DATAFILE 2 TO '/oradata2/undotbs01.dbf';
  SET NEWNAME FOR DATAFILE 3 TO '/oradata3/cwmlite01.dbf';
```

```

SET NEWNAME FOR DATAFILE 4 TO '/oradata4/drsys01';
SET NEWNAME FOR DATAFILE 5 TO '/oradata5/example01.dbf';
SET NEWNAME FOR DATAFILE 6 TO '/oradata6/indx01.dbf';
# Do not set a newname for datafile 7, because it is in the tools tablespace,
# and you are excluding tools from the duplicate database.
SET NEWNAME FOR DATAFILE 8 TO '/oradata7/users01.dbf';
DUPLICATE TARGET DATABASE TO dupdb
  SKIP TABLESPACE tools
  LOGFILE
    GROUP 1 ('/duplogs/redo01a.log',
             '/duplogs/redo01b.log') SIZE 200K REUSE,
    GROUP 2 ('/duplogs/redo02a.log',
             '/duplogs/redo02b.log') SIZE 200K REUSE;
}

```

RMAN DUPLICATE DATABASE From File System to ASM: Example

This duplication example assumes the following:

- You are using recovery catalog database catdb.
- The target database trgt is on host1 and the database files are stored in a non-ASM file system.
- You want to duplicate the target to database dupdb on remote host host2.
- host2 has diskgroup +DISK1.
- You want to store the datafiles for dupdb to +DISK1.
- You want to store two controlfiles in +DISK1.
- The backups and archivelogs created by host1 are accessible by host2.

Create an initialization parameter for auxiliary instance by copying the target database initialization parameter file. Change the parameters as follows:

- Set DB_NAME to the new database name dupdb
- Set CONTROL_FILES to store two copies of the control file in +DISK1
- Make sure DB_FILE_NAME_CONVERT and LOG_FILE_NAME_CONVERT are not set
- Set any other initialization parameters that end in _DEST, such as DB_CREATE_FILE_DEST and DB_CREATE_ONLINE_DEST_n, to reference +DISK1

For example:

```

DB_NAME=dupdb
CONTROL_FILES=+DISK1,+DISK1

```

Create an SPFILE from the parameter file, and start the auxiliary instance:

```

SQL> CONNECT AUXILIARY SYS/oracle@dupdb;
SQL> CREATE SPFILE FROM PFILE=auxiliary instance pfile;
SQL> STARTUP NOMOUNT;

```

Then run the following RMAN commands to duplicate the database:

```

RMAN> CONNECT TARGET /;
RMAN> CONNECT CATALOG rman/cat@catdb;
RMAN> CONNECT AUXILIARY SYS/oracle@dupdb;
RMAN> DUPLICATE TARGET DATABASE TO dupdb
      LOGFILE GROUP 1 ('+DISK1', '+DISK1') SIZE 200K,

```

```
GROUP 2 ('+DISK1', '+DISK1') SIZE 200K;
```

When the DUPLICATE command completes, the duplicate database is created, with datafiles, online logs and control files in ASM disk group +DISK1.

RMAN DUPLICATE DATABASE From ASM to ASM: Example

This duplication example assumes the following:

- You are using recovery catalog database catdb
- The target database trgt is on host1 and contains ASM datafiles and online logs in diskgroup +DISK1
- You want to duplicate the target to database dupdb on remote host host2.
- host2 has diskgroup +DISK2
- You want to store the datafiles for dupdb to +DISK2
- You want to store two controlfiles in +DISK2
- The backups and archivelogs created by host1 are accessible by host2

Create an initialization parameter for auxiliary instance by copying the target database initialization parameter file. Change the parameters as follows:

- Set DB_NAME to the new database name dupdb
- Set CONTROL_FILES to store two copies of the control file in +DISK2
- Set DB_FILE_NAME_CONVERT and LOG_FILE_NAME_CONVERT to convert the datafile and online log file names from +DISK1 to +DISK2
- Set any other initialization parameters that end in _DEST, such as DB_CREATE_FILE_DEST and DB_CREATE_ONLINE_DEST_n, to reference +DISK2

For example:

```
DB_NAME=dupdb
CONTROL_FILES=+DISK2,+DISK2
DB_FILE_NAME_CONVERT=+DISK1,+DISK2
LOG_FILE_NAME_CONVERT=+DISK1,DISK2
```

Create an SPFILE from the parameter file, and start the auxiliary instance:

```
SQL> CONNECT AUXILIARY SYS/oracle@dupdb;
SQL> CREATE SPFILE FROM PFILE=auxiliary instance pfile;
SQL> STARTUP NOMOUNT;
```

Then run the following RMAN commands to duplicate the database:

```
RMAN> CONNECT TARGET /;
RMAN> CONNECT CATALOG rman/cat@catdb;
RMAN> CONNECT AUXILIARY SYS/oracle@dupdb;
RMAN> DUPLICATE TARGET DATABASE TO dupdb;
```

When the DUPLICATE command completes, the duplicate database is created, with datafiles, online logs and control files in ASM disk group +DISK2.

Note: `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT` convert the disk group name portion of the datafile and online log filenames. The new filenames in `+DISK2` are generated by ASM and do not match the original filenames in disk group `+DISK1`.

RMAN DUPLICATE DATABASE at a Past Point in Time: Example

This duplication example assumes the following:

- The target database `trgt` and duplicate database `dupdb` are on different hosts but have exactly the same directory structure.
- You want to name the duplicate database files the same as the target files.
- You are not using a recovery catalog.
- You are using automatic channels for disk and `sbt`, which are already configured.
- You want to recover the duplicate database to one week ago in order to view the data in `prod1` as it appeared at that time (and you have the required backups and logs to recover the duplicate to that point in time).

```
CONNECT TARGET SYS/oracle@trgt
CONNECT AUXILIARY SYS/oracle@dupdb

DUPLICATE TARGET DATABASE TO dupdb
  NOFILENAMECHECK UNTIL TIME 'SYSDATE-7';
```

Duplicating with a Client-Side Parameter File: Example

You can run the `DUPLICATE` command using a client-side parameter file for the auxiliary instance. Assume the following scenario:

- The target host is `host_src` and the duplicate host is `host_dup`.
- The client-side initialization parameter file on `host_dup` is named `/tmp/initTEST.ora`.
- The hosts `host_dup` and `host_src` are linked by a network.

In this scenario, you can run the RMAN client (that is, run the `DUPLICATE` command in an RMAN session) either on `host_src` or `host_dup`.

Running RMAN Client on `host_dup`

If you run the RMAN client on `host_dup`, you can duplicate the database as follows:

```
DUPLICATE
  TARGET DATABASE TO dupdb
  DEVICE TYPE sbt
  PFILE='/tmp/initTEST.ora';
```

Because the initialization parameter file used by the auxiliary instance resides on the same node as the RMAN client, you can reference the parameter file by its path name on the local system.

Running RMAN from `host_src`

You can run the RMAN client on the same host as the target database rather than on the host with the duplicate database. In such a case, the client-side initialization

parameter file needed by the DUPLICATE command is not located on the same node as the RMAN client.

One option is to copy the parameter file from `host_dup` to `host_src`. If the source host and duplicate host have access to the same file system by some means such as NFS, then you can also remotely mount the directory containing the parameter file by some means such as NFS, and access it directly from `host_src`.

The following example starts RMAN on `host_src` and duplicates the database. In this scenario, assume `/net/host_src/tmp` is an NFS mount point for `/net/host_dup/tmp`, so RMAN is able to access the auxiliary parameter file.

```
% rman TARGET SYS/oracle@trgt AUXILIARY SYS/oracle@dupdb
RMAN> DUPLICATE TARGET DATABASE TO dupdb
        DEVICE TYPE sbt PFILE='/net/host_src/tmp/initTEST.ora';
RMAN> EXIT
```

Using DUPLICATE DATABASE to Resynchronize a Duplicate Database: Example

You may want to synchronize a duplicate database with the original database that was duplicated. The solution is to run the DUPLICATE DATABASE command again, essentially re-creating the duplicate database. This technique requires making complete copies of the datafiles of the duplicate database.

In this example, RMAN connects to the database on `host_src` as its target database and the duplicate database on `host_dup` as an auxiliary database. The example uses the DUPLICATE DATABASE command to transfer complete copies of all updated datafiles from the source database to the duplicate site.

This example makes the same assumptions about file names, locations, and the datafiles to be duplicated in ["Duplicating When the Datafiles Use Inconsistent Paths: Example"](#) on page 13-19.

As an initial setup step, the CONFIGURE command is used to set persistent new names for the datafiles. These new names are recorded in the control file and used every time you perform the duplication in the future.

Using CONFIGURE AUXNAME to Name Datafiles at the Duplicate

The script shown in [Example 13-4](#) performs the one-time setup of the names for the datafiles.

Example 13-4 Configuring Auxiliary Names for Datafiles

```
# start RMAN and then connect to the target and auxiliary databases
CONNECT TARGET /;
CONNECT CATALOG rman/cat@catdb;
CONNECT AUXILIARY SYS/oracle@dupdb;

# configure auxiliary names for the datafiles only once
CONFIGURE AUXNAME FOR DATAFILE 1 TO '/oradata1/system01.dbf';
CONFIGURE AUXNAME FOR DATAFILE 2 TO '/oradata2/undotbs01.dbf';
CONFIGURE AUXNAME FOR DATAFILE 3 TO '/oradata3/cwmlite01.dbf';
CONFIGURE AUXNAME FOR DATAFILE 4 TO '/oradata4/drsys01';
CONFIGURE AUXNAME FOR DATAFILE 5 TO '/oradata5/example01.dbf';
CONFIGURE AUXNAME FOR DATAFILE 6 TO '/oradata6/indx01.dbf';
# Do not set a newname for datafile 7, because it is in the tools tablespace,
# and in this example you are excluding tools from the duplicate database.
```

```
CONFIGURE AUXNAME FOR DATAFILE 8 TO '/oradata7/users01.dbf';
```

Synchronizing the Duplicate With the Target Using Auxnames

Run the script in [Example 13–5](#) whenever you want to synchronize the duplicate with the source database. For example, you might run the script daily or weekly.

Example 13–5 Reduplicating the Database

```
# start RMAN and then connect to the target and auxiliary databases
CONNECT TARGET /;
CONNECT CATALOG rman/cat@catdb;
CONNECT AUXILIARY SYS/oracle@dupdb;

# Create the duplicate database. Run the same command periodically
# to re-create the database, thereby keeping the duplicate
# in sync with the target.
DUPLICATE TARGET DATABASE TO dupdb
SKIP TABLESPACE tools
LOGFILE
  GROUP 1 ('/duplogs/redo01a.log',
          '/duplogs/redo01b.log') SIZE 200K REUSE,
  GROUP 2 ('/duplogs/redo02a.log',
          '/duplogs/redo02b.log') SIZE 200K REUSE;
```

Using RMAN Incremental Backups to Refresh a Standby Database

You can create an incremental backup of the target database containing changes to the database since the creation of the duplicate or the previous synchronization. You can apply the incremental backup to the standby database.

Note: This technique cannot be used to update a duplicate database.

RMAN enables you to synchronize a standby database with a primary database by creating an incremental backup at the source database that contains all changed blocks since the duplicate was created or last refreshed. You then apply the incremental backup to the standby database, which updates it with all changes.

This capability facilitates the temporary conversion of a physical standby database into a reporting database, as described in *Oracle Data Guard Concepts and Administration*. In particular, this capability makes it possible to reverse the effects of converting the standby into a reporting database. After the standby database has been used for reporting or testing, Flashback Database can reverse any changes resulting from that work, returning the database to its contents when it was still a standby. An incremental backup created with `BACKUP INCREMENTAL... FROM SCN` can be used to refresh the standby with changes at the primary since the conversion and then managed recovery can resume. The effect is to return the reporting database to its role as standby.

For more details on this scenario, see *Oracle Data Guard Concepts and Administration*.

Using BACKUP INCREMENTAL... FROM SCN

The incremental backup is created at the source database by means of the `BACKUP INCREMENTAL FROM SCN=n` form of the `BACKUP` command. For example:

```

BACKUP DEVICE TYPE SBT INCREMENTAL FROM SCN 750923 DATABASE;
BACKUP INCREMENTAL FROM SCN 750923 DATABASE;
BACKUP DEVICE TYPE DISK INCREMENTAL FROM SCN 750983 DATABASE
    FORMAT '/tmp/incr_standby_%U';

```

RMAN uses the selected SCN as the basis for this incremental backup. For all files being backed up, RMAN includes all data blocks that were changed at SCNs greater than or equal to the FROM SCN in the incremental backup.

Note: ■ RMAN does not consider the incremental backup as part of a backup strategy at the source database. The backup is not suitable for use in a normal RECOVER DATABASE operation at the source database.

- The backup sets produced by this command are written to `*/dbs` by default, even if the flash recovery area or some other backup destination is defined as the default for disk backups.
 - You must create this incremental backup on disk for it to be useful. When you move the incremental backup to the standby, you must catalog it at the standby as described in "[Step 3: Catalog the Incremental Backup Files at the Standby Database](#)" on page 13-26. Backups on tape cannot be cataloged.
-

See Also: *Oracle Database Backup and Recovery Reference* for more details on BACKUP command syntax

Refreshing a Standby Database With INCREMENTAL FROM SCN Backups: Example

This example shows the steps required to update a standby database using incremental backups. The assumption is that you have already activated the standby, performed your tests or other operations at the standby, and then used Flashback Database to undo the effects of those changes. The task here is to refresh the standby with the latest changes to the primary, so that it can resume its role as a standby database.

Step 1: Create the Incremental Backup

Create the needed incremental backup at the source database, using BACKUP with the INCREMENTAL FROM SCN clause.

Assume that the incremental backup to be used in updating the duplicate database is to be created on disk, with the filenames for backup pieces determined by the format `/tmp/incr_for_standby/bkup_%U`.

```

RMAN> BACKUP DEVICE TYPE DISK INCREMENTAL FROM SCN 750983 DATABASE
    FORMAT '/tmp/incr_for_standby/bkup_%U';

```

Step 2: Make the Incremental Backup Accessible at the Standby Database

Make the backup pieces containing the incremental backup available in some directory accessible on the system containing the standby database. For this example, assume that the destination directory is called `/standbydisk1/incrback/` and ensure that it contains nothing besides the incremental backups from Step 1.

Step 3: Catalog the Incremental Backup Files at the Standby Database

Use the RMAN `CATALOG` command to register the backup sets in the RMAN repository at the duplicate. With an RMAN client connected to the standby database and the recovery catalog (if you use one at the standby), mount the standby and run the following command:

```
RMAN> CATALOG START WITH '/standbydisk1/incrback/';
```

The backups are now available for use in recovery of the standby.

Step 4: Apply the Incremental Backup to the Standby Database

Use the RMAN `RECOVER` command with the `NOREDO` option to apply the incremental backup to the standby database. All changed blocks captured in the incremental backup are updated at the standby database, bringing it up to date with the primary database. With an RMAN client connected to the standby database, run the following command:

```
RMAN> RECOVER DATABASE NOREDO;
```

You can now resume managed recovery at the standby. Any redo logs required at the standby with changes since those contained in the incremental are automatically requested from the primary and applied.

Creating Transportable Tablespace Sets from Backup with RMAN

This chapter discusses the use of RMAN to create transportable tablespace sets from backup.

It contains the following sections:

- [About Creating Transportable Tablespace Sets from Backup with RMAN](#)
- [Transportable Tablespace Sets from Backup: Concepts](#)
- [Creating a Transportable Tablespace Set with RMAN: Procedure](#)
- [Troubleshooting RMAN TRANSPORT TABLESPACE](#)

Note: This discussion assumes that you are familiar with transportable tablespaces, as described in *Oracle Database Administrator's Guide*. The procedure described in this chapter is an alternative method of generating a transportable tablespace set, as described in the step-by-step procedure provided for transporting tablespaces.

About Creating Transportable Tablespace Sets from Backup with RMAN

The method of creating transportable tablespace sets described in *Oracle Database Administrator's Guide* requires that the tablespaces to be transported are opened read-only for the duration of the tablespace transport process.

By contrast, when RMAN creates transportable tablespace sets from backup, it does not need access to the live datafiles from the tablespaces to be transported. This eliminates the need to put the tablespace in read-only mode. As a result, database availability is improved, especially for large tablespaces, because the tablespaces to be transported can remain open for writes during the process, and because putting a tablespace into read-only mode can take a long time, depending upon current database activity.

Using RMAN for creating a transportable tablespace set also enables you to specify a target point in time, SCN or restore point during your recovery window and transport the data as it stood at that time rather than as it stands at the time of the TRANSPORT TABLESPACE. If, for example, your backup retention policy guarantees a one week recovery window, and you want to create transportable tablespaces for reporting based upon the contents of your database on the last day of the month, you can use RMAN to do so at any time during the first week of the next month.

The RMAN `TRANSPORT TABLESPACE` command is used to create transportable tablespace sets from RMAN backups.

Note: Even if RMAN is not used for backups of the database, the RMAN `TRANSPORT TABLESPACE` command can still be used to create transportable tablespace sets. However, you must have datafile copies from before the desired SCN for the transportable tablespace set, and you must prepare for `TRANSPORT TABLESPACE` by using the RMAN `CATALOG` command to record the datafile copies and archived log files needed for the operation in the RMAN repository. Once RMAN has a record of all needed backups and logs, it can perform `TRANSPORT TABLESPACE`. See *Oracle Database Backup and Recovery Basics* for details on using `CATALOG`.

See also: *Oracle Database Backup and Recovery Reference*. for reference information on the `TRANSPORT TABLESPACE` command

Transportable Tablespace Sets from Backup: Concepts

This section contains the following topics:

- [When to Use RMAN to Create Transportable Tablespace Sets](#)
- [How RMAN Creates Transportable Tablespace Sets from Backup](#)
- [Limitations of RMAN `TRANSPORT TABLESPACE` Command](#)

When to Use RMAN to Create Transportable Tablespace Sets

In any situation where you need transportable tablespaces, you can use the RMAN `TRANSPORT TABLESPACE` command to create them, as long as you have RMAN backups of your database that can be recovered to the SCN at which you need the tablespaces. Scenarios for using transportable tablespaces are described in *Oracle Database Administrator's Guide*.

However, there are two specific scenarios in which this feature is particularly useful:

- Creating transportable tablespace sets for use with a tablespace repository. For example, you have a database with some tablespaces used for quarterly reporting. You use `TRANSPORT TABLESPACE` to create transportable tablespace sets for those tablespaces quarterly for storage in a tablespace repository. Subsequently, versions of the tablespace can be requested from the repository and attached to some other database for use in the reporting process.

See also: *Oracle Streams Replication Administrator's Guide* for more details on RMAN and tablespace repositories

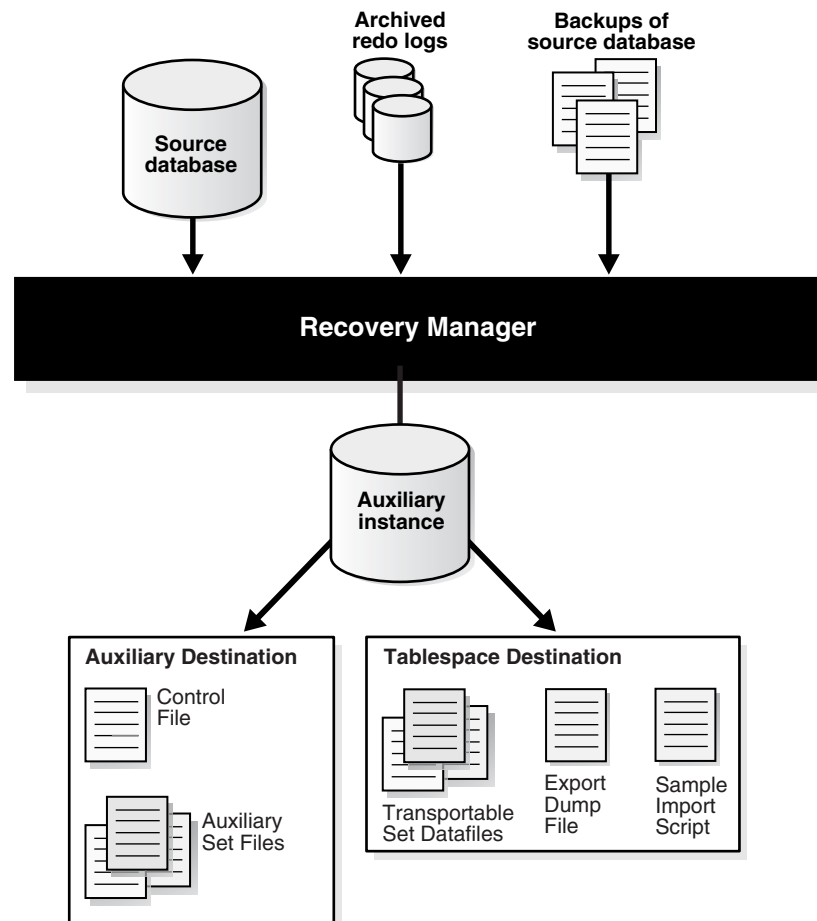
- When preparing to use Streams to keep a destination database synchronized with a source database, you must perform Streams instantiation, to bring the destination database up to a given SCN at which the two databases were known to be synchronized, before you can actually use Streams to move subsequent updates from the source to the target. Creating the transportable tablespace set from backup can be used as part of the process of Streams instantiation.

See also: *Oracle Streams Replication Administrator's Guide* for more details on RMAN and Streams instantiation

How RMAN Creates Transportable Tablespace Sets from Backup

The following diagram illustrates the process of creating transportable tablespace sets from backup using RMAN.

Figure 14–1 RMAN Transportable Tablespace From Backup: Architecture



The elements in the diagram and the process RMAN performs are described in the discussion which follows.

Note: The process by which transportable tablespace sets are created with RMAN is closely related to the process for performing TSPITR, as discussed in [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#). If you are familiar with that process, you will find this one similar.

The figure contains the following entities:

- The **Recovery Manager** client
- The **source database**, containing the tablespaces to be transported
- **Archived redo logs** and **backups** from the source database, used in restore and recovery of the tablespaces being transported
- The **auxiliary instance**, an Oracle database instance which is created by RMAN on the same host as the source database, to perform the restore and recovery of the tablespaces, and, if the process succeeds, cleaned up afterwards.
- The **auxiliary set**, which includes datafiles and other files required for the tablespace transport process but which are not themselves part of the recovery set. The auxiliary set typically includes:
 - A copy of the SYSTEM and SYSAUX tablespaces
 - Undo tablespaces and datafiles containing rollback or undo segments from the source database

The auxiliary instance has other files associated with it, such as its own control file, parameter file, and online logs, but they are not part of the auxiliary set.

- The **auxiliary destination**, a location on disk where auxiliary set files, such as the parameter file, datafiles (other than those of the tablespaces being transported), control files and online logs of the auxiliary instance can be stored while the RMAN `TRANSPORT TABLESPACE` process is running. If the transport process succeeds, these files are deleted.

Note: Specifying an auxiliary destination is optional; however, if you do not use one, you must ensure that locations are specified for all auxiliary instance files, including all datafiles and online log files, using the initialization parameters of the auxiliary instance.

Oracle recommends the use of an auxiliary destination to simplify the use of `TRANSPORT TABLESPACE`. See the rules described in "[Customize Other Auxiliary File Locations in TRANSPORT TABLESPACE](#)" on page 14-12 for details on the extra steps required if you do not use an auxiliary destination.

- The **tablespace destination**, a location on disk which (by default) contains the datafile copies and other output files when the tablespace transport command completes
- The **transportable set**, which consists of the datafiles from the tablespaces to be transported and an **export dump file** (generated using Data Pump Export) for use when plugging in the tablespaces at a destination database
- The sample **import script** (generated by RMAN) for use in plugging the tablespaces in at a target database, and an **export log** containing the output from Data Pump Export

All of these terms will be referenced throughout the remainder of this discussion.

Steps RMAN Performs in Transporting Tablespaces from Backup

RMAN creates transportable tablespaces from backup using a three-phase process:

1. In the startup phase, RMAN constructs an auxiliary database instance. First, an initialization parameter file is automatically created for the instance by RMAN, and used to start the auxiliary instance in `NOMOUNT` mode. Then, RMAN restores a

backup of the source database control file to serve as the auxiliary instance control file, and mounts that control file.

2. Once the auxiliary instance control file is mounted, RMAN restores auxiliary and transportable set datafiles from the backups of the source database, storing the auxiliary datafiles in the selected auxiliary destination and the transportable set files in the tablespace destination. RMAN then performs a `SWITCH` operation at the auxiliary instance, so that the auxiliary instance uses the restored datafiles as the datafiles for the auxiliary instance.

RMAN then performs database point-in-time recovery at the auxiliary instance. This updates auxiliary and transportable set datafiles to their contents as of the target time specified for the `TRANSPORT DATABASE` command. (If no target time is specified, complete recovery is performed.) Archived redo logs are restored from backup as necessary at the auxiliary destination (or other location), and are deleted after they are applied.

Once the recovery is complete, RMAN performs an `OPEN RESETLOGS` operation on the auxiliary database. The datafiles now reflect the tablespace contents as of the target SCN for the tablespace transport operation.

3. Recovery set tablespaces of the auxiliary instance are put into read-only mode, and Data Pump Export is invoked in transportable tablespace mode to generate the export dump file for the recovery set tablespaces.

The dump file is located in the tablespace destination location by default. To specify the location of the dump files, see "[RMAN TRANSPORT TABLESPACE: Specifying Locations for Data Pump Files](#)" on page 14-9.

At this time RMAN also generates the sample Data Pump import script for use when plugging in the transported tablespaces at a target database. The contents of this script are written to a file named `impscript.sql` in the tablespace destination. The commands for the script are also included in the RMAN output for the `TRANSPORT TABLESPACE` command.

If all of these operations are completed successfully, RMAN shuts down the auxiliary instance and deletes all files created during `TRANSPORT TABLESPACE` except for the transportable set files, the Data Pump Export log, and the sample import script.

Note:

- The recovery set datafiles are not automatically converted to the endian format of the destination database by the `TRANSPORT TABLESPACE` process. If necessary, use the `RMAN CONVERT` command, described in [Chapter 15, "RMAN Cross-Platform Transportable Databases and Tablespaces"](#), to convert the datafiles to the endian format of the destination database after creating the transportable set.
 - The sample import script assumes that the files used in plugging the transportable tablespaces at the destination database are stored in the same locations where they were created by the `TRANSPORT TABLESPACE` operation. If files have been moved to new disk locations before being plugged in, you must update the sample script with the new locations of the files before using the script to plug in the transported tablespaces.
-
-

Limitations of RMAN TRANSPORT TABLESPACE Command

The limitations on creating transportable tablespace sets described in *Oracle Database Administrator's Guide* apply to transporting tablespaces from backup, with the exception of the requirement to make the tablespaces read-only.

There are also limitations specific to creating a transportable tablespace set using RMAN:

- You must have a backup of all needed tablespaces (including those in the auxiliary set) and archived redo log files available for use by RMAN that can be recovered to the target point in time for the `TRANSPORT TABLESPACE` operation.

Note: If RMAN is not part of the backup strategy for your database, you can still use `RMAN TRANSPORT TABLESPACE`, as long as the needed datafile copies and archived redo logs are available on disk. Use the `RMAN CATALOG` command to record the datafile copies and archived logs in the RMAN repository. You can then use `TRANSPORT TABLESPACE`. See *Oracle Database Backup and Recovery Basics* for details on using `CATALOG`.

You also have the option of using RMAN to back up your database specifically to create backups for use in creating a transportable tablespace set from backup.

- Because the RMAN process for creating transportable tablespaces from backup uses the Data Pump Export and Import utilities, you cannot use this process if the tablespaces to be transported use XMLTypes. In such a case, you must use the process documented in *Oracle Database Administrator's Guide*.
- Because RMAN creates the automatic auxiliary instance used for restore and recovery on the same host as the source instance, there is some performance overhead during the operation of the `TRANSPORT TABLESPACE` command.
- If you drop a tablespace, then you cannot later use `TRANSPORT TABLESPACE` to include that tablespace in a transportable tablespace set, even if the SCN for `TRANSPORT TABLESPACE` is earlier than the SCN at which the table was dropped.
- If you rename a tablespace, you cannot use `TRANSPORT TABLESPACE` to create a transportable tablespace set as of a point in time before the tablespace was renamed. (RMAN has no information about the previous name of the tablespace.)
- You cannot `TRANSPORT` tables without their associated constraints, or constraints without their associated tables.
- Neither the transportable set nor the auxiliary set datafiles can contain any of the following:
 - Replicated master tables
 - Partial tables
 - Tables with VARRAY columns, nested tables, or external files
 - Snapshot logs and snapshot tables
 - Tablespaces containing undo or rollback segments
 - Tablespaces that contain objects owned by SYS, including rollback segments

If you are performing `TRANSPORT TABLESPACE` without a recovery catalog, the following additional restrictions apply:

- If creating a transportable set with tablespace contents as of a point in time in the past, then the set of tablespaces with undo segments at the time `TRANSPORT TABLESPACE` is executed must be the same as the set of tablespaces with undo segments at the time selected for transport.

Tablespaces including undo segments as of the target SCN for `TRANSPORT TABLESPACE` must be part of the auxiliary set. Unlike the recovery catalog, the RMAN repository in the control file only contains a record of tablespaces that include undo segments at the current time. If the set of tablespaces with undo segments was different at the target time, then `TRANSPORT TABLESPACE` fails.

- If the database has re-used the control file records for the RMAN repository that contained information about backups required for the `TRANSPORT TABLESPACE` process, then the process fails because RMAN cannot locate the required backups. You may be able to use `CATALOG` to add the needed backups to the RMAN repository if they are still available, but if the database is already overwriting control file records you may lose records of other needed backups.

Creating a Transportable Tablespace Set with RMAN: Procedure

The focus in this discussion is on presenting the basics of planning for and performing `TRANSPORT TABLESPACE` operations with RMAN in the most automated case, where RMAN manages the auxiliary instance for you, and then how to customize various aspects of the process. This discussion contains the following sections:

- [Using RMAN `TRANSPORT TABLESPACE`: Basic Scenario](#)
- [RMAN `TRANSPORT TABLESPACE` with `UNTIL` Time or SCN](#)
- [RMAN `TRANSPORT TABLESPACE`: Specifying Locations for Data Pump Files](#)
- [RMAN `TRANSPORT TABLESPACE` with Customized Initialization Parameters](#)
- [Customize Shared Pool Size in RMAN `TRANSPORT TABLESPACE`](#)
- [Customize Auxiliary Control File Location in `TRANSPORT TABLESPACE`](#)
- [Customize Other Auxiliary File Locations in `TRANSPORT TABLESPACE`](#)

Using RMAN `TRANSPORT TABLESPACE`: Basic Scenario

This section describes the use of `TRANSPORT TABLESPACE` in the most automated case.

Note: The process described here is only one part of the process of transporting tablespaces. Before you use `TRANSPORT TABLESPACE`, you must meet the requirements described in *Oracle Database Administrator's Guide*:

- Confirm that tablespace transport is supported between your source and destination platforms
 - Identify a self-contained set of tablespaces to include in the transportable set
-
-

In this scenario, the `AUXILIARY DESTINATION` clause is specified, which causes RMAN to use default values that work for most cases in managing the auxiliary

instance. Only required options are specified. Oracle recommends that you use an auxiliary destination with `TRANSPORT TABLESPACE` to simplify management of auxiliary instance files.

To use RMAN TRANSPORT TABLESPACE:

Start the RMAN client, connecting to the source database and, if used, the recovery catalog. Then enter the `TRANSPORT TABLESPACE` command, specifying the required arguments. For example, to transport the tablespaces `tbs_2` and `tbs_3`, use the `TRANSPORT TABLESPACE` command as follows:

```
transport tablespace tbs_2, tbs_3
  tablespace destination '/disk1/transportdest'
  auxiliary destination '/disk1/auxdest'
;
```

When the `TRANSPORT TABLESPACE` command completes, the following outputs result:

- The transportable set datafiles are left in the location `/disk1/transportdest` with their original names.
- The Data Pump Export dump file for the transportable set is named `dmpfile.dmp`, the export log is named `explog.log` and the sample import script is named `impscript.sql`. All are created in the tablespace destination `/disk1/transportdest`.

Note: If there is already a file under the name of the export dump file in the tablespace destination, then `TRANSPORT TABLESPACE` fails when it calls Data Pump Export. If repeating a previous `TRANSPORT TABLESPACE` operation, make sure you delete the previous output files, including the export dump file.

- The auxiliary set files are removed from `/disk1/auxdest`.

You can now return to the process for transporting tablespaces described in *Oracle Database Administrator's Guide*.

RMAN TRANSPORT TABLESPACE with UNTIL Time or SCN

If you specify a target time with the `TRANSPORT TABLESPACE` command, then during the tablespace transport operation, RMAN will restore the tablespace at the auxiliary instance using backups from prior to the target time and perform point-in-time recovery on the auxiliary database to the specified target time. (Backups and redo logs needed for this point-in-time recovery must be available.)

The target time can be specified using an SCN (in the current incarnation or its ancestors), log sequence number, restore point or time expression. For example:

```
TRANSPORT TABLESPACE tbs_2
  TABLESPACE DESTINATION '/disk1/transportdest'
  AUXILIARY DESTINATION '/disk1/auxdest'
  UNTIL SCN 11379;

TRANSPORT TABLESPACE tbs_2
  TABLESPACE DESTINATION '/disk1/transportdest'
  AUXILIARY DESTINATION '/disk1/auxdest'
  UNTIL RESTORE POINT 'before_upgrade';
```



```

TRANSPORT TABLESPACE tbs_2
  TABLESPACE DESTINATION '/disk1/transportdest'
  AUXILIARY DESTINATION '/disk1/auxdest'
  UNTIL TIME 'SYSDATE-1';

```

RMAN TRANSPORT TABLESPACE: Specifying Locations for Data Pump Files

You can change the names of the Data Pump Export dump file for the transportable set, the sample import script for use at the target database, and the log file generated by Data Pump Export, and the directory to which they are written.

By default, these files are stored in the tablespace destination and named as follows:

- The Data Pump Export dump file is named `dmpfile.dmp`
- The export log file is named `explog.log`
- The sample import script is called `impscript.sql`

You can place the dump file and the export log in a different directory by using the `DATAPUMP DIRECTORY` clause, passing in the name of a database directory object.

Note: The database directory object used by `DATAPUMP DIRECTORY` is not the directory path of an actual file system directory. The value passed corresponds to the `DIRECTORY` command line argument of Data Pump Export. See *Oracle Database Utilities* for more details on the use of directory objects with Data Pump Export.

These files can be renamed using the `DUMP FILE`, `EXPORT LOG` and `IMPORT SCRIPT` clauses of `TRANSPORT TABLESPACE`.

Note:

- The filenames cannot contain full file paths with directory names. If the `DUMP FILE` or `EXPORT LOG` filenames specify file paths, then `TRANSPORT TABLESPACE` fails when it attempts to generate the export dump. Use the `DATAPUMP DIRECTORY` clause to specify a database directory object that identifies a location for the outputs of Data Pump Export.
-

The following example illustrates the use of `TRANSPORT TABLESPACE` with the `DATAPUMP DIRECTORY`, `DUMP FILE`, `EXPORT LOG` and `IMPORT SCRIPT` filenames specified.

Assume that a database directory object has been created as follows, for use with Data Pump Export:

```
SQL> CREATE OR REPLACE DIRECTORY mypumpdir as '/datapumpdest';
```

The following `TRANSPORT TABLESPACE` command illustrates the use of the optional arguments that control file output locations:

```

TRANSPORT TABLESPACE tbs_2
  TABLESPACE DESTINATION '/transportdest'
  AUXILIARY DESTINATION '/auxdest'
  DATAPUMP DIRECTORY mypumpdir
  DUMP FILE 'mydumpfile.dmp'

```

```
IMPORT SCRIPT 'myimportscript.sql'
EXPORT LOG 'myexportlog.log';
```

After a successful run, the auxiliary destination is cleaned up, the Data Pump Export Dump file and the export log are located in the directory referenced by `DATAPUMP DIRECTORY (/datapumpdest/mydumpfile.dmp and /datapumpdest/myexportlog.log)` and the transportable set datafiles are stored in `/transportdest`.

RMAN TRANSPORT TABLESPACE with Customized Initialization Parameters

When RMAN creates the auxiliary instance, it creates an initialization parameter file with default parameters that should work for nearly all `TRANSPORT TABLESPACE` cases, especially if you provide an `AUXILIARY DESTINATION` argument to `TRANSPORT TABLESPACE`.

RMAN `TRANSPORT TABLESPACE` can also use an auxiliary instance parameter file, which can contain values for additional initialization parameters, and values that override the values of those parameters defined in the default initialization parameter file.

Note: The auxiliary instance parameter file is not expected to be a complete initialization parameter file for the auxiliary instance. Any parameters specified are added to or override the default parameters for the auxiliary instance. It is not necessary to specify parameters in the initialization file that you do not intend to override.

Reasons for using an auxiliary instance parameter file include:

- To manage locations for auxiliary instance datafiles (if, for instance, you do not want all auxiliary instance datafiles stored in the same location on disk, but you do not want to specify the location of every file individually).
- To specify control naming for online redo logs using `LOG_FILE_NAME_CONVERT`
- To increase `SHARED_POOL_SIZE` if needed for Data Pump Export.

RMAN TRANSPORT TABLESPACE Default Auxiliary Instance Parameters

RMAN defines the following basic initialization parameters for the automatic auxiliary instance:

- `DB_NAME` - Same as `DB_NAME` of the source database
- `COMPATIBLE` - Same as the compatible setting of the source database
- `DB_UNIQUE_NAME` - Generated, based on the `DB_NAME`, to be unique
- `DB_BLOCK_SIZE` - Same as the `DB_BLOCK_SIZE` of the source database
- `DB_FILES` - Set to same value as `DB_FILES` for the source database
- `SHARED_POOL_SIZE` - Set to 110MB, because Data Pump Export can require this much space or more.
- `LARGE_POOL_SIZE` - Set to 1MB

If the `AUXILIARY DESTINATION` argument to `TRANSPORT TABLESPACE` is used, RMAN also defines:

- `DB_CREATE_FILE_DEST` - Set to the auxiliary destination

- `CONTROL_FILES` - A generated filename in the auxiliary destination

When an auxiliary destination is specified, RMAN uses these two parameters in creating the auxiliary instance online logs and control files in the auxiliary destination.

Note: Overriding one of these basic initialization parameters with an inappropriate value in the auxiliary instance parameter file can cause `TRANSPORT TABLESPACE` to fail. If you encounter a problem, try returning to the default value.

Location of the Auxiliary Instance Parameter File

By default, RMAN looks for the auxiliary instance parameter file at an operating system-dependent location on the host running the RMAN client (which may not be the host running the auxiliary instance). For Unix, the location is: `?:rdbms/admin/params_auxint.ora`, where '?' stands for `ORACLE_HOME` on the host running RMAN. (If no file is found in the default location, RMAN does not generate an error.)

Note: If you want to use the default initialization parameters for the auxiliary instance, check whether an auxiliary instance parameter file exists before running `TRANSPORT TABLESPACE`. The same location is used for an auxiliary instance parameter file in TSPITR, so it is possible that a file remains from a previous TSPITR or `TRANSPORT TABLESPACE` operation. In such a case, you may see unexpected behavior, depending upon the parameters specified.

You can use the RMAN `SET AUXILIARY INSTANCE PARAMETER FILE` command in a `RUN` block before `TRANSPORT TABLESPACE` to specify a different location for the auxiliary instance parameter file. (Note that, as with the default location of the auxiliary instance parameter file, the path specified when using `SET AUXILIARY INSTANCE PARAMETER FILE` is a client-side path.)

Using an Auxiliary Instance Parameter File: Example

For example, you can create a file `/tmp/auxinstparams.ora` on the host running the RMAN client, containing the following initialization parameters:

```
SHARED_POOL_SIZE=150M;
```

You can then use the parameter file with `TRANSPORT TABLESPACE` as in this example:

```
RUN {
SET AUXILIARY INSTANCE PARAMETER FILE TO '/tmp/auxinstparams.ora';
TRANSPORT TABLESPACE tbs_2
    TABLESPACE DESTINATION '/disk1/transportdest'
    AUXILIARY DESTINATION '/disk1/auxdest';
}
```

The `SHARED_POOL_SIZE` parameter specified in `/tmp/auxinstparams.ora` overrides the default value used for `SHARED_POOL_SIZE` in `TRANSPORT TABLESPACE` when RMAN sets up the auxiliary instance.

Customize Shared Pool Size in RMAN TRANSPORT TABLESPACE

If the default shared pool size of 100MB is not sufficient for Data Pump Export to run successfully, then specify a larger value. Use an auxiliary instance parameter file as described in "[RMAN TRANSPORT TABLESPACE with Customized Initialization Parameters](#)" on page 14-10. Include the `SHARED_POOL_SIZE` initialization parameter in your auxiliary instance parameter file, listing the control file locations desired.

For more on the use of the `SHARED_POOL_SIZE` initialization parameter, see *Oracle Database Reference*.

Customize Auxiliary Control File Location in TRANSPORT TABLESPACE

By default, one control file is created for the auxiliary instance, in an operating system dependent location. (In Unix, the default location is `?/dbs/cntrl_@.dbf`, where '?' stands for `ORACLE_HOME` and '@' stands for `ORACLE_SID`. For an automatic auxiliary instance, `ORACLE_SID` is randomly generated by RMAN).

To change the location of the auxiliary instance control file, use an auxiliary instance parameter file as described in "[RMAN TRANSPORT TABLESPACE with Customized Initialization Parameters](#)" on page 14-10. Include the `CONTROL_FILES` initialization parameter in your auxiliary instance parameter file, listing the control file locations desired.

For more on the use of the `CONTROL_FILES` initialization parameter, see *Oracle Database Reference*.

Customize Other Auxiliary File Locations in TRANSPORT TABLESPACE

Several different rules affect the location of auxiliary instance files such as datafiles and online logs, created during `TRANSPORT TABLESPACE`. The simplest method is to use `AUXILIARY DESTINATION` and let RMAN manage all file locations automatically. However, if you need to relocate some or all auxiliary instance files, the options for specifying locations for these files are described here in order of precedence, so that, if several are used, then for each file, the first rule in the list that applies determines the filename.

- [Transport Tablespace with SET NEWNAME for Auxiliary Datafiles](#)
- [Transport Tablespace with CONFIGURE AUXNAME for Auxiliary Datafiles](#)
- [Transport Tablespace with AUXILIARY DESTINATION Parameter](#)
- [Transport Tablespace and Naming Auxiliary Files with Initialization Parameters](#)

Note: If you do not use `AUXILIARY DESTINATION`, then you must use `LOG_FILE_NAME_CONVERT` to specify the location of the online redo log files for the auxiliary instance.

Neither `SET NEWNAME` nor `CONFIGURE AUXNAME` can affect the location of the auxiliary instance online logs, so if you do not use `AUXILIARY DESTINATION` or `LOG_FILE_NAME_CONVERT` RMAN has no information about where to create the online logs.

Transport Tablespace with SET NEWNAME for Auxiliary Datafiles

You can use the `SET NEWNAME` command in a `RUN` block to specify filenames for use in `TRANSPORT TABLESPACE`, as in this example:

```
RUN {
```

```

SET NEWNAME FOR DATAFILE '/oracle/dbs/tbs_12.f'
  TO '/bigdrive/auxdest/tbs_12.f';
SET NEWNAME FOR DATAFILE '/oracle/dbs/tbs_11.f'
  TO '/bigdrive/auxdest/tbs_11.f';
TRANSPORT TABLESPACE tbs_2
  TABLESPACE DESTINATION '/disk1/transportdest'
  AUXILIARY DESTINATION '/disk1/auxdest';
}

```

The SET NEWNAME commands cause these auxiliary instance datafiles to be restored to the locations named instead of /disk1/auxdest.

Note: SET NEWNAME is best used with one-time operations. If you expect to create transportable tablespaces from backup regularly for a particular set of tablespaces, consider using CONFIGURE AUXNAME instead of SET NEWNAME in order to make persistent settings for the location of the auxiliary instance datafiles.

Transport Tablespace with CONFIGURE AUXNAME for Auxiliary Datafiles

You can use the CONFIGURE AUXNAME command to specify persistent locations for recovery set or auxiliary set datafiles. Each datafile for which a CONFIGURE AUXNAME command has been used is restored from backup to the specified location prior to recovery. Recovery set datafiles are left at the specified location after TRANSPORT TABLESPACE is complete, instead of being stored in the tablespace destination. Auxiliary set datafiles are deleted when the operation is complete, unless the operation failed.

For example, consider a database where tablespace tbs_12, containing datafile tbs_12.f, is part of the auxiliary set required when creating a transportable set for a tablespace tbs_11. This CONFIGURE AUXNAME statement sets a persistent non-default location for the auxiliary set datafile '/oracle/dbs/tbs_12.f' during TRANSPORT TABLESPACE:

```
CONFIGURE AUXNAME FOR '/oracle/dbs/tbs_12.f' TO '/disk1/auxdest/tbs_12.f';
```

In the future, you execute the following TRANSPORT TABLESPACE command:

```
TRANSPORT TABLESPACE tbs_1 AUXILIARY DESTINATION '/myauxdest';
```

In such a case, the auxiliary set copy of datafile '/oracle/dbs/tbs_12.f' is restored at the location /disk1/auxdest/tbs_12.f instead of being stored in the location specified by AUXILIARY DESTINATION.

Note: You can view any current CONFIGURE AUXNAME settings using the SHOW AUXNAME command, described in *Oracle Database Backup and Recovery Reference*.

Transport Tablespace with AUXILIARY DESTINATION Parameter

If you use an AUXILIARY DESTINATION argument with TRANSPORT TABLESPACE, as in the example in ["Using RMAN TRANSPORT TABLESPACE: Basic Scenario"](#) on page 14-7, then any auxiliary set file which is not moved to another location using SET NEWNAME or CONFIGURE AUXNAME commands is stored in the auxiliary destination during the TRANSPORT TABLESPACE operation.

Transport Tablespace and Naming Auxiliary Files with Initialization Parameters

You can use the `LOG_FILE_NAME_CONVERT` and `DB_FILE_NAME_CONVERT` initialization parameters in an auxiliary instance parameter file to control the naming for online logs and other database files at the auxiliary instance. To do so, use an auxiliary instance parameter file, and specify the initialization parameters `LOG_FILE_NAME_CONVERT`, which controls naming for online log files, and `DB_FILE_NAME_CONVERT`, which controls naming for other database files at the auxiliary instance.

If no `AUXILIARY DESTINATION` clause is provided to the `TRANSPORT TABLESPACE` command, then these parameters determine the location of any files for which no `CONFIGURE AUXNAME` or `SET NEWNAME` value was provided.

Note:

- If you do not use an `AUXILIARY DESTINATION` clause, then you must use an auxiliary instance parameter file and specify `LOG_FILE_NAME_CONVERT` to generate names for the online logs of the auxiliary instance. Otherwise, RMAN cannot determine a location for the online logs for the automatic auxiliary instance.
 - Neither `LOG_FILE_NAME_CONVERT` nor `DB_FILE_NAME_CONVERT` can be used to generate new Oracle Managed Files (OMF) filenames for files at the auxiliary instance when the original files are stored in Oracle Managed Files. The database must be allowed to manage the generation of unique filenames in each OMF destination. Therefore, you must use an `AUXILIARY DESTINATION` clause to control the location of the online logs, and you must either use `AUXILIARY DESTINATION` clause, `SET NEWNAME` or `CONFIGURE AUXNAME` to specify the location for datafiles.
-
-

See "[RMAN TRANSPORT TABLESPACE with Customized Initialization Parameters](#)" on page 14-10 for details on using `SET AUXILIARY INSTANCE PARAMETER FILE` with `TRANSPORT TABLESPACE`, and *Oracle Database Reference* for more details on the `LOG_FILE_NAME_CONVERT` and `DB_FILE_NAME_CONVERT` initialization parameters.

Troubleshooting RMAN TRANSPORT TABLESPACE

When the `RMAN TRANSPORT TABLESPACE` command fails, the failed auxiliary instance files are left intact in the auxiliary instance destination for troubleshooting.

Common causes for problems during `TRANSPORT TABLESPACE` include the following:

- [Troubleshooting RMAN TRANSPORT TABLESPACE: Insufficient Shared Pool](#)
- [Troubleshooting RMAN TRANSPORT TABLESPACE: Filename Conflicts](#)

Troubleshooting RMAN TRANSPORT TABLESPACE: Insufficient Shared Pool

If you see an error related to shared pool size, create an auxiliary instance parameter file with a larger value for `SHARED_POOL_SIZE`, as described in "[Customize Shared Pool Size in RMAN TRANSPORT TABLESPACE](#)" on page 14-12 .

Troubleshooting RMAN TRANSPORT TABLESPACE: Filename Conflicts

If your `SET NEWNAME`, `CONFIGURE AUXNAME` and `DB_FILE_NAME_CONVERT` cause multiple files in the auxiliary or recovery sets to have the same name, RMAN will report an error during the `TRANSPORT TABLESPACE` command. To correct the problem, use different values for these parameters to ensure that duplicate file names are not created.

RMAN Cross-Platform Transportable Databases and Tablespaces

Oracle supports transportable tablespaces, which are used to move the contents of individual tablespaces between databases. It also supports a transportable database feature, which can be used to recreate an entire database from one platform on another platform as long as the platforms have the same endian order.

This chapter discusses the use of RMAN in transporting tablespaces and entire databases between disparate platforms. This chapter contains the following sections:

- **Cross-Platform Transportable Tablespace: CONVERT DATAFILE or TABLESPACE**

A description of how to use RMAN's `CONVERT DATAFILE` and `CONVERT TABLESPACE` commands when moving transportable tablespaces across platforms with different endian formats

- **Cross-Platform Transportable Database: RMAN CONVERT DATABASE**

A description of how to use RMAN's `CONVERT DATABASE` command to simplify transporting an entire database to a different platform, which must have the same endian format

Cross-Platform Transportable Tablespace: CONVERT DATAFILE or TABLESPACE

Transportable tablespaces allow you to copy an entire tablespace between Oracle databases. A full discussion of transportable tablespaces, their uses, and the different techniques for creating and using them can be found in the *Oracle Database Administrator's Guide*.

Cross-platform transportable tablespaces are transportable tablespaces in which the destination hosts are of a different platform from the source host. Typical uses of cross-platform transportable tablespaces include:

- Publishing structured data as transportable tablespaces for distribution to customers, who can convert the tablespaces for integration into their existing databases regardless of platform
- Moving data from a large data warehouse server to data marts on smaller computers such as Linux-based workstations or servers
- Sharing read-only tablespaces across a heterogeneous cluster where all hosts share the same endian format

When transporting tablespaces between databases where the endian format of the source platform is different from that of the destination platform, the endian format of

the datafiles in the transportable tablespace set must be converted to match the destination platform. This conversion can be performed using the RMAN CONVERT TABLESPACE command (when converting on the source host) or CONVERT DATAFILE command (when converting on the destination host).

Note: Using the RMAN CONVERT command to convert the datafiles of a transportable tablespace set for use on platforms with different endian formats is only one step in using cross-platform transportable tablespaces. Read the discussion of transportable tablespaces in *Oracle Database Administrator's Guide* in its entirety before attempting to use transportable tablespaces or the procedures in this section.

CONVERT does not perform in-place conversion of datafiles. It produces output files in the correct format for use on the destination platform, but does not alter the contents of the source datafiles.

Differences between the conversion process on the source and destination platforms are described in the following discussion. The CONVERT TABLESPACE command must be used on the source platform, while the CONVERT DATAFILE command is used on the destination platform.

This discussion contains the following sections:

- [Using CONVERT TABLESPACE... TO PLATFORM on the Source Platform](#)
- [Using CONVERT DATAFILE... FROM PLATFORM on the Destination Host](#)

Using CONVERT TABLESPACE... TO PLATFORM on the Source Platform

The CONVERT TABLESPACE... TO PLATFORM command is used to convert tablespaces on the source platform to the format of a destination platform. Invoke CONVERT TABLESPACE, specifying the names of one or more tablespaces in the target database to convert and the destination platform for the conversion, as shown in this example:

```
RMAN> CONVERT TABLESPACE ts_1, ts_2...  
      TO PLATFORM 'platform_name';
```

The TO PLATFORM clause is mandatory with CONVERT TABLESPACE. Supported values for *platform_name* can be found in `V$TRANSPORTABLE_PLATFORM`. Note that you can only convert entire tablespaces, on the source platform, not individual datafiles.

Optional parameters for CONVERT TABLESPACE ... TO PLATFORM include:

- PARALLELISM *n*

Used to specify that *n* server sessions should perform the work of conversion in parallel to improve performance. Each datafile is assigned to a single server session for conversion, that is, you cannot improve performance on converting a single datafile by assigning a greater degree of parallelism.

Note: The optimal degree of parallelism to use is a function of the number of effective disk heads available for reading and writing. Setting the degree of parallelism too high for a given number of spindles can actually increase the time required for conversion. It is never useful to specify a degree of parallelism greater than the number of datafiles to be processed.

- *fileNameConversionSpec*
A series of patterns, specified using the `DB_FILE_NAME_CONVERT` argument, used to generate new file names for the converted datafiles, based on the input datafile names.
- `FORMAT formatSpec`
Provides a format used as a template to generate new, unique filenames for the converted datafiles. If no `FORMAT` is specified, then RMAN uses a platform-dependent destination and format.

The full semantics of these parameters are described in the reference entry for the `CONVERT` command in the *Oracle Database Backup and Recovery Reference*.

Rules for Renaming Files with CONVERT TABLESPACE

You can use the `FORMAT` and *fileNameConversionSpec* arguments to control the names of the files generated by the `CONVERT` command, or leave the renaming of converted files to default rules provided by RMAN. The following list describes the order of precedence among these rules:

- Files that match any patterns provided in a *fileNameConversionSpec* clause are renamed based upon that pattern.
- If you specify a `FORMAT` clause, then any file not renamed based upon a *fileNameConversionSpec* pattern are renamed according to the specified *formatSpec*.
- Any file not renamed by *fileNameConversionSpec* or `FORMAT` is assigned a default location based upon the following rules:
 - If the channel used for output has a default `CONFIGURE... FORMAT` setting, that setting is used to generate output file names
 - If a flash recovery area is configured, the converted datafiles are placed in the flash recovery area (though they are not usable backups).
 - Otherwise a platform-specific default `FORMAT` (which includes a `%U` for generating a unique filename) is used.

Note:

- These rules are the same as the rules that determine default locations for `BACKUP AS COPY` backups.
- The *fileNameConversionSpec* method cannot be used to generate output filenames for `CONVERT` when the source files have Oracle Managed Files (OMF) file names (such as `/private/boston/datafile/o1_mf_system_ab12554_.dbf` for a file system using OMF or `+DISK/boston/datafile/system.256.4543080` for ASM) and the destination is an OMF destination.

For complete details on rules governing file naming, see the reference entry for `BACKUP AS COPY` in *Oracle Database Backup and Recovery Reference*.

Converting Tablespaces on the Source Host: Example

Suppose you need to transport tablespaces `finance` and `hr` from a source database running on a Sun Solaris host to a destination database running on a Linux PC. You

plan to store the converted datafiles in the temporary directory `/tmp/transport_linux/` on the source host.

The example assumes that you have carried out the following steps in preparation for the tablespace transport:

- You have set the tablespaces to be transported to be read-only.
- You have looked up the name for the destination platform in `V$TRANSPORTABLE_PLATFORM`.

The database has a list of its own internal names for each platform supporting cross-platform data transport. You may need the exact name of the source or target platform as a parameter to the `CONVERT` command. Query `V$TRANSPORTABLE_PLATFORM` to get the platform name from SQL*Plus as follows:

```
SQL> SELECT PLATFORM_ID, PLATFORM_NAME, ENDIAN_FORMAT
       FROM V$TRANSPORTABLE_PLATFORM
       WHERE UPPER(PLATFORM_NAME) LIKE 'LINUX';
```

The `PLATFORM_NAME` for Linux on a PC is `'Linux IA (32-bit)'`.

Now use `RMAN` to convert the datafiles into the endian format of the destination host. In this example, the `FORMAT` argument controls the name and location of the converted datafiles.

```
% rman TARGET /
RMAN> CONVERT TABLESPACE finance,hr
      TO PLATFORM 'Linux IA (32-bit)'
      FORMAT='/tmp/transport_linux/%U';
```

The result is a set of converted datafiles in the `/tmp/transport_linux/` directory, with data in the correct endian format for the Linux IA (32-bit) platform.

From this point, follow the rest of the general outline for creating a transportable tablespace set:

- Use the export utility to create the export dump file
- Move the converted datafiles from `/tmp/transport_linux/` and the export dump file from the source host to the desired directories on the destination host
- Plug the tablespace into the new database with the Import utility.

Using CONVERT DATAFILE... FROM PLATFORM on the Destination Host

The `CONVERT DATAFILE... FROM PLATFORM` command is used on the destination platform to convert tablespaces from the format of a source platform into a format that can be plugged into a database on the destination platform.

Note: You cannot use `CONVERT TABLESPACE` on the destination platform. Until the datafiles are transported into the destination database, the datafiles are not associated with a tablespace name in the database, so `RMAN` cannot translate the tablespace name into a list of datafiles. Therefore, you must use `CONVERT DATAFILE` and identify the datafiles by filename.

`CONVERT DATAFILE` takes as arguments the names of one or more datafiles to convert, and the name of the source platform for the datafiles, as shown in this example:

```
RMAN> CONVERT DATAFILE datafile_1, datafile_2...
```

```
FROM PLATFORM 'platform_name';
```

The destination platform is the platform on which the destination database is running.

The value provided for FROM PLATFORM must match the actual format of the datafiles to be converted, or RMAN returns an error. Supported values for *platform_name* can be found in V\$TRANSPORTABLE_PLATFORM.

Note: The FROM PLATFORM clause is optional with CONVERT DATAFILE. If omitted, however, it is assumed that the datafiles to be converted are already in the format of the destination platform. The effect of CONVERT DATAFILE without FROM PLATFORM is to copy datafiles from one location to another without changing their format, and without recording the copies created in the RMAN repository as backups of the copied datafiles.

The primary use of CONVERT DATAFILE without FROM PLATFORM is in copying datafiles into and out of Automated Storage Management (ASM) disk groups. Operating system-level tools cannot be used to read or write files into ASM, but RMAN provides the required functionality. See ["Using RMAN CONVERT to Copy Files Between ASM and Non-ASM Storage"](#) on page 15-16 for more details on this procedure.

The PARALLELISM, FORMAT, and *fileNameConversionSpec* arguments, described in ["Using CONVERT TABLESPACE... TO PLATFORM on the Source Platform"](#) on page 15-2, behave the same on the source and destination platforms.

Converting Datafiles on the Destination Platform: Example

Suppose you need to transport tablespaces *finance* (datafiles 'fin/fin01.dbf' and 'fin/fin02.dbf') and *hr* (datafiles 'hr/hr01.dbf' and 'hr/hr02.dbf') from a source database running on a Sun Solaris host to a destination database running on a Linux PC host. You plan to perform conversion on the destination host.

When the datafiles are plugged into the destination database, you plan to store them in */orahome/dbs* and preserve the current directory structure (that is, datafiles for the *hr* tablespace go in the */orahome/dbs/hr* subdirectory, and datafiles for the *finance* tablespace go in the */orahome/dbs/fin* directory).

The example assumes that you have carried out the following steps in preparation for converting the tablespaces and plugging them in at the destination:

- You have set the source tablespaces to be transported to be read-only, used the Original Export utility to create the export dump file (named, in our example, *expdat.dmp*), and copied *expdat.dmp* and the unconverted datafiles of the transportable tablespace set to the destination host, in the */tmp/transport_solaris/* directory. You have preserved the subdirectory structure from the files' original location, that is, the datafiles are stored as:
 - */tmp/transport_solaris/fin/fin01.dbf*
 - */tmp/transport_solaris/fin/fin02.dbf*
 - */tmp/transport_solaris/hr/hr01.dbf*
 - */tmp/transport_solaris/hr/hr02.dbf*

Now use RMAN's CONVERT command to convert the datafiles to be transported to the destination host's format and store the converted datafiles in */orahome/dbs*.

```
% rman TARGET /
RMAN> CONVERT DATAFILE
    '/tmp/transport_solaris/fin/fin01.dbf',
    '/tmp/transport_solaris/fin/fin02.dbf',
    '/tmp/transport_solaris/hr/hr01.dbf',
    '/tmp/transport_solaris/hr/hr02.dbf'
DB_FILE_NAME_CONVERT
    '/tmp/transport_solaris/fin','/orahome/dbs/fin',
    '/tmp/transport_solaris/hr','/orahome/dbs/hr'
```

Note the following:

- The `FORMAT` argument controls the name and location of the converted datafiles.
- When converting on the destination platform, you must specify the source platform using the `FROM PLATFORM` argument. Otherwise, RMAN will assume that the source platform is the same as the platform of the host performing the conversion.

The result is a set of converted datafiles in the `/orahome/dbs/` directory, named thus:

- `/orahome/dbs/fin/fin01.dbf`
- `/orahome/dbs/fin/fin02.dbf`
- `/orahome/dbs/hr/hr01.dbf`
- `/orahome/dbs/hr/hr02.dbf`

From this point, follow the rest of the general outline for tablespace transport. Plug the converted tablespaces into the new database with the import utility, and make the tablespaces read-write if applicable.

Restrictions on CONVERT TABLESPACE and CONVERT DATAFILE

Note the following restrictions on `CONVERT TABLESPACE` and `CONVERT DATAFILE`:

- Both source and destination databases must be running with the `COMPATIBLE` initialization parameter set to 10.0 or higher.
- Not all combinations of source and destination platforms are supported. To determine whether your source and destination platforms are supported, query `V$TRANSPORTABLE_PLATFORM`. If both the source and destination platforms are listed in this view, then `CONVERT` can be used to prepare datafiles from one platform for use on the other.
- A tablespace must be made read-write at least once in release 10g before it can be transported to another platform using `CONVERT`. Hence, any read-only tablespaces (or previously transported tablespaces) from a previous release must first be made read-write before they can be transported to another platform.
- RMAN does not process user datatypes that require endian conversions.
- Prior to release 10g, CLOBs were created with a variable-width character set and stored in an endian-dependent format. The `CONVERT` command does not perform conversions on these CLOBs. Instead, RMAN captures the endian format of each LOB column and propagates it to the target database. Subsequent reads of this data by the SQL layer interpret the data correctly based upon either endian format, and write the data out in an endian-independent format if the tablespace is writeable.

CLOBs created in Oracle Database Release 10g are stored in character set AL16UTF16, which is platform-independent.

Cross-Platform Transportable Database: RMAN CONVERT DATABASE

The RMAN CONVERT DATABASE command is used to automate the movement of an entire database from one platform (the source platform) to another (the destination platform). Provided that the source and destination platforms are of the same endian format, RMAN automates most of the steps of creating a new database on the destination platform, which contains the same data as the source database, and which has, with a few exceptions, the same settings as the source database.

The CONVERT DATABASE command is one step in a longer process for transporting databases across platforms, which is described in the following sections. As with transporting tablespaces, the RMAN CONVERT DATABASE step in the process can be performed either on the source platform or the destination platform. The overall transport process changes depending upon whether you perform the CONVERT DATABASE step on the source or destination platform.

Files automatically transported to the destination platform during cross-platform database transport include:

- Datafiles that belong to permanent tablespaces.

Note: In spite of the fact that the endian formats for the source and destination platform are the same, the datafiles for a transportable database must undergo a conversion process and cannot simply be copied directly from one platform to another, as is possible with transporting tablespaces.

Unlike transporting tablespaces across platforms, transporting entire databases requires that certain types of blocks, such as blocks in undo segments, be reformatted to ensure compatibility with the destination platform.

- If a PFILE is used, it is transported. If an SPFILE is used, a PFILE is generated based on the SPFILE and transported, and a new SPFILE is created at the destination based on the settings in the PFILE.

Note: In most cases, some parameters in the PFILE require manual updating for the new database. For example, you may change the DB_NAME, as well as parameters such as CONTROL_FILES that indicate the locations of files on the destination host.

This section contains the following topics:

- [Restrictions on Cross-Platform Transportable Database](#)
- [Preparing for CONVERT DATABASE: Using the DBMS_TDB Package](#)
- [CONVERT DATABASE, Converting Datafiles on the Source Platform](#)
- [CONVERT DATABASE. Converting Datafiles on the Destination Host](#)

Restrictions on Cross-Platform Transportable Database

The principal restriction on cross-platform transportable database is that the source and destination platform must share the same endian format. For example, while you can transport a database from Microsoft Windows to Linux for x86 (both little-endian), or from HP-UX to AIX (both big-endian), you cannot transport a whole database from HP-UX to Linux for x86 using this feature. You can, however, create a new database on a destination platform manually, and transport needed tablespaces from the source database using cross-platform transportable tablespace as described in "[Cross-Platform Transportable Tablespace: CONVERT DATAFILE or TABLESPACE](#)" on page 15-1.

Also note that some parts of the database cannot be transported directly:

- Redo log files and control files from the source database are not transported. New control files and redo log files are created for the new database during the transport process, and an OPEN RESETLOGS is performed once the new database is created.

Note: The control file for the converted database does not contain a copy of the RMAN repository information from the source database. Backups from the source database cannot be used with the converted database.

- BFILEs are not transported. RMAN provides a list of objects using the BFILE datatype in the output for the `CONVERT DATABASE` command, but users must copy the BFILEs themselves and fix their locations on the destination database.
- Tempfiles belonging to locally managed temporary tablespaces are not transported. The temporary tablespace will be re-created on the target platform when the transport script is run.
- External tables and directories are not transported. RMAN provides a list of affected objects as part of the output of the `CONVERT DATABASE` command, but users must redefine these on the destination platform. See *Oracle Database Administrator's Guide* for more information on managing external tables and directories.
- Password files are not transported. If a password file was used with the source database, the output of `CONVERT DATABASE` includes a list of all usernames and their associated privileges. Create a new password file on the destination database using this information. See *Oracle Database Security Guide* for more information on managing password files.

Performing Cross-Platform Database Transport

The process for using cross-platform transportable database differs depending upon whether you choose to perform the conversion of database files on the source or destination system. This section contains the following topics:

- [Preparing for CONVERT DATABASE: Using the DBMS_TDB Package](#)
- [Using the RMAN CONVERT DATABASE Command](#)

Preparing for CONVERT DATABASE: Using the DBMS_TDB Package

The `DBMS_TDB` PL/SQL package defines two functions used in preparing for a `CONVERT DATABASE` operation. Their use is described in the following sections:

- [Using DBMS_TDB.CHECK_DB to Check Database State](#)
- [Using DBMS_TDB.CHECK_EXTERNAL to Identify External Objects](#)

See also: *PL/SQL Packages and Types Reference* for more details about the DBMS_TDB package

Note: Each of these subprograms is best run with SERVEROUTPUT set to ON, so that the descriptive output of the subprogram is visible.

Using DBMS_TDB.CHECK_DB to Check Database State

DBMS_TDB.CHECK_DB checks whether a database can be transported to a desired destination platform, and whether the current state of the database permits transport. It can be called without arguments to see if any condition at the source database prevents transport. It can also be called with one or both of the following arguments:

Table 15–1 CHECK_DB Procedure Parameters

Parameter	Description
target_platform_name	The name of the destination platform, as it appears in V\$DB_TRANSPORTABLE_PLATFORM. This parameter is optional, but is required when the skip_option parameter is used. If omitted, it is assumed that the destination platform is compatible with the source platform, and only the conditions in Table 15–2, "Conditions Tested by CHECK_DB Preventing Use of CONVERT DATABASE" not related to platform compatibility are tested.
skip_option	Specifies which, if any, parts of the database to skip when checking whether the database can be transported. Supported values (of type NUMBER) are: <ul style="list-style-type: none"> ■ SKIP_NONE (or 0), which checks all tablespaces ■ SKIP_OFFLINE (or 2), which skips checking datafiles in offline tablespaces ■ SKIP_READONLY (or 3), which skips checking datafiles in read-only tablespaces

DBMS_TDB.CHECK_DB returns TRUE if the source database can be transported using CONVERT DATABASE, and FALSE otherwise.

Make sure your database is open in read-only mode, then call DBMS_TDB.CHECK_DB with appropriate parameters.

If SERVEROUTPUT is ON, and DBMS_TDB.CHECK_DB returns FALSE, then the output includes the reason why the database cannot be transported. Possible conditions preventing the use of CONVERT DATABASE and their resolution are listed in the following table:

Table 15–2 Conditions Tested by CHECK_DB Preventing Use of CONVERT DATABASE

Condition	Action
Unrecognized target platform name.	Check V\$DB_TRANSPORTABLE_PLATFORM for recognized platform names.

Table 15–2 (Cont.) Conditions Tested by CHECK_DB Preventing Use of CONVERT

Condition	Action
Target platform has a different endian format.	Conversion is not supported.
Database is not open read-only.	Open database read-only and retry.
There are active or in-doubt transactions in the database.	Open the database read-write. After the active transactions are rolled back and the in-doubt transactions are resolved, open the database read-only and retry. This can happen if users flashback the database and open it read only. The active transactions will be rolled back when the database is opened read-write.
Deferred transaction rollback needs to be done.	Open the database read-write and bring online the necessary tablespaces. Once the deferred transaction rollback is complete, open the database read-only and retry.
Database compatibility version is below 10.	Change the init.ora COMPATIBLE parameter to 10 or higher, open the database read-only and retry.
Some tablespaces have not been open read-write with compatibility version is 10 or higher.	Change the init.ora COMPATIBLE parameter to 10 or higher. Then open the affected tablespaces read-write. Then shut down the database, open it read-only, and retry.

This example illustrates the use of CHECK_DB on a 32-bit Linux platform for transporting a database to 32-bit Windows, skipping read-only tablespaces, with a database that is currently open read-write.

```
SQL> set serveroutput on
SQL> declare
    db_ready boolean;
begin
    db_ready := dbms_tdb.check_db('Microsoft Windows IA (32-bit)', dbms_
tdb.skip_readonly);
end;
/
```

Database is not open READ ONLY. Please open database READ ONLY and retry.

PL/SQL procedure successfully completed.

If you call DBMS_TDB.CHECK_DB and no messages are displayed indicating conditions preventing transport before the PL/SQL procedure successfully completed message, then your database is ready for transport.

Using DBMS_TDB.CHECK_EXTERNAL to Identify External Objects

DBMS_TDB.CHECK_EXTERNAL must be used to identify any external tables, directories or BFILES. RMAN cannot automate the transport of such files.

DBMS_TDB.CHECK_EXTERNAL takes no parameters. With SERVEROUTPUT set to ON, the output of DBMS_TDB.CHECK_EXTERNAL lists the external tables, directories and BFILES of your database.

The following example shows how to call DBMS_TDB.CHECK_EXTERNAL:

```
SQL> set serveroutput on
```

```

SQL> declare
      external boolean;
begin
  /* value of external is ignored, but with SERVEROUTPUT set to ON
   * dbms_tdb.check_external displays report of external objects
   * on console */
  external := dbms_tdb.check_external;
end;

```

If there are no external objects, then this procedure completes with no output. If there are external objects, however, the output will be similar to the following example:

```

The following external tables exist in the database:
SH.SALES_TRANSACTIONS_EXT
The following directories exist in the database:
SYS.DATA_PUMP_DIR, SYS.MEDIA_DIR, SYS.DATA_FILE_DIR, SYS.LOG_FILE_DIR
The following BFILES exist in the database:
PM.PRINT_MEDIA

PL/SQL procedure successfully completed.

```

Using the RMAN CONVERT DATABASE Command

The process for using CONVERT DATABASE is different depending upon whether the conversion is performed on the source platform or the destination platform. This section includes the following topics:

- [CONVERT DATABASE, Converting Datafiles on the Source Platform](#)
- [CONVERT DATABASE. Converting Datafiles on the Destination Host](#)

CONVERT DATABASE, Converting Datafiles on the Source Platform

When the RMAN CONVERT DATABASE step is to be performed on the source platform, the process for transporting databases across platforms works as follows:

- In preparation for transporting the database, the source database must be opened read-only.

```

SQL> STARTUP MOUNT;
SQL> ALTER DATABASE OPEN READ ONLY;

```

- Use the CHECK_DB function in the DBMS_TDB package as described in "[Preparing for CONVERT DATABASE: Using the DBMS_TDB Package](#)" on page 15-8 to ensure that no conditions exist that would prevent the transport of the database, such as incorrect compatibility settings, in-doubt or active transactions, or incompatible endian formats between the source platform and the desired destination platform.

```

set serveroutput on
declare
  db_ready boolean;
begin
  /* db_ready is ignored, but with SERVEROUTPUT set to ON any
   * conditions preventing transport will be output to console */
  db_ready := dbms_tdb.check_db('Microsoft Windows IA (32-bit)',
    dbms_tdb.skip_none);
end;

```

- DBMS_TDB.CHECK_EXTERNAL must be used to identify any external objects:

```

SQL> set serveroutput on

```

```
SQL> declare
    external boolean;
begin
    /* value of external is ignored, but with SERVEROUTPUT set to ON
    * dbms_tdb.check_external displays report of external objects
    * on console */
    external := dbms_tdb.check_external;
end;
```

Typical output of `DBMS_TDB.CHECK_EXTERNAL` is shown in this example:

```
The following external tables exist in the database:
SH.SALES_TRANSACTIONS_EXT
The following directories exist in the database:
SYS.DATA_PUMP_DIR, SYS.MEDIA_DIR, SYS.DATA_FILE_DIR, SYS.LOG_FILE_DIR
The following BFILES exist in the database:
PM.PRINT_MEDIA
```

PL/SQL procedure successfully completed.

- When the database is ready for transport, the `RMAN CONVERT DATABASE` command is run, specifying a destination platform and how to name the output files. RMAN produces the files needed to move the database to the destination system, including the following:

- A complete copy of the datafiles of the database, ready to be transported
- A PFILE for use with the new database on the destination platform, containing settings used in the PFILE or SPFILE from the source database. Several entries at the top of the PFILE should be edited when the database is moved to the destination platform:

```
# Please change the values of the following parameters:
control_files           = "/tmp/convertedb/cf_D-NEWDBT_id-1778429277_
00gb9u2s"
db_recovery_file_dest   = "/tmp/convertedb/orcva"
db_recovery_file_dest_size= 10737418240
instance_name           = "NEWDBT"
service_names           = "NEWDBT.regress.rdbms.dev.us.oracle.com"
plsql_native_library_dir = "/tmp/convertedb/plsqlnld1"
db_name                 = "NEWDBT"
```

- A **transport script**, which contains SQL statements used to create the new database on the destination platform

The following example demonstrates the use of `CONVERT DATABASE` on the source platform, along with its outputs. Output related to the transport script and the parameter file for the new database is highlighted.

```
RMAN> CONVERT DATABASE NEW DATABASE 'newdb'
    transport script '/tmp/convertedb/transportscript'
    to platform 'Microsoft Windows IA (32-bit)'
    db_file_name_convert '/disk1/oracle/dbs' '/tmp/convertedb'
    ;
```

```
Starting convert at 25-JAN-05
using channel ORA_DISK_1
```

```
External table SH.SALES_TRANSACTIONS_EXT found in the database
```

```
Directory SYS.DATA_PUMP_DIR found in the database
Directory SYS.MEDIA_DIR found in the database
```

```

Directory SYS.DATA_FILE_DIR found in the database
Directory SYS.LOG_FILE_DIR found in the database

BFILE PM.PRINT_MEDIA found in the database

User SYS with SYSDBA and SYSOPER privilege found in password file
User OPER with SYSDBA privilege found in password file
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00001 name=/disk1/oracle/dbs/tbs_01.f
converted datafile=/tmp/convertdb/tbs_01.f
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:15
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00002 name=/disk1/oracle/dbs/tbs_ax1.f
converted datafile=/tmp/convertdb/tbs_ax1.f
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:03
.
.
.
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00016 name=/disk1/oracle/dbs/tbs_52.f
converted datafile=/tmp/convertdb/tbs_52.f
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
Run SQL script /tmp/convertdb/transportscript on the target platform
to create database
Edit init.ora file init_00gb3vfv_1_0.ora.This PFILE will be used to
create the database on the target platform
To recompile all PL/SQL modules, run utlirp.sql and utlrp.sql on
the target platform
To change the internal database identifier, use DBNEWID Utility
Finished backup at 25-JAN-05

RMAN>

```

When CONVERT DATABASE completes, the source database may be opened read-write again. Then, all of the files produced must then be copied to the destination host.

- Place the datafiles in the desired locations on the destination host. If the path to the datafiles is different on the destination, then edit the transport script to refer to the new datafile locations. Also edit the PFILE to change any settings for the destination database.
- Then execute the transport script in SQL*Plus to create the new database on the destination host.

```
SQL> @transportscript
```

When the transport script finishes, the creation of the new database is complete.

CONVERT DATABASE. Converting Datafiles on the Destination Host

You may choose to convert the datafiles of the database being transported on the destination platform instead of the source platform. Reasons for doing this include:

- Avoiding any performance overhead on the source host due to the conversion process.
- Distributing a database from one source system to multiple recipients on several different platforms.

In such a case, the preparations for the transport process are the same as in ["CONVERT DATABASE, Converting Datafiles on the Source Platform"](#) on page 15-11.

You must still open the database read-only, use `DBMS_TDB.CHECK_DB` to identify any conditions that prevent transport, and use `DBMS_TDB.CHECK_EXTERNAL` to identify external objects.

The remaining steps are:

- Run the `RMAN CONVERT DATABASE` command on the source platform specifying the `ON TARGET PLATFORM` argument. When used in this manner, the command syntax is as follows:

```
CONVERT DATABASE ON TARGET PLATFORM
  CONVERT SCRIPT '/tmp/convertedb/convertscript.rman'
  TRANSPORT SCRIPT '/tmp/convertedb/transportscript.sql'
  new database 'newdb'
  FORMAT '/tmp/convertedb/%U'
```

As with `CONVERT DATABASE` on the source platform, `CONVERT DATABASE ON TARGET PLATFORM` produces a transport script containing SQL*Plus commands to create a new database on the destination platform, and a PFILE for the new database containing the same settings as the source database.

`CONVERT DATABASE ON TARGET PLATFORM` also generates a **convert script** containing `RMAN CONVERT DATAFILE` commands for each of the datafiles of the database being transported. The source datafiles must be copied unconverted to some temporary location at the destination, and then the convert script must be run at the destination to actually convert the datafiles to a format usable by the destination host. A typical convert script contains commands like the following:

```
RUN {

  CONVERT DATAFILE '/disk1/oracle/dbs/tbs_01.f'
  FROM PLATFORM 'Linux IA (32-bit)'
  FORMAT '/tmp/convertedb/data_D-TV_I-1778429277_TS-SYSTEM_FNO-1_7qgb9u2s';

  CONVERT DATAFILE '/disk1/oracle/dbs/tbs_ax1.f'
  FROM PLATFORM 'Linux IA (32-bit)'
  FORMAT '/tmp/convertedb/data_D-TV_I-1778429277_TS-SYSAUX_FNO-2_7rgb9u2s';

  CONVERT DATAFILE '/disk1/oracle/dbs/tbs_03.f'
  FROM PLATFORM 'Linux IA (32-bit)'
  FORMAT '/tmp/convertedb/data_D-TV_I-1778429277_TS-SYSTEM_FNO-17_7sgb9u2s';
.
.
.

  CONVERT DATAFILE '/disk1/oracle/dbs/tbs_51.f'
  FROM PLATFORM 'Linux IA (32-bit)'
  FORMAT '/tmp/convertedb/data_D-TV_I-1778429277_TS-TBS_5_FNO-15_8egb9u2u';

  CONVERT DATAFILE '/disk1/oracle/dbs/tbs_52.f'
  FROM PLATFORM 'Linux IA (32-bit)'
  FORMAT '/tmp/convertedb/data_D-TV_I-1778429277_TS-TBS_5_FNO-16_8fgb9u2u';

}
```

One `CONVERT DATAFILE` command is generated for each datafile to be converted.

Note that `CONVERT DATABASE ON TARGET PLATFORM` does not produce converted datafile copies.

If the filesystem containing the datafiles of the source database is accessible from the destination system using the same path names, then you can use the `convert`

script on the destination host without any changes. The CONVERT DATAFILE commands in the script produce datafile copies in the locations specified by FORMAT, converted for use with the new database. (Once the convert script has been run, the source database can be opened for read-write access again.)

Otherwise, while the datafiles are still read-only, copy them to a temporary location. (As soon as copies of the datafiles are made, the source database can be opened read-write again.) If necessary, move the copies of the datafiles to a temporary location on the destination host, and then update each CONVERT DATAFILE command in the convert script to use the temporary location of each datafile as input and the FORMAT parameter of each CONVERT command to specify the desired final location of the datafiles of the transported database.

This example shows the use of CONVERT DATABASE ON TARGET PLATFORM on the source host, with typical output:

```

RMAN> convert database on target platform
convert script '/tmp/convertdb/convertscript-target'
      transport script '/tmp/convertdb/transportscript-target'
      new database 'newdbt'
      format '/tmp/convertdb/%U'
      ;

```

```

Starting convert at 28-JAN-05
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=39 devtype=DISK

```

```

External table SH.SALES_TRANSACTIONS_EXT found in the database

```

```

Directory SYS.DATA_PUMP_DIR found in the database
Directory SYS.MEDIA_DIR found in the database
Directory SYS.DATA_FILE_DIR found in the database
Directory SYS.LOG_FILE_DIR found in the database

```

```

BFILE PM.PRINT_MEDIA found in the database

```

```

User SYS with SYSDBA and SYSOPER privilege found in password file
User OPER with SYSDBA privilege found in password file
channel ORA_DISK_1: starting to check datafiles
input datafile fno=00001 name=/disk1/oracle/dbs/tbs_01.f
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile fno=00002 name=/disk1/oracle/dbs/tbs_ax1.f
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile fno=00017 name=/disk1/oracle/dbs/tbs_03.f
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00

```

```

.
.
.

```

```

channel ORA_DISK_1: starting to check datafiles
input datafile fno=00015 name=/disk1/oracle/dbs/tbs_51.f
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile fno=00016 name=/disk1/oracle/dbs/tbs_52.f
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00

```

```

Run SQL script /tmp/convertdb/transportscript-target on the target platform to
create database

```

```

Edit init.ora file /tmp/convertdb/init_00gb9u2s_1_0.ora. This PFILE will be

```

```
used to create the database on the target platform
Run RMAN script /tmp/convertdb/convertscript-target on target platform to
convert datafiles
To recompile all PL/SQL modules, run utlirp.sql and utlpr.sql on the target
platform
To change the internal database identifier, use DBNEWID Utility
Finished backup at 28-JAN-05
```

Run the convert script on the target platform to prepare the datafiles, and make any desired changes to the parameter file. Then run the transport script to create the new database, as described in ["CONVERT DATABASE, Converting Datafiles on the Source Platform"](#) on page 15-11.

When the transport script completes, the new database is created.

Using RMAN CONVERT to Copy Files Between ASM and Non-ASM Storage

Native operating system commands such as Linux `cp` or Windows `COPY` cannot write or read files in Automated Storage Management (ASM) storage. Because RMAN can read and write ASM files, it can be used to copy datafiles into and out of ASM storage, or between ASM disk groups.

If your goal is to migrate parts or all of your database or flash recovery area into ASM storage, see [Chapter 16, "Migrating Databases To and From ASM with Recovery Manager"](#). If, however, you just to create copies of some datafiles from non-ASM to ASM storage on the same platform, you can use the `CONVERT` command, without specifying a source or destination platform.

The methods available for specifying the output filenames for `CONVERT` described in ["Rules for Renaming Files with CONVERT TABLESPACE"](#) on page 15-3 apply to this use of `CONVERT` as well.

Note: Using `CONVERT`, without specifying a source or destination platform, produces image copies of the datafiles passed to `CONVERT` in the locations you specify.

The files produced by using `CONVERT` in this manner are identical to the files that would be produced by `BACKUP AS COPY` of the same datafiles to the same destination. However, `CONVERT` does not record the datafile copies produced in the RMAN repository. The assumption is that these copies are not being created as part of a backup strategy.

Because these copies are not recorded in the repository, RMAN does not try to use them in future restore and recovery operations, or consider them when determining whether the retention policy is satisfied.

Use `BACKUP AS COPY` to create image copy backups that RMAN can use in future restore and recovery operations.

Copying Datafiles To ASM Using CONVERT DATAFILE: Example

In this example, two datafiles are copied from the non-ASM storage location `/disk1/dbs` to the ASM disk group `diskgroup` using the `CONVERT DATAFILE` command:

```
RMAN> convert datafile
```



```

'/disk1/dbs/my_tbs_f1.df',
'/disk1/dbs/my_tbs_f2.df'
format '+diskgroup';

```

CONVERT DATAFILE can be used in this manner to copy datafiles even if they are not part of the connected target database.

Copying Datafiles From ASM Using CONVERT DATAFILE: Example

This example illustrates copying a single datafile out of ASM storage using CONVERT DATAFILE:

```
RMAN> convert datafile "+DATAFILE/tbs_21.f" format "/tmp/conv_df_%U";
```

```

Starting backup at 03-JUN-05
using channel ORA_DISK_1
channel ORA_DISK_1: starting datafile conversion
input filename="+DATAFILE/tbs_21.f"
converted datafile=/tmp/conv_df_data_D-L2_I-2786301554_TS-TBS_2_FNO-6_18gm2hft
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:02
Finished backup at 03-JUN-05

```

Copying Tablespaces From ASM With CONVERT TABLESPACE: Example

In this example, the datafiles of tablespace tbs_2 (stored in ASM) are copied from an ASM storage location to non-ASM storage using the CONVERT TABLESPACE command:

```
RMAN> convert tablespace tbs_2 format '/tmp/tbs_2_%U.df';
```

```

Starting backup at 03-JUN-05
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=20 devtype=DISK
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00006 name="+DATAFILE/tbs_21.f"
converted datafile=/tmp/tbs_2_data_D-L2_I-2786301554_TS-TBS_2_FNO-6_11gm2fq9.df
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00007 name="+DATAFILE/tbs_22.f"
converted datafile=/tmp/tbs_2_data_D-L2_I-2786301554_TS-TBS_2_FNO-7_12gm2fqa.df
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00019 name="+DATAFILE/tbs_25.f"
converted datafile=/tmp/tbs_2_data_D-L2_I-2786301554_TS-TBS_2_FNO-19_13gm2fqb.df
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00009 name="+DATAFILE/tbs_23.f"
converted datafile=/tmp/tbs_2_data_D-L2_I-2786301554_TS-TBS_2_FNO-9_14gm2fqc.df
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
channel ORA_DISK_1: starting datafile conversion
input datafile fno=00010 name="+DATAFILE/tbs_24.f"
converted datafile=/tmp/tbs_2_data_D-L2_I-2786301554_TS-TBS_2_FNO-10_15gm2fqd.df
channel ORA_DISK_1: datafile conversion complete, elapsed time: 00:00:01
Finished backup at 03-JUN-05

```

Migrating Databases To and From ASM with Recovery Manager

This chapter describes how to migrate part or all of your database into and out of ASM storage using Recovery Manager. It includes the following topics:

- [Migrating a Database into ASM](#)
- [Migrating the Flash Recovery Area to ASM](#)
- [Migrating a Database from ASM to Non-ASM Storage](#)
- [PL/SQL Scripts Used in Migrating to ASM Storage](#)

Migrating a Database into ASM

To take advantage of Automatic Storage Management with an existing database you must migrate that database into ASM. This migration is performed using Recovery Manager (RMAN) even if you are not using RMAN for your primary backup and recovery strategy.

A database can be moved from non-ASM disk storage directly into ASM, assuming you have enough disk space to hold the entire database both in non-ASM storage and in ASM storage. If you do not have enough disk space to have a complete copy of your database in ASM and another in non-ASM storage simultaneously, you can modify the procedure described here to back the database up to tape, create an ASM disk group that uses the old disk space, and then restore the database from tape into ASM.

Note: Enterprise Manager provides a GUI-based option for migration of a database to ASM storage. See *Oracle Database 2 Day DBA* for details.

Limitation on ASM Migration with Transportable Tablespaces

The procedure described here does not work for transportable (foreign) tablespaces. Such tablespaces need to be made read-write and imported into the database, before they can be migrated into ASM using this procedure.

Preparing to Migrate a Database to ASM

There are several steps required to prepare your database for migration and collect useful information you will need later, before you start the actual migration process.

Determine Names of Database Files

Obtain the filenames of the control files, datafiles, and online redo logs for your database. This information will be useful if you decide to migrate back to old (non-ASM) storage later. Information about datafiles is available by querying `V$DATAFILE`, online log file names in `V$LOGFILE`, and the control file names can be found in the `CONTROL_FILES` initialization parameter.

Generate RMAN Command File to Undo ASM Migration

If you need to migrate your database back to non-ASM storage later, this process will be simplified if you generate an RMAN command file now with the necessary commands to perform this migration. Even if you make changes to your database later, such as adding datafiles, the command file you create now will serve as a useful starting point.

There is a PL/SQL script described in "[Generating ASM-to-Non-ASM Storage Migration Script](#)" on page 16-10 which generates the necessary RMAN commands for you. Run this script and save the output as part of the permanent records you keep for your database.

Disk-Based Migration of a Database to ASM

If you have enough disk space that you can have both your entire non-ASM database and your ASM disk group on disk at the same time, you can do the migration directly without using tapes.

The procedure differs slightly between primary and standby databases. A number of the steps described in this procedure apply only in one or the other case. There are also a few steps where the procedure is different depending upon whether you are using a recovery catalog. The steps that vary are identified as necessary in the description of the process.

The process described in this section is intended to minimize downtime for your database. The example assumes that the destination ASM disk group for database files is `+DISK`, and a separate ASM disk group `+FRA` will be used to store the flash recovery area.

Note: During the migration process all flashback logs are discarded. As a result, any guaranteed restore points in the database become unusable. You should drop all guaranteed restore points before performing the migration.

To perform the migration, carry out the following steps:

1. Back up your database files as copies to the ASM disk group.

```
BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE
      FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';
```

You can perform this backup with multiple channels to improve performance, depending upon your hardware configuration. For example:

```
run {
  allocate channel dev1 type disk;
  allocate channel dev2 type disk;
  allocate channel dev3 type disk;
  allocate channel dev4 type disk;
  BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE
```

```

        FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';
    }

```

To ensure that the backup can also be made consistent, archive the current redo log after the backup:

```

RMAN> sql 'alter system archive log current';

```

Note: This backup may take a long time, depending upon the size of your database. If there has been a lot of activity on the database during the time the backup was created, you may wish to use the following procedure to create an incremental backup of the database afterwards, to refresh the copy with changes since the migration process started. If so, use the following script:

```

RMAN> backup incremental level 1 for recover of copy
        with tag 'ORA_ASM_MIGRATION' database ;
RMAN> recover copy of database with tag 'ORA_ASM_MIGRATION';

```

This minimizes the time required for the media recovery performed just before the copy of the database in ASM is opened at the end of the migration process. You may also want to perform this step using multiple channels, if using them improves performance in your environment.

2. Create a copy of the SPFILE in the ASM disk group. In this example, the SPFILE for the migrated database will be stored as +DISK/spfile.

If the database is using an SPFILE already, then run these commands:

```

run {
    BACKUP AS BACKUPSET SPFILE;
    RESTORE SPFILE TO "+DISK/spfile";
}

```

If you are not using an SPFILE, then use CREATE SPFILE from SQL*Plus to create the new SPFILE in ASM. For example, if your parameter file is called /private/init.ora, use the following command:

```

SQL> create spfile='+DISK/spfile' from pfile='/private/init.ora';

```

3. If this is standby database, stop managed recovery mode.

```

SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL;

```

Perform a consistent shutdown of the database.

```

RMAN> SHUTDOWN IMMEDIATE;

```

4. At this point, if you want the option of easily returning the database to non-ASM storage later, make copies of your current control file and all online logs. This command backs up the current control file to a non-ASM location:

```

RMAN> STARTUP MOUNT;
RMAN> BACKUP AS COPY CURRENT CONTROLFILE FORMAT '/disk1/pre-ASM-controfile.cf';

```

Note: RMAN cannot be used to backup your online logs. You must use operating-system commands to copy them.

- Now create an init.ora specifying the location of the new SPFILE, and start the instance with it. For example, create /tmp/pfile.ora with the following contents:

```
SPFILE=+DISK/spfile
```

Now start the database in NOMOUNT:

```
SQL> startup nomount PFILE="/tmp/pfile.ora";
```

- The next step is to migrate the control file to ASM.

In SQL*Plus, change the CONTROL_FILES initialization parameter using the following command:

```
SQL> alter system set control_files='+DISK/ct1.f','+FRA/ct2.f' scope=spfile
sid='*';
```

- Now specify the location of the flash recovery area by setting DB_RECOVERY_FILE_DEST and DB_RECOVERY_FILE_DEST_SIZE. Assuming that the desired size of the flash recovery area is 100 gigabytes, enter the following commands in SQL*Plus to set the parameters:

```
SQL> alter system set DB_RECOVERY_FILE_DEST_SIZE=100G SID='*';
SQL> alter system set DB_RECOVERY_FILE_DEST='+FRA' SID='*';
```

- Shut down and startup in NOMOUNT again, so that the changed parameters take effect. (The CONTROL_FILES parameter change only takes effect upon a restart because it is a static parameter.)

Then, use RMAN to actually create the new control files in ASM. For example, assuming that one of your original control file locations was /private/ct1.f, use the following command:

```
RMAN> shutdown immediate;
RMAN> startup nomount PFILE='/tmp/pfile.ora'; #using ASM SPFILE now
RMAN> restore controlfile from '/private/ct1.f';
RMAN> alter database mount;
RMAN> switch database to copy;
RMAN> recover database;
```

- The next step is to migrate your tempfiles to ASM. You must use a SET NEWNAME command for each tempfile to direct it to ASM, then a SWITCH to make the new names take effect.

```
RMAN > run {
    set newname for tempfile 1 to '+DISK'
    set newname for tempfile 2 to '+DISK';
    ...
    switch tempfile all;
}
```

The new tempfiles are created when you open the database.

- Disable logging for Flashback Database, and then re-enable it again to start creating flashback logs in the new ASM flash recovery area. For example:

```
SQL> ALTER DATABASE FLASHBACK OFF;
SQL> ALTER DATABASE FLASHBACK ON;
```

Note: Flashback logs cannot be migrated. All data in the flashback logs is lost.

11. The change tracking file cannot be migrated. You can only disable change tracking, then re-enable it, specifying an ASM disk location for the change tracking file:

```
SQL> alter database disable block change tracking;
SQL> ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE '+DISK';
```

12. At this point, if the database is a primary database, then open the database.

```
SQL> ALTER DATABASE OPEN;
```

For a standby database, resume managed recovery mode:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE;
```

13. For a primary database, migrating the online logs is performed by adding new log group members in ASM, and then dropping the old members. The easiest way to perform this step is to use the PL/SQL script in "[Migrating Online Logs of Primary Database to ASM](#)" on page 16-10.

For a standby database, you can follow similar steps to the script to drop the old standby redo logs and add new ones in the +DISK disk group, but the online redo logs cannot be migrated until the database is opened as a primary.

At this point the migration is complete. Your database and flash recovery area are stored in ASM. You may wish to move your existing flash recovery area backups using the process described in "[Migrating Existing Backups to ASM Flash Recovery Area](#)" on page 16-9.

Migrating the Flash Recovery Area to ASM

This section describes procedures to use if you already have backups on disk in non-ASM disk storage and you want to start using ASM to store your flash recovery area. You must set the initialization parameters related to the flash recovery area to refer to the ASM disk location. Then you can migrate existing backups, online logs, to the new ASM flash recovery area

- [Setting Initialization Parameters for Flash Recovery Area in ASM](#)
- [Migrating the Control File to an ASM Flash Recovery Area](#)
- [Migrating Existing Backups to ASM Flash Recovery Area](#)

Setting Initialization Parameters for Flash Recovery Area in ASM

The following process moves the flash recovery area from non-ASM disk storage to an ASM disk group named +FRA.

Note:

- If you have already migrated all of these files to ASM storage using the procedure in ["Disk-Based Migration of a Database to ASM"](#) on page 16-2 you do not need to perform this step.
- Before changing the location of the flash recovery area, you should drop any guaranteed restore points. Flashback logs used to support guaranteed restore points are stored in the flash recovery area.

Specify the location of the flash recovery area by setting `DB_RECOVERY_FILE_DEST` and `DB_RECOVERY_FILE_DEST_SIZE`. (For this example, assume the intended size of the flash recovery area is 100 gigabytes.) If you are using an SPFILE then in SQL*Plus enter the following commands:

```
SQL> alter system set DB_RECOVERY_FILE_DEST_SIZE=100G SID='*';
SQL> alter system set DB_RECOVERY_FILE_DEST='+FRA' SID='*';
```

If you are using a PFILE, then shut down the database, edit the above parameters in the PFILE with the new values for `DB_RECOVERY_FILE_DEST` and `DB_RECOVERY_FILE_DEST_SIZE` and restart the instance.

Migrating the Control File to an ASM Flash Recovery Area

Note: If you have already migrated all of these files to ASM storage using the procedure in ["Disk-Based Migration of a Database to ASM"](#) on page 16-2 you do not need to perform this step.

In this example, it is assumed that you have already set the initialization parameters for a flash recovery area in ASM storage, using the process in ["Setting Initialization Parameters for Flash Recovery Area in ASM"](#) on page 16-5. It is also assumed that one control file is already stored in `+DISK/ct1.f`, the other in non-ASM storage. The goal is to move the non-ASM control file to the flash recovery area and store it as `+FRA/ct2.f`.

- In SQL*Plus, bring the database to NOMOUNT:

```
SQL> SHUTDOWN IMMEDIATE;
SQL> STARTUP NOMOUNT;
```

Change the `CONTROL_FILES` initialization parameter to refer to the new location. If you are using an SPFILE, use the following command:

```
SQL> alter system set control_files='+DISK/ct1.f','+FRA/ct2.f' scope=spfile
sid='*';
```

If using a PFILE, edit the PFILE with the new for the `CONTROL_FILES` initialization parameter.

Shut down and startup in NOMOUNT again, so that the changed `CONTROL_FILES` parameter takes effect.

- Then, use RMAN to actually create the new control files in ASM.

```
RMAN> restore controlfile from '+DISK/ct1.f';
RMAN> alter database mount;
```


- If you were using flashback logging before to support flashback database, you can re-enable it now. For example:

```
SQL> ALTER DATABASE FLASHBACK ON;
```

Note: If you have already migrated all of these files to ASM storage using the procedure in "[Disk-Based Migration of a Database to ASM](#)" on page 16-2 you do not need to perform this step.

The following procedure changes the database configuration so that the flash recovery area is used for all future backups.

1. The first step is to change the initialization parameters for the database to store the flash recovery area in ASM, as described in [Setting Initialization Parameters for Flash Recovery Area in ASM](#) on page 16-5.
2. If this database is a primary database and your online logs, control file or archived redo logs are in the flash recovery area, then perform a consistent shutdown of your database. For example:

```
SQL> SHUTDOWN IMMEDIATE
```

If this database is a standby database and your standby online logs, control file, or archive logs are in recovery area, then stop managed recovery mode and shut down the database.

3. Modify the initialization parameter file of the target database as follows:
 - Set DB_RECOVERY_FILE_DEST to the desired ASM disk group.
 - Modify DB_RECOVERY_FILE_DEST_SIZE if you need to change the size of the flash recovery area.
4. If you shut down the database in step 2, then bring the database to a NOMOUNT state. For example:

```
RMAN> STARTUP NOMOUNT
```

If the old flash recovery area has copy of the current control file, then restore control file from the old DB_RECOVERY_FILE_DEST and mount the database again.

```
RMAN> RESTORE CONTROLFILE FROM 'filename_of_old_control_file';
RMAN> ALTER DATABASE MOUNT;
```

5. The next step is to migrate the control file from the old flash recovery area to the new flash recovery area. In this example, one control file is stored as +DISK/ct1.f, the other as +FRA/ct2.f.

In SQL*Plus, change the CONTROL_FILES initialization parameter using the following command:

```
SQL> alter system set control_files='+DISK/ct1.f','+FRA/ct2.f' scope=spfile
sid='*';
```

6. If you were using flashback logging before to support flashback database, you can re-enable it now. For example:

```
SQL> ALTER DATABASE FLASHBACK ON;
```

At this point, all future files that are directed to the flash recovery area are created in the new ASM flash recovery area location.

Changing Flashback Log Location to ASM Flash Recovery Area

Note: If you have already migrated all of these files to ASM storage using the procedure in ["Disk-Based Migration of a Database to ASM"](#) on page 16-2 you do not need to perform this step.

In this example, it is assumed that you have already set the initialization parameters for a flash recovery area in ASM storage, using the process in ["Setting Initialization Parameters for Flash Recovery Area in ASM"](#) on page 16-5.

Because the actual flashback logs cannot be migrated, the only step required to move the location of flashback logs to the new ASM flash recovery area is to disable and then enable flashback logging. After a clean shutdown, mount the database and run the following commands in SQL*Plus:

```
SQL> ALTER DATABASE FLASHBACK OFF;
SQL> ALTER DATABASE FLASHBACK ON;
```

Future flashback logs will be created in the new flash recovery area. The old flashback logs are automatically deleted from non-ASM storage.

Migrating Online Logs to ASM Flash Recovery Area

Note: If you have already migrated all of these files to ASM storage using the procedure in ["Disk-Based Migration of a Database to ASM"](#) on page 16-2 you do not need to perform this step.

In this example, it is assumed that you have already set the initialization parameters for a flash recovery area in ASM storage, using the process in ["Setting Initialization Parameters for Flash Recovery Area in ASM"](#) on page 16-5.

For a primary database, migrating the online logs is performed by adding new log group members in ASM, and then dropping the old members. The database must be open to perform this task.

The easiest way to perform this step is to use a PL/SQL script based upon the one in ["Migrating Online Logs of Primary Database to ASM"](#) on page 16-10. However, change the script so that it does not specify the disk group name.

For example, change the code that creates the online logs from:

```
stmt := 'alter database add logfile thread ' ||
        rlcRec.thr || ' '+DISK' size ' ||
        rlcRec.bytes_k || 'K';
```

to:

```
stmt := 'alter database add logfile thread ' ||
        rlcRec.thr || ' size ' ||
        rlcRec.bytes_k || 'K';
```

Also change the code that creates the standby logs from:

```
stmt := 'alter database add standby logfile thread ' ||
```

```
rlcRec.thr || ' '+DISK'' size ' ||
rlcRec.bytes_k || 'K';
```

to the following code:

```
stmt := 'alter database add standby logfile thread ' ||
rlcRec.thr || ' size ' ||
rlcRec.bytes_k || 'K';
```

For a standby database, you can follow similar steps to the script to drop the old standby redo logs and add new ones in the +FRA disk group, but the online redo logs cannot be migrated until the database is opened as a primary.

Once you have run your script, the migration of online logs is complete.

Migrating Existing Backups to ASM Flash Recovery Area

In this example, it is assumed that you have already set the initialization parameters for a flash recovery area in ASM storage, using the process in ["Setting Initialization Parameters for Flash Recovery Area in ASM"](#) on page 16-5.

Note: After you configure the database to change the location of the flash recovery area, backups created in the old flash recovery area location remain in their old location, still count against the total disk quota of the flash recovery area, are deleted from the old flash recovery area as space is required for other files, and can still be managed by RMAN and used in RMAN recovery operations. There is no need to move existing backups to the new ASM flash recovery area, unless you need the disk space used by those files for other purposes.

If you do need to free the space taken up by leftover non-ASM flash recovery area files, your options include backing them up to tape (for example, by using `BACKUP RECOVERY AREA DELETE INPUT`) or moving the backups from the old flash recovery area location to the new one, as described in this section.

To back up the existing archived redo log files to the new flash recovery area, use this command:

```
RMAN> BACKUP AS COPY ARCHIVELOG ALL DELETE INPUT;
```

To move backup sets to the new flash recovery area, use this command:

```
RMAN> BACKUP DEVICE TYPE DISK BACKUPSET ALL DELETE INPUT;
```

To move all datafile copies to the new flash recovery area location, use this command:

```
RMAN> BACKUP AS COPY DATAFILECOPY ALL DELETE INPUT;
```

At this point, all backups have been moved from the old flash recovery area to the new one.

Migrating a Database from ASM to Non-ASM Storage

Migrating a database back from ASM storage to non-ASM storage is similar to the original migration to ASM. The general outline of the process can be based on the steps in ["Disk-Based Migration of a Database to ASM"](#) on page 16-2. The primary changes to these steps are to modify each step to reference file locations in non-ASM

storage. For example, the command that initially created the datafile backups in ASM that become the live datafiles for the migrated database was:

```
BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE
FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';
```

You could use the following command to move them back:

```
BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE
FORMAT '/non-asm/df/%U' TAG 'ORA_NON_ASM_MIGRATION';
```

Similar modifications can be applied to the other steps in the migration process and to the PL/SQL scripts used during migration.

PL/SQL Scripts Used in Migrating to ASM Storage

The following PL/SQL scripts perform tasks which arise in the migration scenarios described in this chapter.

Generating ASM-to-Non-ASM Storage Migration Script

You can use the following PL/SQL script to generate a series of RMAN commands that you can use to migrate your database back from ASM to non-ASM disk storage.

```
set serveroutput on;
declare
  cursor df is select file#, name from v$datafile;
begin
  dbms_output.put_line('run');
  dbms_output.put_line('{}');
  for dfrec in df loop
    dbms_output.put_line('set newname for datafile ' ||
      dfrec.file# || ' to '' ' || dfrec.name || '' ');
  end loop;
  dbms_output.put_line('restore database;');
  dbms_output.put_line('switch all;');
  dbms_output.put_line('{}');
end;
```

Run this PL/SQL script and save the output into a file. The result is an RMAN script which you can save to a file and later run as a command file in the RMAN client to migrate your datafiles back out of ASM storage to their original non-ASM locations. Even if you later add or delete datafiles, this script provides a useful starting point for a migration script that will work for the new database.

Migrating Online Logs of Primary Database to ASM

The following PL/SQL script can be used to migrate the online redo log groups into ASM, as part of migrating a database or a flash recovery area into ASM. For each online redo log group, the script adds a log file stored in ASM, archives the current redo logs, and then drops the non-ASM log file.

```
declare
  cursor rlc is
    select group# grp, thread# thr, bytes/1024 bytes_k, 'NO' srl
      from v$log
    union
    select group# grp, thread# thr, bytes/1024 bytes_k, 'YES' srl
      from v$standby_log
```

```
        order by 1;
    stmt      varchar2(2048);
    swtstmt   varchar2(1024) := 'alter system switch logfile';
    ckpstmt   varchar2(1024) := 'alter system checkpoint global';
begin
    for rlcRec in rlc loop
        if (rlcRec.srl = 'YES') then
            stmt := 'alter database add standby logfile thread ' ||
                rlcRec.thr || ' '+DISK' size ' ||
                rlcRec.bytes_k || 'K';
            execute immediate stmt;
            stmt := 'alter database drop standby logfile group ' || rlcRec.grp;
            execute immediate stmt;
        else
            stmt := 'alter database add logfile thread ' ||
                rlcRec.thr || ' '+DISK' size ' ||
                rlcRec.bytes_k || 'K';
            execute immediate stmt;
            begin
                stmt := 'alter database drop logfile group ' || rlcRec.grp;
                dbms_output.put_line(stmt);
                execute immediate stmt;
            exception
                when others then
                    execute immediate swtstmt;
                    execute immediate ckpstmt;
                    execute immediate stmt;
            end;
        end if;
    end loop;
end;
```


Part IV

Performing User-Managed Backup and Recovery

The following chapters describe how to perform backup and recovery when using a user-managed backup and recovery strategy, that is, one that does not depend upon Recovery Manager. This part of the book contains these chapters:

- [Chapter 17, "Making User-Managed Backups"](#)
- [Chapter 18, "Performing User-Managed Database Flashback and Recovery"](#)
- [Chapter 19, "Advanced User-Managed Recovery Scenarios"](#)
- [Chapter 20, "Performing User-Managed TSPITR"](#)
- [Chapter 21, "Troubleshooting User-Managed Media Recovery"](#)

Making User-Managed Backups

This chapter describes methods of backing up an Oracle database in a user-managed backup and recovery strategy, that is, a strategy that does not depend on using Recovery Manager (RMAN).

This chapter contains the following sections:

- [Querying V\\$ Views to Obtain Backup Information](#)
- [Making User-Managed Backups of the Whole Database](#)
- [Making User-Managed Backups of Offline Tablespaces and Datafiles](#)
- [Making User-Managed Backups of Online Tablespaces and Datafiles](#)
- [Making User-Managed Backups of the Control File](#)
- [Making User-Managed Backups of Archived Redo Logs](#)
- [Making User-Managed Backups in SUSPEND Mode](#)
- [Making User-Managed Backups to Raw Devices](#)
- [Verifying User-Managed Backups](#)
- [Making Logical Backups with Oracle Export Utilities](#)
- [Making User-Managed Backups of Miscellaneous Oracle Files](#)
- [Keeping Records of Current and Backup Database Files](#)

Querying V\$ Views to Obtain Backup Information

Before making a backup, you must identify all the files in your database and decide what to back up. Several V\$ views can provide the necessary information.

Listing Database Files Before a Backup

Use V\$DATAFILE, V\$LOGFILE and V\$CONTROLFILE to identify the datafiles, log files and control files for your database. This same procedure works whether you named these files manually or allowed Oracle Managed Files to name them.

To list datafiles, online redo logs, and control files:

1. Start SQL*Plus and query V\$DATAFILE to obtain a list of datafiles. For example, enter:

```
SQL> SELECT NAME FROM V$DATAFILE;
```

You can also join the V\$TABLESPACE and V\$DATAFILE views to obtain a listing of datafiles along with their associated tablespaces:

```
SELECT t.NAME "Tablespace", f.NAME "Datafile"
FROM V$TABLESPACE t, V$DATAFILE f
WHERE t.TS# = f.TS#
ORDER BY t.NAME;
```

2. Obtain the filenames of online redo log files by querying the V\$LOGFILE view. For example, issue the following query:

```
SQL> SELECT MEMBER FROM V$LOGFILE;
```

3. Obtain the filenames of the current control files by querying the V\$CONTROLFILE view. For example, issue the following query:

```
SQL> SELECT NAME FROM V$CONTROLFILE;
```

Note that you only need to back up one copy of a multiplexed control file.

4. If you plan to take a control file backup with the ALTER DATABASE BACKUP CONTROLFILE TO 'filename' statement, then save a list of all datafiles and online redo log files with the control file backup. Because the current database structure may not match the database structure at the time a given control file backup was created, saving a list of files recorded in the backup control file can aid the recovery procedure.

Determining Datafile Status for Online Tablespace Backups

To check whether a datafile is part of a current online tablespace backup, query the V\$BACKUP view.

This view is useful only for user-managed online tablespace backups, because neither RMAN backups nor offline tablespace backups require the datafiles of a tablespace to be in backup mode.

The V\$BACKUP view is most useful when the database is open. It is also useful immediately after an instance failure because it shows the backup status of the files at the time of the failure. Use this information to determine whether you have left any tablespaces in backup mode.

V\$BACKUP is not useful if the control file currently in use is a restored backup or a new control file created after the media failure occurred. A restored or re-created control file does not contain the information the database needs to populate V\$BACKUP accurately. Also, if you have restored a backup of a file, this file's STATUS in V\$BACKUP reflects the backup status of the older version of the file, not the most current version. Thus, this view can contain misleading data about restored files.

For example, the following query displays which datafiles are currently included in a tablespace that has been placed in backup mode:

```
SELECT t.name AS "TB_NAME", d.file# as "DF#", d.name AS "DF_NAME", b.status
FROM V$DATAFILE d, V$TABLESPACE t, V$BACKUP b
WHERE d.TS#=t.TS#
AND b.FILE#=d.FILE#
AND b.STATUS='ACTIVE';
```

The following sample output shows that the tools and users tablespaces currently have ACTIVE status:

TB_NAME	DF#	DF_NAME	STATUS
-----	-----	-----	-----

TOOLS	7	/oracle/oradata/trgt/tools01.dbf	ACTIVE
USERS	8	/oracle/oradata/trgt/users01.dbf	ACTIVE

In the `STATUS` column, `NOT ACTIVE` indicates that the file is not currently in backup mode (that is, you have not executed the `ALTER TABLESPACE . . . BEGIN BACKUP` or `ALTER DATABASE BEGIN BACKUP` statement), whereas `ACTIVE` indicates that the file is currently in backup mode.

Making User-Managed Backups of the Whole Database

You can make a whole database backup of all files in a database after the database has been shut down with the `NORMAL`, `IMMEDIATE`, or `TRANSACTIONAL` options. A whole database backup taken while the database is open or after an instance failure or `SHUTDOWN ABORT` is inconsistent. In such cases, the files are inconsistent with respect to the checkpoint SCN.

You can make a whole database backup if a database is operating in either `ARCHIVELOG` or `NOARCHIVELOG` mode. If you run the database in `NOARCHIVELOG` mode, however, the backup must be consistent; that is, you must shut down the database cleanly before the backup.

The set of backup files that results from a consistent whole database backup is consistent because all files are checkpointed to the same SCN. You can restore the consistent database backup without further recovery. After restoring the backup files, you can perform additional recovery steps to recover the database to a more current time if the database is operated in `ARCHIVELOG` mode. Also, you can take inconsistent whole database backups if your database is in `ARCHIVELOG` mode.

Control files play a crucial role in database restore and recovery. For databases running in `ARCHIVELOG` mode, Oracle recommends that you back up control files with the `ALTER DATABASE BACKUP CONTROLFILE TO 'filename'` statement.

See Also: ["Making User-Managed Backups of the Control File"](#) on page 17-10 for more information about backing up control files

Making Consistent Whole Database Backups

This section describes how to back up the database with an operating system utility.

To make a consistent whole database backup:

1. If the database is open, use SQL*Plus to shut down the database with the `NORMAL`, `IMMEDIATE`, or `TRANSACTIONAL` options.
2. Use an operating system utility to make backups of all datafiles as well as all control files specified by the `CONTROL_FILES` parameter of the initialization parameter file. Also, back up the initialization parameter file and other Oracle product initialization files. To find these files, do a search for `*.ora` starting in your Oracle home directory and recursively search all of its subdirectories.

For example, you can back up the datafiles, control files and archived logs to `/disk2/backup` as follows:

```
% cp $ORACLE_HOME/oradata/trgt/*.dbf /disk2/backup
% cp $ORACLE_HOME/oradata/trgt/arch/* /disk2/backup/arch
```

3. Restart the database. For example, enter:

```
SQL> STARTUP
```

See Also: *Oracle Database Administrator's Guide* for more information on starting up and shutting down a database

Making User-Managed Backups of Offline Tablespaces and Datafiles

Note the following guidelines when backing up offline tablespaces:

- You cannot offline the `SYSTEM` tablespace or a tablespace with active rollback segments. The following procedure cannot be used for such tablespaces.
- Assume that a table is in tablespace `Primary` and its index is in tablespace `Index`. Taking tablespace `Index` offline while leaving tablespace `Primary` online can cause errors when DML is issued against the indexed tables located in `Primary`. The problem only manifests when the access method chosen by the optimizer needs to access the indexes in the `Index` tablespace.

To back up offline tablespaces:

1. Before beginning a backup of a tablespace, identify the tablespace's datafiles by querying the `DBA_DATA_FILES` view. For example, assume that you want to back up the `users` tablespace. Enter the following in SQL*Plus:

```
SELECT TABLESPACE_NAME, FILE_NAME
FROM SYS.DBA_DATA_FILES
WHERE TABLESPACE_NAME = 'USERS';
```

TABLESPACE_NAME	FILE_NAME
USERS	/oracle/oradata/trgt/users01.dbf

In this example, `/oracle/oradata/trgt/users01.dbf` is a fully specified filename corresponding to the datafile in the `users` tablespace.

2. Take the tablespace offline using normal priority if possible because it guarantees that you can subsequently bring the tablespace online without having to recover it. For example:

```
SQL> ALTER TABLESPACE users OFFLINE NORMAL;
```

3. Back up the offline datafiles. For example:

```
% cp /oracle/oradata/trgt/users01.dbf /d2/users01_'date +%m_%d_%y' '.dbf
```

4. Bring the tablespace online. For example:

```
ALTER TABLESPACE users ONLINE;
```

Note: If you took the tablespace offline using temporary or immediate priority, then you cannot bring the tablespace online unless you perform tablespace recovery.

5. Archive the unarchived redo logs so that the redo required to recover the tablespace backup is archived. For example, enter:

```
ALTER SYSTEM ARCHIVE LOG CURRENT;
```

Making User-Managed Backups of Online Tablespaces and Datafiles

You can back up all or only specific datafiles of an online tablespace while the database is open. The procedure differs depending on whether the online tablespace is read/write or read-only.

Note: You should not back up temporary tablespaces.

Making User-Managed Backups of Online Read/Write Tablespaces

You must put a read/write tablespace in backup mode to make user-managed datafile backups when the tablespace is online and the database is open. The `ALTER TABLESPACE ... BEGIN BACKUP` statement places a tablespace in backup mode. In backup mode, the database copies whole changed data blocks into the redo stream. After you take the tablespace out of backup mode with the `ALTER TABLESPACE . . . END BACKUP` or `ALTER DATABASE END BACKUP` statement, the database advances the datafile header to the current database checkpoint.

When restoring a datafile backed up in this way, the database asks for the appropriate set of redo log files to apply if recovery be needed. The redo logs contain all changes required to recover the datafiles and make them consistent.

To back up online read/write tablespaces in an open database:

1. Before beginning a backup of a tablespace, identify all of the datafiles in the tablespace with the `DBA_DATA_FILES` data dictionary view. For example, assume that you want to back up the `users` tablespace. Enter the following:

```
SELECT TABLESPACE_NAME, FILE_NAME
FROM SYS.DBA_DATA_FILES
WHERE TABLESPACE_NAME = 'USERS';
```

TABLESPACE_NAME	FILE_NAME
-----	-----
USERS	/oracle/oradata/trgt/users01.dbf
USERS	/oracle/oradata/trgt/users02.dbf

2. Mark the beginning of the online tablespace backup. For example, the following statement marks the start of an online backup for the tablespace `users`:

```
SQL> ALTER TABLESPACE users BEGIN BACKUP;
```

Caution: If you do not use `BEGIN BACKUP` to mark the beginning of an online tablespace backup and wait for that statement to complete before starting your copies of online tablespaces, or then the datafile copies produced are not usable for subsequent recovery operations. Attempting to recover such a backup is risky and can return errors that result in inconsistent data. For example, the attempted recovery operation can issue a "fuzzy files" warning, and can lead to an inconsistent database that you cannot open.

3. Back up the online datafiles of the online tablespace with operating system commands. For example, UNIX users might enter:

```
% cp /oracle/oradata/trgt/users01.dbf /d2/users01_'date +%m_%d_%y' '.dbf
% cp /oracle/oradata/trgt/users02.dbf /d2/users02_'date +%m_%d_%y' '.dbf
```

4. After backing up the datafiles of the online tablespace, run the SQL statement `ALTER TABLESPACE` with the `END BACKUP` option. For example, the following statement ends the online backup of the tablespace `users`:

```
SQL> ALTER TABLESPACE users END BACKUP;
```

5. Archive the unarchived redo logs so that the redo required to recover the tablespace backup is archived. For example, enter:

```
SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;
```

Caution: If you fail to take the tablespace out of backup mode, then Oracle continues to write copies of data blocks in this tablespace to the online logs, causing performance problems. Also, you will receive an `ORA-01149` error if you try to shut down the database with the tablespaces still in backup mode.

Making Multiple User-Managed Backups of Online Read/Write Tablespaces

When backing up several online tablespaces, you can back them up either serially or in parallel. Use either of the following procedures depending on your needs.

Backing Up Online Tablespaces in Parallel

You can simultaneously create datafile copies of multiple tablespaces requiring backups in backup mode. Note, however, that by putting all tablespaces in online mode at once, you can generate large redo logs if there is heavy update activity on the affected tablespaces, because the redo must contain a copy of each changed data block in each changed datafile. Be sure to consider the size of the likely redo before using the procedure outlined here.

To back up online tablespaces in parallel:

1. Prepare all online tablespaces for backup by issuing all necessary `ALTER TABLESPACE` statements at once. For example, put tablespaces `users`, `tools`, and `index` in backup mode as follows:

```
SQL> ALTER TABLESPACE users BEGIN BACKUP;
SQL> ALTER TABLESPACE tools BEGIN BACKUP;
SQL> ALTER TABLESPACE indx BEGIN BACKUP;
```

If you are backing up all tablespaces, you might want to use this command:

```
SQL> ALTER DATABASE BEGIN BACKUP;
```

2. Back up all files of the online tablespaces. For example, a UNIX user might back up datafiles with the `*.dbf` suffix as follows:

```
% cp $ORACLE_HOME/oradata/trgt/*.dbf /disk2/backup/
```

3. Take the tablespaces out of backup mode as in the following example:

```
SQL> ALTER TABLESPACE users END BACKUP;
SQL> ALTER TABLESPACE tools END BACKUP;
SQL> ALTER TABLESPACE indx END BACKUP;
```

Again, if you are handling all datafiles at once you can use the `ALTER DATABASE` command instead of `ALTER TABLESPACE`:

```
SQL> ALTER DATABASE END BACKUP;
```

4. Archive the online redo logs so that the redo required to recover the tablespace backups will be available for later media recovery. For example, enter:

```
SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;
```

Backing Up Online Tablespaces Serially

You can place all tablespaces requiring online backups in backup mode one at a time. Oracle recommends the serial backup option because it minimizes the time between `ALTER TABLESPACE . . . BEGIN/END BACKUP` statements. During online backups, more redo information is generated for the tablespace because whole data blocks are copied into the redo log.

To back up online tablespaces serially:

1. Prepare a tablespace for online backup. For example, to put tablespace `users` in backup mode enter the following:

```
SQL> ALTER TABLESPACE users BEGIN BACKUP;
```

In this case you probably do not want to use `ALTER DATABASE BEGIN BACKUP` to put all tablespaces in backup mode simultaneously, because of the unnecessary volume of redo log information generated for tablespaces in online mode.

2. Back up the datafiles in the tablespace. For example, enter:

```
% cp /oracle/oradata/trgt/users01.dbf /d2/users01_'date +%m_%d_%y' '.dbf
```

3. Take the tablespace out of backup mode. For example, enter:

```
SQL> ALTER TABLESPACE users END BACKUP;
```

4. Repeat this procedure for each remaining tablespace.
5. Archive the unarchived redo logs so that the redo required to recover the tablespace backups is archived. For example, enter:

```
SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;
```

Ending a Backup After an Instance Failure or SHUTDOWN ABORT

The following situations can cause a tablespace backup to fail and be incomplete:

- The backup completed, but you did not run the `ALTER TABLESPACE . . . END BACKUP` statement.
- An instance failure or `SHUTDOWN ABORT` interrupted the backup.

Whenever crash recovery is required, if a datafile is in backup mode when an attempt is made to open it, then the database will not open the database until either a recovery command is issued, or the datafile is taken out of backup mode.

For example, the database may display a message such as the following at startup:

```
ORA-01113: file 12 needs media recovery
ORA-01110: data file 12: '/oracle/dbs/tbs_41.f'
```

If the database indicates that the datafiles for multiple tablespaces require media recovery because you forgot to end the online backups for these tablespaces, then so long as the database is mounted, running the `ALTER DATABASE END BACKUP` statement takes all the datafiles out of backup mode simultaneously.

In high availability situations, and in situations when no DBA is monitoring the database, the requirement for user intervention is intolerable. Hence, you can write a crash recovery script that does the following:

1. Mounts the database
2. Runs the ALTER DATABASE END BACKUP statement
3. Runs ALTER DATABASE OPEN, allowing the system to come up automatically

An automated crash recovery script containing ALTER DATABASE END BACKUP is especially useful in the following situations:

- All nodes in an Oracle Real Application Clusters (RAC) configuration fail.
- One node fails in a **cold failover cluster** (that is, a cluster that is *not* a RAC configuration in which the secondary node must mount and recover the database when the first node fails).

Alternatively, you can take the following manual measures after the system fails with tablespaces in backup mode:

- Recover the database and avoid issuing END BACKUP statements altogether.
- Mount the database, then run ALTER TABLESPACE . . . END BACKUP for each tablespace still in backup mode.

Ending Backup Mode with the ALTER DATABASE END BACKUP Statement

You can run the ALTER DATABASE END BACKUP statement when you have multiple tablespaces still in backup mode. The primary purpose of this command is to allow a crash recovery script to restart a failed system without DBA intervention. You can also perform the following procedure manually.

To take tablespaces out of backup mode simultaneously:

1. Mount but do not open the database. For example, enter:

```
SQL> STARTUP MOUNT
```

2. If performing this procedure manually (that is, not as part of a crash recovery script), query the V\$BACKUP view to list the datafiles of the tablespaces that were being backed up before the database was restarted:

```
SQL> SELECT * FROM V$BACKUP WHERE STATUS = 'ACTIVE';
FILE#          STATUS          CHANGE#          TIME
-----
          12 ACTIVE          20863 25-NOV-02
          13 ACTIVE          20863 25-NOV-02
          20 ACTIVE          20863 25-NOV-02
3 rows selected.
```

3. Issue the ALTER DATABASE END BACKUP statement to take all datafiles currently in backup mode out of backup mode. For example, enter:

```
SQL> ALTER DATABASE END BACKUP;
```

You can use this statement only when the database is mounted but not open. If the database is open, use ALTER TABLESPACE . . . END BACKUP or ALTER DATABASE DATAFILE . . . END BACKUP for each affected tablespace or datafile.

Caution: Do not use ALTER DATABASE END BACKUP if you have restored any of the affected files from a backup.

Ending Backup Mode with the SQL*Plus RECOVER Command

The ALTER DATABASE END BACKUP statement is not the only way to respond to a failed online backup: you can also run the SQL*Plus RECOVER command. This method is useful when you are not sure whether someone has restored a backup, because if someone has indeed restored a backup, then the RECOVER command brings the backup up to date. Only run the ALTER DATABASE END BACKUP or ALTER TABLESPACE . . . END BACKUP statement if you are sure that the files are current.

Note: The RECOVER command method is slow because the database must scan redo generated from the beginning of the online backup.

To take tablespaces out of backup mode with the RECOVER command:

1. Mount the database. For example, enter:

```
SQL> STARTUP MOUNT
```

2. Recover the database as normal. For example, enter:

```
SQL> RECOVER DATABASE
```

3. Use the V\$BACKUP view to confirm that there are no active datafiles:

```
SQL> SELECT * FROM V$BACKUP WHERE STATUS = 'ACTIVE';
FILE#      STATUS          CHANGE#      TIME
-----
0 rows selected.
```

See Also: [Chapter 18, "Performing User-Managed Database Flashback and Recovery"](#) for information on recovering a database

Making User-Managed Backups of Read-Only Tablespaces

When backing up an online read-only tablespace, you can simply back up the online datafiles. You do not have to place the tablespace in backup mode because the database is not permitting changes to the datafiles.

If the set of read-only tablespaces is self-contained, then in addition to backing up the tablespaces with operating system commands, you can also export the tablespace metadata with the transportable tablespace functionality. In the event of a media error or a user error (such as accidentally dropping a table in the read-only tablespace), you can transport the tablespace back into the database.

See Also: *Oracle Database Administrator's Guide* to learn how to transport tablespaces

To back up online read-only tablespaces in an open database:

1. Query the DBA_TABLESPACES view to determine which tablespaces are read-only. For example, run this query:

```
SELECT TABLESPACE_NAME, STATUS
FROM DBA_TABLESPACES
WHERE STATUS = 'READ ONLY';
```

2. Before beginning a backup of a read-only tablespace, identify all of the tablespace's datafiles by querying the DBA_DATA_FILES data dictionary view. For example, assume that you want to back up the history tablespace:

```
SELECT TABLESPACE_NAME, FILE_NAME
FROM SYS.DBA_DATA_FILES
WHERE TABLESPACE_NAME = 'HISTORY';
```

TABLESPACE_NAME	FILE_NAME
HISTORY	/oracle/oradata/trgt/history01.dbf
HISTORY	/oracle/oradata/trgt/history02.dbf

3. Back up the online datafiles of the read-only tablespace with operating system commands. You do not have to take the tablespace offline or put the tablespace in backup mode because users are automatically prevented from making changes to the read-only tablespace. For example:

```
% cp $ORACLE_HOME/oradata/trgt/history*.dbf /disk2/backup/
```

Note: When restoring a backup of a read-only tablespace, take the tablespace offline, restore the datafiles, then bring the tablespace online. A backup of a read-only tablespace is still usable if the read-only tablespace is made read/write after the backup, but the restored backup will require recovery.

4. Optionally, export the metadata in the read-only tablespace. By using the transportable tablespace feature, you can quickly restore the datafiles and import the metadata in case of media failure or user error. For example, export the metadata for tablespace `history` as follows:

```
% exp TRANSPORT_TABLESPACE=y TABLESPACES=(history) FILE=/disk2/backup/hs.dmp
```

See Also: *Oracle Database Reference* for more information about the `DBA_DATA_FILES` and `DBA_TABLESPACES` views

Making User-Managed Backups of the Control File

Back up the control file of a database after making a structural modification to a database operating in ARCHIVELOG mode. To back up a database's control file, you must have the ALTER DATABASE system privilege.

Backing Up the Control File to a Binary File

The primary method for backing up the control file is to use a SQL statement to generate a binary file. A binary backup is preferable to a trace file backup because it contains additional information such as the archived log history, offline range for read-only and offline tablespaces, and backup sets and copies (if you use RMAN). Note that binary control file backups do not include tempfile entries.

To back up the control file after a structural change:

1. Make the desired change to the database. For example, you may create a new tablespace:

```
CREATE TABLESPACE tbs_1 DATAFILE 'file_1.f' SIZE 10M;
```

2. Back up the database's control file, specifying a filename for the output binary file. The following example backs up a control file to `/disk1/backup/cf.bak`:

```
ALTER DATABASE BACKUP CONTROLFILE TO '/disk1/backup/cf.bak' REUSE;
```

Specify the REUSE option to make the new control file overwrite one that currently exists.

Backing Up the Control File to a Trace File

To back up the control file to a trace file, mount or open the database and issue the following SQL statement:

```
ALTER DATABASE BACKUP CONTROLFILE TO TRACE;
```

If you specify neither the RESETLOGS nor NORESETLOGS option in the SQL statement, then the resulting trace file contains versions of the control file for both RESETLOGS and NORESETLOGS options. Tempfile entries are included in the output using "ALTER TABLESPACE... ADD TEMPFILE" statements.

See Also: ["Recovery of Read-Only Files with a Re-Created Control File"](#) on page 19-5 for special issues relating to read-only, offline normal, and temporary files included in CREATE CONTROLFILE statements

Backing Up the Control File to a Trace File: Example

Assume that you want to generate a script that re-creates the control file for the sales database. The database has these characteristics:

- Three threads are enabled, of which thread 2 is public and thread 3 is private.
- The redo logs are multiplexed into three groups of two members each.
- The database has the following datafiles:
 - /diska/prod/sales/db/filea.dbf (offline datafile in online tablespace)
 - /diska/prod/sales/db/database1.dbf (online in SYSTEM tablespace)
 - /diska/prod/sales/db/fileb.dbf (only file in read-only tablespace)

You issue the following statement to create a trace file containing a CREATE CONTROLFILE ... NORESETLOGS statement:

```
ALTER DATABASE BACKUP CONTROLFILE TO TRACE NORESETLOGS;
```

You then edit the trace file to create a script that creates a new control file for the sales database based on the control file that was current when you generated the trace file. To avoid recovering offline normal or read-only tablespaces, edit them out of the CREATE CONTROLFILE statement in the trace file. When you open the database with the re-created control file, the dictionary check code will mark these omitted files as MISSING. You can run an ALTER DATABASE RENAME FILE statement renames them back to their original filenames.

For example, you can edit the CREATE CONTROLFILE ... NORESETLOGS script in the trace file as follows, renaming files labeled MISSING:

```
# The following statements will create a new control file and use it to open the
# database. Log history and RMAN metadata will be lost. Additional logs may be
# required for media recovery of offline datafiles. Use this only if the current
# version of all online logs are available.
```

```
STARTUP NOMOUNT
CREATE CONTROLFILE REUSE DATABASE SALES NORESETLOGS ARCHIVELOG
MAXLOGFILES 32
```

```

MAXLOGMEMBERS 2
MAXDATAFILES 32
MAXINSTANCES 16
MAXLOGHISTORY 1600
LOGFILE
  GROUP 1
    '/diska/prod/sales/db/log1t1.dbf',
    '/diskb/prod/sales/db/log1t2.dbf'
  ) SIZE 100K
  GROUP 2
    '/diska/prod/sales/db/log2t1.dbf',
    '/diskb/prod/sales/db/log2t2.dbf'
  ) SIZE 100K,
  GROUP 3
    '/diska/prod/sales/db/log3t1.dbf',
    '/diskb/prod/sales/db/log3t2.dbf'
  ) SIZE 100K
DATAFILE
  '/diska/prod/sales/db/database1.dbf',
  '/diskb/prod/sales/db/filea.dbf'
;

# This datafile is offline, but its tablespace is online. Take the datafile
# offline manually.
ALTER DATABASE DATAFILE '/diska/prod/sales/db/filea.dbf' OFFLINE;

# Recovery is required if any datafiles are restored backups,
# or if the most recent shutdown was not normal or immediate.
RECOVER DATABASE;

# All redo logs need archiving and a log switch is needed.
ALTER SYSTEM ARCHIVE LOG ALL;

# The database can now be opened normally.
ALTER DATABASE OPEN;

# The backup control file does not list read-only and normal offline tablespaces
# so that Oracle can avoid performing recovery on them. Oracle checks the data
# dictionary and finds information on these absent files and marks them
# 'MISSINGxxxx'. It then renames the missing files to acknowledge them without
# having to recover them.
ALTER DATABASE RENAME FILE 'MISSING0002'
  TO '/diska/prod/sales/db/fileb.dbf';

```

Making User-Managed Backups of Archived Redo Logs

To save disk space in your primary archiving location, you may want to back up archived logs to tape or to an alternative disk location. If you archive to multiple locations, then only back up one copy of each log sequence number.

To back up archived redo logs:

1. To determine which archived redo log files that the database has generated, query V\$ARCHIVED_LOG. For example, run the following query:

```

SELECT THREAD#, SEQUENCE#, NAME
FROM V$ARCHIVED_LOG;

```

2. Back up one copy of each log sequence number by using an operating system utility. This example backs up all logs in the primary archiving location to a disk devoted to log backups:

```
% cp $ORACLE_HOME/oracle/trgt/arch/* /disk2/backup/arch
```

See Also: *Oracle Database Reference* for more information about the data dictionary views

Making User-Managed Backups in SUSPEND Mode

This section contains the following topics:

- [About the Suspend/Resume Feature](#)
- [Making Backups in a Suspended Database](#)

About the Suspend/Resume Feature

Some third-party tools allow you to mirror a set of disks or logical devices, that is, maintain an exact duplicate of the primary data in another location, and then **split the mirror**. Splitting the mirror involves separating the copies so that you can use them independently.

With the `SUSPEND/RESUME` functionality, you can suspend I/O to the database, then split the mirror and make a backup of the split mirror. By using this feature, which complements the backup mode functionality, you can suspend database I/O so that no new I/O can be performed. You can then access the suspended database to make backups without I/O interference.

You do not need to use `SUSPEND/RESUME` to make split mirror backups in most cases, although it is necessary if your system requires the database cache to be free of **dirty buffers** before a volume can be split. Some RAID devices benefit from suspending writes while the split operation is occurring; your RAID vendor can advise you on whether your system would benefit from this feature.

The `ALTER SYSTEM SUSPEND` statement suspends the database by halting I/Os to datafile headers, datafiles, and control files. When the database is suspended, all pre-existing I/O operations can complete; however, any new database I/O access attempts are queued.

The `ALTER SYSTEM SUSPEND` and `ALTER SYSTEM RESUME` statements operate on the database and not just the instance. If the `ALTER SYSTEM SUSPEND` statement is entered on one system in a RAC configuration, then the internal locking mechanisms propagate the halt request across instances, thereby suspending I/O operations for all active instances in a given cluster.

Making Backups in a Suspended Database

After a successful database suspension, you can back up the database to disk or break the mirrors. Because suspending a database does not guarantee immediate termination of I/O, Oracle recommends that you precede the `ALTER SYSTEM SUSPEND` statement with a `BEGIN BACKUP` statement so that the tablespaces are placed in backup mode.

You must use conventional user-managed backup methods to back up split mirrors. `RMAN` cannot make database backups or copies because these operations require reading the datafile headers. After the database backup is finished or the mirrors are re-silvered, then you can resume normal database operations using the `ALTER SYSTEM RESUME` statement.

Backing up a suspended database without splitting mirrors can cause an extended database outage because the database is inaccessible during this time. If backups are taken by splitting mirrors, however, then the outage is nominal. The outage time depends on the size of cache to flush, the number of datafiles, and the time required to break the mirror.

Note the following restrictions for the `SUSPEND/RESUME` feature:

- In a RAC configuration, you should not start a new instance while the original nodes are suspended.
- No checkpoint is initiated by the `ALTER SYSTEM SUSPEND` or `ALTER SYSTEM RESUME` statements.
- You cannot issue `SHUTDOWN` with `IMMEDIATE`, `NORMAL`, or `TRANSACTIONAL` options while the database is suspended.
- Issuing `SHUTDOWN ABORT` on a database that was already suspended reactivates the database. This prevents media recovery or crash recovery from hanging.

To make a split mirror backup in SUSPEND mode:

1. Place the database tablespaces in backup mode. For example, to place tablespace `users` in backup mode enter:

```
ALTER TABLESPACE users BEGIN BACKUP;
```

If you are backing up all of the tablespaces for your database, you can instead use:

```
ALTER DATABASE BEGIN BACKUP;
```

2. If your mirror system has problems with splitting a mirror while disk writes are occurring, then suspend the database. For example, issue the following:

```
ALTER SYSTEM SUSPEND;
```

3. Check to make sure that the database is suspended by querying `V$INSTANCE`. For example:

```
SELECT DATABASE_STATUS FROM V$INSTANCE;
```

```
DATABASE_STATUS
-----
SUSPENDED
```

4. Split the mirrors at the operating system or hardware level.
5. End the database suspension. For example, issue the following statement:

```
ALTER SYSTEM RESUME;
```

6. Check to make sure that the database is active by querying `V$INSTANCE`. For example, enter:

```
SELECT DATABASE_STATUS FROM V$INSTANCE;
```

```
DATABASE_STATUS
-----
ACTIVE
```

7. Take the specified tablespaces out of backup mode. For example, enter the following to take tablespace `users` out of backup mode:

```
ALTER TABLESPACE users END BACKUP;
```

8. Copy the control file and archive the online redo logs as usual for a backup.

Caution: Do not use the `ALTER SYSTEM SUSPEND` statement as a substitute for placing a tablespace in backup mode.

See Also: *Oracle Database Administrator's Guide* for more information about the `SUSPEND/RESUME` feature, and *Oracle Database SQL Reference* for more information about the `ALTER SYSTEM` statement

Making User-Managed Backups to Raw Devices

A **raw device** is a disk or partition that does not have a file system. In other words, a raw device can contain only a single file. Backing up files on raw devices poses operating system specific issues. The following sections discuss some of these issues on two of the most common operating systems supporting Oracle: UNIX and Windows.

See Also: *Oracle Real Application Clusters Installation and Configuration Guide* for a general overview of raw devices as they relate to Oracle Real Application Clusters

Backing Up to Raw Devices on UNIX

When backing up to or from raw devices, the UNIX `dd` command is the most common backup utility. See your operating system specific documentation for complete details about this utility.

Using `dd` effectively requires specifying the correct options, based on your database. Details about your database that affect the options you use for `dd` are listed in the following table.

Data	Explanation
Block size	You can specify the size of the buffer that <code>dd</code> uses to copy data. For example, you can specify that <code>dd</code> should copy data in units of 8 KB or 64 KB. Note that the block size for <code>dd</code> need not correspond to either the Oracle block size or the operating system block size: it is merely the size of the buffer used by <code>dd</code> when making the copy.
Raw offset	On some systems, the beginning of the file on the raw device is reserved for use by the operating system. This storage space is called the raw offset . Oracle should not back up or restore these bytes.
Size of Oracle block 0	At the beginning of every Oracle file, the operating system-specific code places an Oracle block called block 0 . The generic Oracle code does not recognize this block, but the block is included in the size of the file on the operating system. Typically, this block is the same size as the other Oracle blocks in the file.

The information in the preceding table enables you to set the `dd` options specified in [Table 17-1](#).

Table 17–1 Options for dd Command

This option ...	Specifies ...
if	The name of the input file, that is, the file that you are reading.
of	The name of the output file, that is, the file to which you are writing.
bs	The buffer size used by dd to copy data.
skip	The number of dd buffers to skip on the input raw device if a raw offset exists. For example, if you are backing up a file on a raw device with a 64 KB raw offset, and the dd buffer size is 8 KB, then you can specify skip=8 so that the copy starts at offset 64 KB.
seek	The number of dd buffers to skip on the output raw device if a raw offset exists. For example, if you are backing up a file onto a raw device with a 64 KB raw offset, and the dd buffer size is 8 KB, then you can specify skip=8 so that the copy starts at offset 64 KB.
count	The number of blocks on the input raw device for dd to copy. It is best to specify the exact number of blocks to copy when copying from raw device to file system, otherwise any extra space at the end of the raw volume that is not used by the Oracle datafile is copied to the file system. Remember to include block 0 in the total size of the input file. For example, if the dd block size is 8 KB, and you are backing up a 30720 KB datafile, then you can set count=3841. This value for count actually backs up 30728 KB: the extra 8 KB are for Oracle block 0.

Because a raw device can be the input or output device for a backup, you have four possible scenarios for the backup. The possible options for dd depend on which scenario you choose, as illustrated in [Table 17–2](#).

Table 17–2 Scenarios Involving dd Backups

Backing Up from ...	Backing Up to ...	Options Specified for dd Command
Raw device	Raw device	if, of, bs, skip, seek, count
Raw device	File system	if, of, bs, skip, count
File system	Raw device	if, of, bs, seek
File system	File system	if, of, bs

Backing Up with the dd utility on UNIX: Examples

For these examples of dd utility usage, assume the following:

- You are backing up a 30720 KB datafile.
- The beginning of the datafile has a block 0 of 8 KB.
- The raw offset is 64 KB.
- You set the dd block size to 8 KB when a raw device is involved in the copy.

In the following example, you back up from one raw device to another raw device:

```
% dd if=/dev/rsd1b of=/dev/rsd2b bs=8k skip=8 seek=8 count=3841
```

In the following example, you back up from a raw device to a file system:

```
% dd if=/dev/rsd1b of=/backup/df1.dbf bs=8k skip=8 count=3841
```

In the following example, you back up from a file system to a raw device:

```
% dd if=/backup/df1.dbf of=/dev/rsd2b bs=8k seek=8
```


In the following example, you back up from a file system to a file system, and so can set the block size to a high value to boost I/O performance:

```
% dd if=/oracle/dbs/df1.dbf of=/backup/df1.dbf bs=1024k
```

Backing Up to Raw Devices on Windows

Like UNIX, Windows supports raw disk partitions in which the database can store datafiles, online logs, and control files. Each raw partition is assigned either a drive letter or physical drive number and does not contain a file system. As in UNIX, each raw partition on NT is mapped to a single file.

NT differs from UNIX in the naming convention for Oracle files. On NT, raw datafile names are formatted as follows:

```
\\.\drive_letter:
\\.\PHYSICALDRIVEdrive_number
```

For example, the following are possible raw filenames:

```
\\.\G:
\\.\PHYSICALDRIVE3
```

Note that you can also create aliases to raw filenames. The standard Oracle database installation provides a SETLINKS utility that can create aliases such as \\.\Datafile12 that point to filenames such as \\.\PHYSICALDRIVE3.

The procedure for making user-managed backups of raw datafiles is basically the same as for copying files on an NT file system, except that you should use the Oracle OCOPI utility rather than the NT-supplied copy.exe or ntbakup.exe utilities. OCOPI supports 64-bit file I/O, physical raw drives, and raw files. Note that OCOPI cannot back up directly to tape.

To display online documentation for OCOPI, enter OCOPI by itself at the Windows prompt. Sample output follows:

```
Usage of OCOPI:
  ocopy from_file [to_file [a | size_1 [size_n]]]
  ocopy -b from_file to_drive
  ocopy -r from_drive to_dir
```

Note the important OCOPI options described in the following table.

This option ...	Specifies ...
b	Splits the input file into multiple output files. This option is useful for backing up to devices that are smaller than the input file.
r	Combines multiple input files and writes to a single output file. This option is useful for restoring backups created with the -b option.

Backing Up with OCOPI: Example

In this example, assume the following:

- Datafile 12 is mounted on the \\.\G: raw partition.
- The C: drive mounts a file system.
- The database is open.

To back up the datafile on the raw partition `\\.\G:` to a local file system, you can run the following command at the prompt after placing datafile 12 in backup mode:

```
OCOPY "\\.\G:" C:\backup\datafile12.bak
```

Specifying the `-b` and `-r` Options for `OCOPY`: Example

In this example, assume the following:

- `\\.\G:` is a raw partition containing datafile 7
- The `A:` drive is a removable disk drive.
- The database is open.

To back up the datafile onto drive `A:`, you can execute the following command at the NT prompt after placing datafile 7 in backup mode:

```
# first argument is filename, second argument is drive  
OCOPY -b "\\.\G:" A:\
```

When drive `A:` fills up, you can use another disk. In this way, you can divide the backup of datafile 7 into multiple files.

Similarly, to restore the backup, take the tablespace containing datafile 7 offline and run this command:

```
# first argument is drive, second argument is directory  
OCOPY -r A:\ "\\.\G:"
```

Verifying User-Managed Backups

You should periodically verify your backups to ensure that they are usable for recovery. This section contains the following topics:

- [Testing the Restore of Backups](#)
- [Running the `DBVERIFY` Utility](#)

Testing the Restore of Backups

The best way to test the usability of backups is to restore them to a separate host and attempt to open the database, performing media recovery if necessary. This option requires that you have a separate host available for the restore procedure.

See Also:

- ["Restoring Datafiles with Operating System Utilities"](#) on page 18-4
- ["Restoring Archived Redo Logs with Operating System Utilities"](#) on page 18-5
- ["Restoring Control Files"](#) on page 18-6
- ["Performing Complete User-Managed Media Recovery"](#) on page 18-15 to learn how to recover files with `SQL*Plus`

Running the `DBVERIFY` Utility

The `DBVERIFY` program is an external command-line utility that performs a physical data structure integrity check on an offline datafile. Use `DBVERIFY` primarily when

you need to ensure that a user-managed backup of a datafile is valid before it is restored or as a diagnostic aid when you have encountered data corruption problems.

The name and location of DBVERIFY is dependent on your operating system. For example, to perform an integrity check on datafile `tbs_52.f` on UNIX, you can run the `dbv` command as follows:

```
% dbv file=tbs_52.f
```

Sample `dbv` output follows:

```
DBVERIFY - Verification starting : FILE = users01.dbf
```

```
DBVERIFY - Verification complete
```

```
Total Pages Examined          : 250
Total Pages Processed (Data)  : 1
Total Pages Failing (Data)    : 0
Total Pages Processed (Index): 0
Total Pages Failing (Index)   : 0
Total Pages Processed (Other): 2
Total Pages Processed (Seg)   : 0
Total Pages Failing (Seg)     : 0
Total Pages Empty             : 247
Total Pages Marked Corrupt    : 0
Total Pages Influx            : 0
```

See Also: *Oracle Database Utilities* to learn about DBVERIFY

Making Logical Backups with Oracle Export Utilities

Oracle import and export utilities move Oracle data in and out of Oracle databases. Export utilities write exported database objects to operating system files in an Oracle-proprietary format. Import utilities can read the files produced by export utilities and re-create database objects. Logical exports of data can be a useful supplement to physical database backups in some situations, especially in backing up recovery catalog databases.

There are two sets of Oracle database import and export utilities: Original Import and Export (which were used in previous releases) and Data Pump Import and Export (new for Oracle Database Release 10g). The Data Pump utilities offer better performance and more complete support of features of Oracle Database Release 10g. Whichever export tool you use to export objects to a file, you must use the corresponding import tool to import the objects from the file, that is, you cannot use the Data Pump Import utility to read the output of the Original Export utility, or the Original Import to read the output of the Data Pump Export utility.

See Also: *Oracle Database Utilities* for complete documentation of the Oracle import and export utilities, including a comparison of their capabilities.

Making User-Managed Backups of Miscellaneous Oracle Files

Always back up initialization parameter files, networking and configuration files, and password files. If a media failure destroys these files, then you may have difficulty re-creating your environment. For example, if you back up the database and server parameter file but do not back up the networking files, then you can restore and

recover the database but will not be able to authenticate users through Oracle Net until you re-create the networking files.

As a rule, you should back up miscellaneous Oracle files after changing them. For example, if you add or change the net service names that can be used to access the database, then create a new backup of the `tnsnames.ora` file.

The easiest way to find configuration files is to start in the Oracle home directory and do a recursive search for all files ending in the `.ora` extension. For example, on UNIX you can run this command:

```
% find $ORACLE_HOME -name "*.ora" -print
```

You must use third-party utilities to back up the configuration files. For example, you can use the UNIX `cp` command to back up the `tnsnames.ora` and `listener.ora` files as follows:

```
% cp $ORACLE_HOME/network/admin/tnsnames.ora /d2/tnsnames'date +%m_%d_%y'.ora
% cp $ORACLE_HOME/network/admin/listener.ora /d2/listener'date +%m_%d_%y'.ora
```

You can also use an operating system utility to back up the server parameter file. Although the database does not depend on the existence of a particular version of the server parameter file to be started, you should keep relatively current backups of this file so that you do not lose changes made to the file. Note that if you lose the server parameter file, you can always create a new one or start the instance with a client-side initialization parameter file (PFILE).

Keeping Records of Current and Backup Database Files

One of the most important aspects of user-managed backup and recovery is keeping records of all current database files as well as the backups of these files. For example, you should have records for the location of the following files:

- Datafiles and control files
- Online and archived redo logs (note that online logs are never backed up)
- Initialization parameter files
- Password files
- Networking-related files

Recording the Locations of Datafiles, Control Files, and Online Redo Logs

The following useful SQL script displays the location of all control files, datafiles, and online redo log files for the database:

```
SELECT NAME FROM V$DATAFILE
UNION ALL
SELECT MEMBER FROM V$LOGFILE
UNION ALL
SELECT NAME FROM V$CONTROLFILE;
```

See Also: *Oracle Database Reference* for more information on the V\$ views

Recording the Locations of Archived Redo Logs

You can determine the location of the default archived log destinations by executing the following SQL script:

```
SELECT NAME, VALUE
FROM V$PARAMETER
WHERE NAME LIKE log_archive_dest%
AND VALUE IS NOT NULL;
```

```
NAME                                VALUE
-----                                -----
log_archive_dest_1                   LOCATION=/oracle/work/arc_dest/arc
log_archive_dest_state_1             enable
```

Determine the format for archived logs by running SHOW as follows:

```
SHOW PARAMETER LOG_ARCHIVE_FORMAT
```

To see a list of all the archived logs recorded in the control file, issue this query:

```
SELECT NAME FROM V$ARCHIVED_LOG;
```

Recording the Locations and Dates of Backup Files

It is not enough to merely record the location of backup files: you must correlate the backups with the original files. If possible, name the backups with names derived from the name of the primary file. Whatever naming system you use, keep a table containing the relevant information. For example, you could keep the following table as a record of database file locations in case of a restore emergency.

Datafile #	Tbs	Current Datafiles	Backup Datafiles
0 (cf)	0 (cf)	/oracle/oradata/trgt/control01.dbf	/d2/control01_10_31_02.dbf
1	SYSTEM	/oracle/oradata/trgt/system01.dbf	/d2/system01_10_31_02.dbf
2	undo	/oracle/oradata/trgt/undo01.dbf	/d2/undo01_10_31_02.dbf
3	cwmlite	/oracle/oradata/trgt/cwmlite01.dbf	/d2/cwmlite01_10_31_02.dbf
4	drsys	/oracle/oradata/trgt/drsys01.dbf	/d2/drsys01_10_31_02.dbf

Performing User-Managed Database Flashback and Recovery

This chapter describes how to restore and recover a database, and use the flashback features of Oracle, when using a user-managed backup and recovery strategy, that is, a strategy that does not depend upon using Recovery Manager.

This chapter includes the following topics:

- [User-Managed Flashback Features of Oracle](#)
- [About User-Managed Restore Operations](#)
- [Determining Which Datafiles Require Recovery](#)
- [Restoring Datafiles and Archived Redo Logs](#)
- [Restoring Control Files](#)
- [About User-Managed Media Recovery](#)
- [Performing Complete User-Managed Media Recovery](#)
- [Performing User-Managed Database Point-in-Time Recovery](#)
- [Opening the Database with the RESETLOGS Option](#)
- [Recovering a Database in NOARCHIVELOG Mode](#)
- [Controlling Parallel Media Recovery](#)

User-Managed Flashback Features of Oracle

Oracle's flashback features, which let you undo damage to your database after logical data corruption, include the following:

- Oracle Flashback Database, which returns your entire database to a previous state without requiring you to restore files from backup;
- Oracle Flashback Table, which returns one or more tables to their contents at a previous time;
- Oracle Flashback Drop, which undoes the effects of the `DROP TABLE` operation;
- Oracle Flashback Query, which is used to query the contents of the database at a past time;
- Oracle Flashback Version Query, which lets you view past states of data;
- Oracle Flashback Transaction Query, which is used to review transactions affecting a table over time.

All of these operations are available within SQL*Plus, and none of them require the use of Recovery Manager. More details about using the flashback features of Oracle in data recovery situations are provided in *Oracle Database Backup and Recovery Basics*.

Performing Flashback Database with SQL*Plus

The SQL*Plus `FLASHBACK DATABASE` command performs the same function as the RMAN `FLASHBACK DATABASE` command: it returns the database to a prior state.

Note that using Flashback Database requires that you create a flash recovery area for your database and enable the collection of flashback logs. See *Oracle Database Backup and Recovery Basics* for more details about how the Flashback Database feature works, requirements for using Flashback Database, and how to enable collection of flashback logs required for Flashback Database. The requirements and preparations are the same whether you use RMAN or user-managed backup and recovery.

To perform the `FLASHBACK DATABASE` operation in SQL*Plus:

1. Query the target database to determine the range of possible flashback SCNs. The following SQL*Plus queries show you the latest and earliest SCN in the flashback window:

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

```
SQL> SELECT OLDEST_FLASHBACK_SCN, OLDEST_FLASHBACK_TIME  
FROM V$FLASHBACK_DATABASE_LOG;
```

2. Use other flashback features if necessary, to identify the SCN or time of the unwanted changes to your database.
3. Start SQL*Plus with administrator privileges, and run the `FLASHBACK DATABASE` statement to return the database to a prior `TIMESTAMP` or `SCN`. For example:

```
FLASHBACK DATABASE TO SCN 46963;  
FLASHBACK DATABASE TO TIMESTAMP (SYSDATE-1/24);  
FLASHBACK DATABASE TO TIMESTAMP timestamp'2002-11-05 14:00:00';  
FLASHBACK DATABASE  
TO TIMESTAMP to_timestamp('2002-11-11 16:00:00', 'YYYY-MM-DD HH24:MI:SS');
```

Open the database read-only to examine the results of the Flashback Database operation. When the operation completes, you can open the database read-only and perform some queries to make sure you have recovered the data you need. If you find that you need to perform Flashback Database again to a different target time, then use `RECOVER DATABASE` to return the database back to the present time, and then try another `FLASHBACK DATABASE` statement.

If you are satisfied with the results of Flashback Database, then you can re-open your database with the `RESETLOGS` option. If appropriate, you can also use an Oracle export utility like Data Pump Export to save lost data, use `RECOVER DATABASE` to return the database to the present, and re-import the lost object.

About User-Managed Restore Operations

To restore a file is to replace it with a backup file. Typically, you restore a file when a media failure or user error has damaged or deleted the original file. The following files are candidates for restore operations:

- Datafiles and control files
- Archived redo logs

- Server parameter file

In each case, the loss of a primary file and the restore of a backup has the following implications for media recovery.

If you lose . . .	Then . . .
One or more datafiles	You must restore them from a backup and perform media recovery. Recovery is required whenever the checkpoint SCN in the datafile header does not match the checkpoint SCN for the datafile that is recorded in the control file.
All copies of the current control file	You must restore a backup control file and then open the database with the <code>RESETLOGS</code> option. If you do not have a backup, then you can attempt to re-create the control file. If possible, use the script included in the <code>ALTER DATABASE BACKUP CONTROLFILE TO TRACE</code> output. Additional work may be required to match the control file structure with the current database structure.
One copy of a multiplexed control file	Copy one of the intact multiplexed control files into the location of the damaged or missing control file and open the database. If you cannot copy the control file to its original location, then edit the initialization parameter file to reflect a new location or remove the damaged control file. Then, open the database.
One or more archived logs required for media recovery	You must restore backups of these archived logs for recovery to proceed. You can restore either to the default or nondefault location. If you do not have backups, then you must performing incomplete recovery up to an SCN before the first missing redo log and open <code>RESETLOGS</code> .
The server parameter file	If you have a backup of the server parameter file, then restore it. Alternatively, if you have a backup of the client-side initialization parameter file, then you can restore a backup of this file, start the instance, and then re-create the server parameter file.

Note: Restore and recovery of Oracle-managed files is no different from restore and recovery of user-named files.

Determining Which Datafiles Require Recovery

You can use the dynamic performance view `V$RECOVER_FILE` to determine which files to restore in preparation for media recovery. This view lists all files that need to be recovered and explains why they need to be recovered.

1. If you are planning to perform complete recovery rather than point-in-time recovery, you can recover only those datafiles which require recovery, rather than the whole database. (Note that for point-in-time recovery, you must restore and recover all datafiles, unless you perform tablespace point-in-time recovery as described in [Chapter 20, "Performing User-Managed TSPITR"](#). You can also use Flashback Database as described in ["User-Managed Flashback Features of Oracle"](#) on page 18-1, but this affects all datafiles and returns the entire database to a past time.)

You can query `V$RECOVER_FILE` to list datafiles requiring recovery by datafile number with their status and error information.

```
SELECT FILE#, ERROR, ONLINE_STATUS, CHANGE#, TIME
       FROM V$RECOVER_FILE;
```

Note: You cannot use V\$RECOVER_FILE with a control file restored from backup or a control file that was re-created after the time of the media failure affecting the datafiles. A restored or re-created control file does not contain the information needed to update V\$RECOVER_FILE accurately.

You can also perform useful joins using the datafile number and the V\$DATAFILE and V\$TABLESPACE views, to get the datafile and tablespace names. Use the following SQL*Plus commands to format the output of the query:

```
COL DF# FORMAT 999
COL DF_NAME FORMAT A35
COL TBSP_NAME FORMAT A7
COL STATUS FORMAT A7
COL ERROR FORMAT A10
COL CHANGE# FORMAT 99999999
SELECT r.FILE# AS df#, d.NAME AS df_name, t.NAME AS tbsp_name,
       d.STATUS, r.ERROR, r.CHANGE#, r.TIME
FROM V$RECOVER_FILE r, V$DATAFILE d, V$TABLESPACE t
WHERE t.TS# = d.TS#
AND d.FILE# = r.FILE#
;
```

The ERROR column identifies the problem for each file requiring recovery.

See Also: *Oracle Database Reference* for information about the V\$ views

Restoring Datafiles and Archived Redo Logs

This section contains the following topics:

- [Restoring Datafiles with Operating System Utilities](#)
- [Restoring Archived Redo Logs with Operating System Utilities](#)

Restoring Datafiles with Operating System Utilities

If a media failure permanently damages one or more datafiles of a database, then you must restore backups of these datafiles before you can recover the damaged files. If you cannot restore a damaged datafile to its original location (for example, you must replace a disk, so you restore the files to an alternate disk), then you must indicate the new locations of these files to the control file.

If you are restoring a database file on a raw disk or partition, then the procedure is basically the same as when restoring to a file on a file system. However, be aware of the naming conventions for files on raw devices (which differ depending on the operating system), and use an operating system utility that supports raw devices.

See Also: ["Making User-Managed Backups to Raw Devices"](#) on page 17-15 for an overview of considerations when backing up and restoring files on raw devices

To restore backup datafiles to their default location:

1. Determine which datafiles to recover by using the techniques described in ["Determining Which Datafiles Require Recovery"](#) on page 18-3.

2. If the database is open, then take the tablespaces containing the inaccessible datafiles offline. For example, enter:

```
ALTER TABLESPACE users OFFLINE IMMEDIATE;
```

3. Copy backups of the damaged datafiles to their default location using operating system commands. For example, to restore users01.dbf you might issue:

```
% cp /disk2/backup/users01.dbf $ORACLE_HOME/oradata/trgt/users01.dbf
```

4. Recover the affected tablespace. For example, enter:

```
RECOVER TABLESPACE users
```

5. Bring the recovered tablespace online. For example, enter:

```
ALTER TABLESPACE users ONLINE;
```

Restoring Archived Redo Logs with Operating System Utilities

All archived redo logs generated between the time a restored backup was created and the target recovery time are required for the pending recovery. The archived logs will eventually need to be on disk so that they are available to the database.

To restore necessary archived redo logs:

1. To determine which archived redo log files are needed, query V\$ARCHIVED_LOG and V\$RECOVERY_LOG. V\$ARCHIVED_LOG lists filenames for all archived logs. V\$RECOVERY_LOG lists only the archived redo logs that the database needs to perform media recovery. It also includes the probable names of the files, using LOG_ARCHIVE_FORMAT.

Note:

V\$RECOVERY_LOG is only populated when media recovery is required for a datafile. Hence, this view is not useful in the case of a planned recovery, such as recovery from a user error.

If a datafile requires recovery, but no backup of the datafile exists, then you need all redo generated starting from the time when the datafile was added to the database.

2. If space is available, then restore the required archived redo log files to the location specified by LOG_ARCHIVE_DEST_1. The database locates the correct log automatically when required during media recovery. For example, enter:

```
% cp /disk2/arch/* $ORACLE_HOME/oradata/trgt/arch
```

3. If sufficient space is not available at the location indicated by the archiving destination initialization parameter, restore some or all of the required archived redo log files to an alternate location. Specify the location before or during media recovery using the LOGSOURCE parameter of the SET statement in SQL*Plus or the RECOVER . . . FROM parameter of the ALTER DATABASE statement in SQL. For example, enter:

```
SET LOGSOURCE /tmp # set location using SET statement
DATABASE RECOVER FROM '/tmp'; # set location in RECOVER statement
```

4. After an archived log is applied, and after making sure that a copy of each archived log group still exists in offline storage, delete the restored copy of the archived redo log file to free disk space. For example:

```
% rm /tmp/*.dbf
```

See Also: *Oracle Database Reference* for more information about the data dictionary views, and ["About User-Managed Media Recovery"](#) on page 18-11 for an overview of log application during media recovery

Restoring Control Files

This section contains the following topics:

- [Restore Lost Copy of a Multiplexed Control File](#)
- [Restore Control File from Backup After Loss of All Current Control Files](#)
- [Create New Control File After Losing All Current and Backup Control Files](#)

Restore Lost Copy of a Multiplexed Control File

Use the following procedures to recover a database if a permanent media failure has damaged one or more control files of a database and at least one control file has *not* been damaged by the media failure.

Copying a Multiplexed Control File to a Default Location

If the disk and file system containing the lost control file are intact, then you can simply copy one of the intact control files to the location of the missing control file. In this case, you do not have to alter the `CONTROL_FILES` initialization parameter setting.

To replace a damaged control file by copying a multiplexed control file:

1. If the instance is still running, then shut it down:

```
SHUTDOWN ABORT
```
2. Correct the hardware problem that caused the media failure. If you cannot repair the hardware problem quickly, then proceed with database recovery by restoring damaged control files to an alternative storage device, as described in ["Copying a Multiplexed Control File to a Nondefault Location"](#) on page 18-6.
3. Use an intact multiplexed copy of the database's current control file to copy over the damaged control files. For example, to replace `bad_cf.f` with `good_cf.f`, you might enter:

```
% cp /oracle/good_cf.f /oracle/dbs/bad_cf.f
```

4. Start a new instance and mount and open the database. For example, enter:

```
STARTUP
```

Copying a Multiplexed Control File to a Nondefault Location

Assuming that the disk and file system containing the lost control file are not intact, then you cannot copy one of the good control files to the location of the missing control file. In this case, you must alter the `CONTROL_FILES` initialization parameter to indicate a new location for the missing control file.

To restore a control file to a nondefault location:

1. If the instance is still running, then shut it down:

```
SHUTDOWN ABORT
```

2. If you cannot correct the hardware problem that caused the media failure, then copy the intact control file to alternative locations. For example, to copy a good version of `control01.dbf` to a new disk location you might issue:

```
% cp $ORACLE_HOME/oradata/trgt/control01.dbf /new_disk/control01.dbf
```

3. Edit the parameter file of the database so that the `CONTROL_FILES` parameter reflects the current locations of all control files and excludes all control files that were not restored. Assume the initialization parameter file contains:

```
CONTROL_FILES='/oracle/oradata/trgt/control01.dbf','/bad_disk/control02.dbf'
```

Then, you can edit it as follows:

```
CONTROL_FILES='/oracle/oradata/trgt/control01.dbf','/new_disk/control02.dbf'
```

4. Start a new instance and mount and open the database. For example:

```
STARTUP
```

Restore Control File from Backup After Loss of All Current Control Files

Use the following procedures to restore a backup control file if a permanent media failure has damaged all control files of a database and you have a backup of the control file. When a control file is inaccessible, you can start the instance, but not mount the database. If you attempt to mount the database when the control file is unavailable, then you receive this error message:

```
ORA-00205: error in identifying control file, check alert log for more info
```

You cannot mount and open the database until the control file is accessible again. If you restore a backup control file, then you must open `RESETLOGS`.

As indicated in the following table, the procedure for restoring the control file depends on whether the online redo logs are available.

Table 18–1 Scenarios When Control Files Are Lost

Status of Online Logs	Status of Datafiles	Restore Procedure
Available	Current	If the online logs contain redo necessary for recovery, then restore a backup control file and apply the logs during recovery. You must specify the filename of the online logs containing the changes in order to open the database. After recovery, open <code>RESETLOGS</code> .
Unavailable	Current	If the online logs contain redo necessary for recovery, then re-create the control file. Because the online redo logs are inaccessible, open <code>RESETLOGS</code> (when the online logs are accessible it is not necessary to <code>OPEN RESETLOGS</code> after recovery with a created control file).
Available	Backup	Restore a backup control file, perform complete recovery, and then open <code>RESETLOGS</code> .
Unavailable	Backup	Restore a backup control file, perform incomplete recovery, and then open <code>RESETLOGS</code> .

Restoring a Backup Control File to the Default Location

If possible, restore the control file to its original location. In this way, you avoid having to specify new control file locations in the initialization parameter file.

To restore a backup control file to its default location:

1. If the instance is still running, shut it down:

```
SHUTDOWN ABORT
```

2. Correct the hardware problem that caused the media failure.
3. Restore the backup control file to all locations specified in the `CONTROL_FILES` parameter. For example, if `ORACLE_HOME/oradata/trgt/control01.dbf` and `ORACLE_HOME/oradata/trgt/control02.dbf` are the control file locations listed in the server parameter file, then use an operating system utility to restore the backup control file to these locations:

```
% cp /backup/control01.dbf ORACLE_HOME/oradata/trgt/control01.dbf
% cp /backup/control02.dbf ORACLE_HOME/oradata/trgt/control02.dbf
```

4. Start a new instance and mount the database. For example, enter:

```
STARTUP MOUNT
```

5. Begin recovery by executing the `RECOVER` command with the `USING BACKUP CONTROLFILE` clause. Specify `UNTIL CANCEL` if you are performing incomplete recovery. For example, enter:

```
RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL
```

6. Apply the prompted archived logs. If you then receive another message saying that the required archived log is missing, it probably means that a necessary redo record is located in the online redo logs. This situation can occur when unarchived changes were located in the online logs when the instance crashed.

For example, assume that you see the following:

```
ORA-00279: change 55636 generated at 11/08/2002 16:59:47 needed for thread 1
ORA-00289: suggestion : /oracle/work/arc_dest/arcr_1_111.arc
ORA-00280: change 55636 for thread 1 is in sequence #111
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

You can specify the name of an online redo log and press Enter (you may have to try this a few times until you find the correct log):

```
ORACLE_HOME/oradata/redo01.dbf
Log applied.
Media recovery complete.
```

If the online logs are inaccessible, then you can cancel recovery without applying them. If all datafiles are current, and if redo in the online logs is required for recovery, then you cannot open the database without applying the online logs. If the online logs are inaccessible, then you must re-create the control file, using the procedure described in ["Create New Control File After Losing All Current and Backup Control Files"](#) on page 18-9.

7. Open the database with the `RESETLOGS` option after finishing recovery:

```
ALTER DATABASE OPEN RESETLOGS;
```

Restoring a Backup Control File to a Nondefault Location

If you cannot restore the control file to its original place because the media damage is too severe, then you must specify new control file locations in the server parameter file. A valid control file must be available in all locations specified by the `CONTROL_FILES` initialization parameter. If not, then the database prevents you from mounting the database.

To restore a control file to a nondefault location:

Follow the steps in ["Restoring a Backup Control File to the Default Location"](#) on page 18-8, except after step 2 add the following step:

Edit all locations specified in the `CONTROL_FILES` initialization parameter to reflect the new control file locations. For example, if the control file locations listed in the server parameter file are as follows, and both locations are inaccessible:

```
CONTROL_FILES='/oracle/oradata/trgt/control01.dbf',
              '/oracle/oradata/trgt/control01.dbf'
```

Then, you can edit the initialization parameter file as follows:

```
CONTROL_FILES='/good_disk/control01.dbf', '/good_disk/control02.dbf'
```

Create New Control File After Losing All Current and Backup Control Files

If all control files have been lost in a permanent media failure, but all online redo log members remain intact, then you can recover the database after creating a new control file. The advantage of this tactic is that you are *not* required to open the database with the `RESETLOGS` option.

Depending on the existence and currency of a control file backup, you have the options listed in the following table for generating the text of the `CREATE CONTROLFILE` statement. Note that changes to the database are recorded in the `alert_SID.log`, so check this log when deciding which option to choose.

Table 18–2 Options for Creating the Control File

If you . . .	Then . . .
Executed <code>ALTER DATABASE BACKUP CONTROLFILE TO TRACE NORESETLOGS</code> after you made the last structural change to the database, and if you have saved the SQL command trace output	Use the <code>CREATE CONTROLFILE</code> statement from the trace output as-is.
Performed your most recent execution of <code>ALTER DATABASE BACKUP CONTROLFILE TO TRACE</code> before you made a structural change to the database	Edit the output of <code>ALTER DATABASE BACKUP CONTROLFILE TO TRACE</code> to reflect the change. For example, if you recently added a datafile to the database, then add this datafile to the <code>DATAFILE</code> clause of the <code>CREATE CONTROLFILE</code> statement.
Backed up the control file with the <code>ALTER DATABASE BACKUP CONTROLFILE TO filename</code> statement (not the <code>TO TRACE</code> option)	Use the control file copy to obtain SQL output. Create a temporary database instance, mount the backup control file, and then run <code>ALTER DATABASE BACKUP CONTROLFILE TO TRACE NORESETLOGS</code> . If the control file copy predated a recent structural change, then edit the trace to reflect the change.
Do not have a control file backup in either <code>TO TRACE</code> format or <code>TO filename</code> format	Execute the <code>CREATE CONTROLFILE</code> statement manually (refer to <i>Oracle Database SQL Reference</i>).

Note: If your character set is not the default US7ASCII, then you must specify the character set as an argument to the CREATE CONTROLFILE statement. The database character set is written to the alert log at startup. The character set information is also recorded in the BACKUP CONTROLFILE TO TRACE output.

To create a new control file:

1. Start the database in NOMOUNT mode. For example, enter:

```
STARTUP NOMOUNT
```

2. Create the control file with the CREATE CONTROLFILE statement, specifying the NORESETLOGS option (refer to [Table 18-2](#) for options). The following example assumes that the character set is the default US7ASCII:

```
CREATE CONTROLFILE REUSE DATABASE SALES NORESETLOGS ARCHIVELOG
  MAXLOGFILES 32
  MAXLOGMEMBERS 2
  MAXDATAFILES 32
  MAXINSTANCES 16
  MAXLOGHISTORY 1600
LOGFILE
  GROUP 1 (
    '/diska/prod/sales/db/log1t1.dbf',
    '/diskb/prod/sales/db/log1t2.dbf'
  ) SIZE 100K
  GROUP 2 (
    '/diska/prod/sales/db/log2t1.dbf',
    '/diskb/prod/sales/db/log2t2.dbf'
  ) SIZE 100K,
DATAFILE
  '/diska/prod/sales/db/database1.dbf',
  '/diskb/prod/sales/db/filea.dbf';
```

After creating the control file, the instance mounts the database.

3. Recover the database as normal (*without* specifying the USING BACKUP CONTROLFILE clause):

```
RECOVER DATABASE
```

4. Open the database after recovery completes (RESETLOGS option not required):

```
ALTER DATABASE OPEN;
```

5. Immediately back up the control file. The following SQL statement backs up a database's control file to /backup/control01.dbf:

```
ALTER DATABASE BACKUP CONTROLFILE TO '/backup/control01.dbf' REUSE;
```

See Also: ["Backing Up the Control File to a Trace File"](#) on page 17-11, and ["Re-Creating Datafiles When Backups Are Unavailable: Scenario"](#) on page 19-3

About User-Managed Media Recovery

To perform recovery, Oracle recommends that you use the `RECOVER SQL` statement in SQL*Plus. You can also use the SQL statement `ALTER DATABASE RECOVER`, but the `RECOVER` statement is simpler in most cases.

Preconditions of Performing User-Managed Recovery

To start any type of media recovery, you must adhere to the following restrictions:

- You must have administrator privileges.
- All recovery sessions must be compatible.
- One session cannot start complete media recovery while another performs incomplete media recovery.
- You cannot start media recovery if you are connected to the database through a shared server process.

Applying Logs Automatically with the RECOVER Command

Oracle recommends that you use the SQL*Plus `RECOVER` command rather than the `ALTER DATABASE RECOVER` statement to perform media recovery. In almost all cases, the SQL*Plus method is easier.

When using SQL*Plus to perform media recovery, the easiest strategy is to perform automatic recovery. Automatic recovery initiates recovery without manually prompting SQL*Plus to apply each individual archived log.

When using SQL*Plus, you have two options for automating the application of the default filenames of archived redo logs needed during recovery:

- Issuing `SET AUTORECOVERY ON` before issuing the `RECOVER` command
- Specifying the `AUTOMATIC` keyword as an option of the `RECOVER` command

In either case, no interaction is required when you issue the `RECOVER` command if the necessary files are in the correct locations with the correct names. The filenames used when you use automatic recovery are derived from the concatenated values of `LOG_ARCHIVE_FORMAT` with `LOG_ARCHIVE_DEST_n`, where `n` is the highest value among all enabled, local destinations.

For example, assume the following initialization parameter settings are in effect in the database instance:

```
LOG_ARCHIVE_DEST_1 = "LOCATION=/arc_dest/loc1/"
LOG_ARCHIVE_DEST_2 = "LOCATION=/arc_dest/loc2/"
LOG_ARCHIVE_DEST_STATE_1 = DEFER
LOG_ARCHIVE_DEST_STATE_2 = ENABLE
LOG_ARCHIVE_FORMAT = arch_%t_%s_%r.arc
```

In this case, SQL*Plus automatically suggests the filename `/arc_dest/loc2/arch_%t_%s_%r.arc` (where `%t` is the thread, `%s` is the sequence and `%r` is the resetlogs ID).

If you run `SET AUTORECOVERY OFF`, which is the default option, then you must enter the filenames manually, or accept the suggested default filename by pressing the Enter key.

Automating Recovery with SET AUTORECOVERY

Run the `SET AUTORECOVERY ON` command to enable on automatic recovery.

To automate the recovery using SET AUTORECOVERY:

1. Restore a backup of the offline datafiles. This example restores an inconsistent backup of all datafiles with an operating system utility:

```
% cp /backup/datafiles/*.dbf $ORACLE_HOME/oradata/trgt/
```

2. Ensure the database is mounted. For example, if the database is shut down, run:

```
STARTUP MOUNT
```

3. Enable automatic recovery. For example, in SQL*Plus run:

```
SET AUTORECOVERY ON
```

4. Recover the desired datafiles. This example recovers the whole database:

```
RECOVER DATABASE
```

The database automatically suggests and applies the necessary archived logs.

5. Open the database. For example:

```
ALTER DATABASE OPEN;
```

Note: After issuing the SQL*Plus RECOVER command, you can view all files that have been considered for recovery in the V\$RECOVERY_FILE_STATUS view. You can access status information for each file in the V\$RECOVERY_STATUS view. These views are not accessible after you terminate the recovery session.

Automating Recovery with the AUTOMATIC Option of the RECOVER Command

Besides using SET AUTORECOVERY to turn on automatic recovery, you can also simply specify the AUTOMATIC keyword in the RECOVER command.

To automate the recovery with the RECOVER AUTOMATIC command:

1. Restore a backup of the offline datafiles. This example restores a backup of all datafiles:

```
% cp /backup/datafiles/*.dbf $ORACLE_HOME/oradata/trgt/
```

2. Ensure the database is mounted. For example, if the database is shut down, run:

```
STARTUP MOUNT
```

3. Recover the desired datafiles by specifying the AUTOMATIC keyword. This example performs automatic recovery on the whole database:

```
RECOVER AUTOMATIC DATABASE
```

The database automatically suggests and applies the necessary archived logs.

4. Open the database. For example:

```
ALTER DATABASE OPEN;
```

If you use an Oracle Real Application Clusters configuration, and if you are performing incomplete recovery or using a backup control file, then the database can only compute the name of the first archived redo log file from the *first* redo thread. You may have to manually apply the first log file from the other redo threads. After

the first log file in a given thread has been supplied, the database can suggest the names of the subsequent logs in this thread.

See Also: Your platform-specific Oracle documentation for examples of log file application

Recovering When Archived Logs Are in the Default Location

Recovering when the archived logs are in their default location is the simplest case. As a log is needed, the database suggests the filename. If you are running nonautomatic media recovery with SQL*Plus, then the output is displayed in this format:

```
ORA-00279: Change ##### generated at DD/MM/YY HH:MM:SS needed for thread#
ORA-00289: Suggestion : logfile
ORA-00280: Change ##### for thread # is in sequence #
Specify log: [<RET> for suggested | AUTO | FROM logsource | CANCEL ]
```

For example, SQL*Plus displays output similar to the following:

```
ORA-00279: change 53577 generated at 11/26/02 19:20:58 needed for thread 1
ORA-00289: suggestion : /oracle/oradata/trgt/arch/arcr_1_802.arc
ORA-00280: change 53577 for thread 1 is in sequence #802
Specify log: [<RET> for suggested | AUTO | FROM logsource | CANCEL ]
```

Similar messages are returned when you use an ALTER DATABASE . . . RECOVER statement. However, no prompt is displayed.

The database constructs suggested archived log filenames by concatenating the current values of the initialization parameters LOG_ARCHIVE_DEST_1 (where *n* is the highest value among all enabled, local destinations) and LOG_ARCHIVE_FORMAT and using log history data from the control file. The following are possible settings:

```
LOG_ARCHIVE_DEST_1 = 'LOCATION = /oracle/oradata/trgt/arch/'
LOG_ARCHIVE_FORMAT = arcr_%t_%s.arc
```

```
SELECT NAME FROM V$ARCHIVED_LOG;
```

```
NAME
```

```
-----
/oracle/oradata/trgt/arch/arcr_1_467.arc
/oracle/oradata/trgt/arch/arcr_1_468.arc
/oracle/oradata/trgt/arch/arcr_1_469.arc
```

Thus, if all the required archived log files are mounted at the LOG_ARCHIVE_DEST_1 destination, and if the value for LOG_ARCHIVE_FORMAT is never altered, then the database can suggest and apply log files to complete media recovery automatically.

Recovering When Archived Logs Are in a Nondefault Location

Performing media recovery when archived logs are not in their default location adds an extra step. You have the following mutually exclusive options:

- Edit the LOG_ARCHIVE_DEST_1 parameter that specifies the location of the archived redo logs, then recover as usual.
- Use the SET statement in SQL*Plus to specify the nondefault log location before recovery, or the LOGFILE parameter of the RECOVER command

Resetting the Archived Log Destination

You can edit the initialization parameter file or issue `ALTER SYSTEM` statements to change the default location of the archived redo logs.

To change the default archived log location before recovery:

1. Use an operating system utility to restore the archived logs to a nondefault location. For example, enter:

```
% cp /backup/arch/* /tmp/
```

2. Change the value for the archive log parameter to the nondefault location. You can issue `ALTER SYSTEM` statements while the instance is started, or edit the initialization parameter file and then start the database instance. For example, while the instance is shut down edit the parameter file as follows:

```
LOG_ARCHIVE_DEST_1 = 'LOCATION=/tmp/'  
LOG_ARCHIVE_FORMAT = arcr_%t_%s.arc
```

3. Using SQL*Plus, start a new instance by specifying the edited initialization parameter file, and then mount the database. For example, enter:

```
STARTUP MOUNT
```

4. Begin media recovery as usual. For example, enter:

```
RECOVER DATABASE
```

Overriding the Archived Log Destination

In some cases, you may want to override the current setting for the archiving destination parameter as a source for redo log files.

To recover archived logs in a nondefault location with SET LOGSOURCE:

1. Using an operating system utility, copy the archived redo logs to an alternative location. For example, enter:

```
% cp $ORACLE_HOME/oradata/trgt/arch/* /tmp
```

2. Specify the alternative location within SQL*Plus for the recovery operation. Use the `LOGSOURCE` parameter of the `SET` statement or the `RECOVER . . . FROM` clause of the `ALTER DATABASE` statement. For example, start SQL*Plus and run:

```
SET LOGSOURCE "/tmp"
```

3. Recover the offline tablespace. For example, to recover the offline tablespace `users` do the following:

```
RECOVER AUTOMATIC TABLESPACE users
```

4. Alternatively, you can avoid running `SET LOGSOURCE` and simply run:

```
RECOVER AUTOMATIC TABLESPACE users FROM "/tmp"
```

Note: Overriding the redo log source does not affect the archive redo log destination for online redo logs groups being archived.

Responding to Unsuccessful Application of Redo Logs

If you are using SQL*Plus's recovery options (not SQL statements), then each time the database successfully applies a redo log file, the following message is returned:

```
Log applied.
```

You are then prompted for the next log in the sequence or, if the most recently applied log is the last required log, terminates recovery.

If the suggested file is incorrect or you provide an incorrect filename, then the database returns an error message. For example, you may see something like:

```
ORA-00308: cannot open archived log "/oracle/oradata/trgt/arch/arcr_1_811.arc"
ORA-27037: unable to obtain file status
SVR4 Error: 2: No such file or directory
Additional information: 3
```

Recovery cannot continue until the required redo log is applied. If the database returns an error message after supplying a log filename, then the following responses are possible.

Error	Possible Cause	Solution
ORA-27037: unable to obtain file status	Entered wrong filename. Log is missing.	Reenter correct filename. Restore backup archived redo log.
ORA-27047: unable to read the header block of file	The log may have been partially written or become corrupted.	If you can locate an uncorrupted or complete log copy, then apply the intact copy and continue recovery. If no copy of the log exists and you know the time of the last valid redo entry, then you use incomplete recovery. Restore backups and restart recovery.

Interrupting User-Managed Media Recovery

If you start media recovery and must then interrupt it, for example, because a recovery operation must end for the night and resume the next morning, then take either of the following actions:

- Enter the word `CANCEL` when prompted for a redo log file.
- Use your operating system's interrupt signal if you must terminate when recovering an individual datafile, or when automated recovery is in progress.

After recovery is canceled, you can resume it later with the `RECOVER` command. Recovery resumes where it left off when it was canceled.

Performing Complete User-Managed Media Recovery

When you perform complete recovery, you recover the backups to the current SCN. You can either recover the whole database at once or recover individual tablespaces or datafiles. Because you do not have to open the database with the `RESETLOGS` option after complete recovery as you do after incomplete recovery, you have the option of recovering some datafiles at one time and the remaining datafiles later.

This section describes the steps necessary to complete media recovery operations, and includes the following topics:

- [Performing Closed Database Recovery](#)
- [Performing Datafile Recovery in an Open Database](#)

See Also: *Oracle Database Backup and Recovery Basics* for basic information about media recovery concepts, which apply in both user-managed and RMAN-based backup and recovery.

See Also: *Oracle Database Backup and Recovery Basics* to familiarize yourself with fundamental recovery concepts and strategies

Performing Closed Database Recovery

This section describes steps to perform complete recovery while the database is not open. You can recover either all damaged datafiles in one operation, or perform individual recovery of each damaged datafile in separate operations.

Perform the media recovery in the following stages:

1. Prepare for closed database recovery as described in "[Preparing for Closed Database Recovery](#)" on page 18-16.
2. Restore the necessary files as described in "[Restoring Backups of the Damaged or Missing Files](#)" on page 18-16.
3. Recover the restored datafiles as described in "[Recovering the Database](#)" on page 18-17.

Preparing for Closed Database Recovery

In this stage, you shut down the instance and inspect the media device that is causing the problem.

To prepare for closed database recovery:

1. If the database is open, then shut it down. For example:

```
SHUTDOWN IMMEDIATE
```
2. If you are recovering from a media error, then correct it if possible. If the hardware problem that caused the media failure was temporary, and if the data was undamaged (for example, a disk or controller power failure), then no media recovery is required: simply start the database and resume normal operations. If you cannot repair the problem, then proceed to the next stage.

Restoring Backups of the Damaged or Missing Files

In this stage, you restore all necessary backups.

To restore the necessary files:

1. Determine which datafiles to recover by using the techniques described in "[Determining Which Datafiles Require Recovery](#)" on page 18-3.
2. If the files are permanently damaged, then identify the most recent backups for the damaged files. Restore *only* the datafiles damaged by the media failure: do not restore any undamaged datafiles or any online redo log files.

For example, if `ORACLE_HOME/oradata/trgt/users01.dbf` is the only damaged file, then you may determine that `/backup/users01_10_24_02.dbf` is the most recent backup of this file. If you do not have a backup of a specific

datafile, then you may be able to create an empty replacement file that can be recovered.

3. Use an operating system utility to restore the files to their default location or to a new location. Restore the necessary files as described in ["Restoring Datafiles and Archived Redo Logs"](#) on page 18-4. For example, a UNIX user restoring `users01.dbf` to its default location might enter:

```
% cp /backup/users01_10_24_02.dbf $ORACLE_HOME/oradata/trgt/users01.dbf
```

Use the following guidelines when determining where to restore datafile backups.

If . . .	Then . . .
The hardware problem is repaired and you can restore the datafiles to their default locations	Restore the datafiles to their default locations and begin media recovery.
The hardware problem persists and you cannot restore datafiles to their original locations	Restore the datafiles to an alternative storage device. Indicate the new location of these files in the control file with <code>ALTER DATABASE RENAME FILE</code> . Use the operation described in "Renaming and Relocating Datafiles" in the <i>Oracle Database Administrator's Guide</i> , as necessary.

Recovering the Database

In the final stage, you recover the datafiles that you have restored.

To recover the restored datafiles:

1. Connect to the database with administrator privileges, then start a new instance and mount, but do not open, the database. For example, enter:

```
STARTUP MOUNT
```

2. Obtain the datafile names and statuses of all datafiles by checking the list of datafiles that normally accompanies the current control file or querying the `V$DATAFILE` view. For example, enter:

```
SELECT NAME, STATUS FROM V$DATAFILE;
```

3. Ensure that all datafiles of the database are online. All datafiles of the database requiring recovery must be online unless an offline tablespace was taken offline normally or is part of a read-only tablespace. For example, to guarantee that a datafile named `/oracle/dbs/tbs_10.f` is online, enter the following:

```
ALTER DATABASE DATAFILE '/oracle/dbs/tbs_10.f' ONLINE;
```

If a specified datafile is already online, then the database ignores the statement. If you prefer, create a script to bring all datafiles online at once as in the following:

```
SPOOL onlineall.sql
SELECT 'ALTER DATABASE DATAFILE ''' || name || ''' ONLINE;' FROM V$DATAFILE;
SPOOL OFF
```

```
SQL> @onlineall
```

4. Issue the statement to recover the database, tablespace, or datafile. For example, enter one of the following `RECOVER` command:

```
RECOVER DATABASE # recovers whole database
RECOVER TABLESPACE users # recovers specific tablespace
```

```
RECOVER DATAFILE '?/oradata/trgt/users01.dbf'; # recovers specific datafile
```

Follow these guidelines when deciding which statement to execute:

If you want to . . .	Then . . .
Recover all damaged files in one step	Execute <code>RECOVER DATABASE</code>
Recover an individual tablespace	Execute <code>RECOVER TABLESPACE</code>
Recover an individual damaged datafile	Execute <code>RECOVER DATAFILE</code>
Prevent parallelization of recovery on multiple-CPU systems	Refer to " Controlling Parallel Media Recovery " on page 18-29

- If you choose not to automate the application of archived logs, then you must accept or reject each prompted log. If you automate recovery, then the database applies the logs automatically. Recovery continues until all required archived and online redo logs have been applied to the restored datafiles.

- The database notifies you when media recovery is complete:

```
Media recovery complete.
```

If no archived redo log files are required for complete media recovery, then the database applies all necessary online redo log files and terminates recovery.

- After recovery terminates, open the database for use:

```
ALTER DATABASE OPEN;
```

See Also: "[About User-Managed Media Recovery](#)" on page 18-11 for more information about applying redo log files

Performing Datafile Recovery in an Open Database

It is possible for a media failure to occur while the database remains open, leaving the undamaged datafiles online and available for use. Damaged datafiles—but not the tablespaces that contain them—are automatically taken offline if the database writer is unable to write to them. Queries that cannot read damaged files return errors, but the datafiles are not taken offline because of the failed queries. For example, you may run a query and see output such as:

```
ERROR at line 1:
ORA-01116: error in opening database file 3
ORA-01110: data file 11: '/oracle/oradata/trgt/cwmlite02.dbf'
ORA-27041: unable to open file
SVR4 Error: 2: No such file or directory
Additional information: 3
```

The procedure in this section cannot be used to perform complete media recovery on the datafiles of the `SYSTEM` tablespace while the database is open. If the media failure damages datafiles of the `SYSTEM` tablespace, then the database automatically shuts down.

Perform media recovery in these stages:

- Prepare the database for recovery by making sure it is open and taking the tablespaces requiring recovery offline, as described in "[Preparing for Open Database Recovery](#)" on page 18-19.

2. Restore the necessary files in the affected tablespaces as described in ["Restoring Backups of the Inaccessible Datafiles"](#) on page 18-19.
3. Recover the affected tablespaces as described in ["Recovering Offline Tablespaces in an Open Database"](#) on page 18-19.

See Also:

- ["Determining Which Datafiles Require Recovery"](#) on page 18-3
- ["Performing Closed Database Recovery"](#) on page 18-16 for procedures for proceeding with complete media recovery of the SYSTEM tablespace

Preparing for Open Database Recovery

In this stage, you take affected tablespaces offline and inspect the media device that is causing the problem.

To prepare for datafile recovery when the database is open:

1. If the database is open when you discover that recovery is required, take all tablespaces containing damaged datafiles offline. For example, if tablespace `users` and `tools` contain damaged datafiles, enter:

```
ALTER TABLESPACE users OFFLINE TEMPORARY;
ALTER TABLESPACE tools OFFLINE TEMPORARY;
```

2. Correct the hardware problem that caused the media failure. If the hardware problem cannot be repaired quickly, proceed with database recovery by restoring damaged files to an alternative storage device.

Restoring Backups of the Inaccessible Datafiles

In this stage, you restore all necessary backups in the offline tablespaces.

To restore datafiles in an open database:

1. If files are permanently damaged, then restore the most recent backup files of *only* the datafiles damaged by the media failure. Do not restore undamaged datafiles, online redo logs, or control files. If the hardware problem is fixed and the datafiles can be restored to their original locations, then do so. Otherwise, restore the datafiles to an alternative storage device.

Note: In some circumstances, if you do not have a backup of a specific datafile, you can use `ALTER DATABASE CREATE DATAFILE` to create an empty replacement file that is recoverable.

2. If you restored one or more damaged datafiles to alternative locations, update the control file of the database to reflect the new datafile names. For example, to change the filename of the datafile in tablespace `users` you might enter:

```
ALTER DATABASE RENAME FILE '?/oradata/trgt/users01.dbf' TO
'/disk2/users01.dbf';
```

See Also: *Oracle Database SQL Reference* for more information about `ALTER DATABASE RENAME FILE`

Recovering Offline Tablespaces in an Open Database

In the final stage, you recover the datafiles in the offline tablespaces.

To recover offline tablespaces in an open database:

1. Connect to the database with administrator privileges, and start offline tablespace recovery of all damaged datafiles in one or more offline tablespaces using one step. For example, recover `users` and `tools`:

```
RECOVER TABLESPACE users, tools;
```

2. The database begins the roll forward phase of media recovery by applying the necessary redo logs (archived and online) to reconstruct the restored datafiles. Unless the applying of files is automated with `RECOVER AUTOMATIC` or `SET AUTORECOVERY ON`, the database prompts for each required redo log file.

Recovery continues until all required archived logs have been applied to the datafiles. The online redo logs are then automatically applied to the restored datafiles to complete media recovery. If no archived redo logs are required for complete media recovery, then the database does not prompt for any. Instead, all necessary online redo logs are applied, and media recovery is complete.

3. When the damaged tablespaces are recovered up to the moment that media failure occurred, bring the offline tablespaces online. For example, to bring tablespaces `users` and `tools` online, issue the following statements:

```
ALTER TABLESPACE users ONLINE;  
ALTER TABLESPACE tools ONLINE;
```

See Also: *Oracle Database Administrator's Guide* for more information about creating datafiles

Performing User-Managed Database Point-in-Time Recovery

- [Preparing for Incomplete Recovery](#)
- [Restoring Datafiles Before Performing Incomplete Recovery](#)
- [Performing Cancel-Based Incomplete Recovery](#)
- [Performing Time-Based or Change-Based Incomplete Recovery](#)

Note: If your database is affected by seasonal time changes (for example, daylight savings time), then you may experience a problem if a time appears twice in the redo log and you want to recover to the second, or later time. To handle time changes, perform cancel-based or change-based recovery.

Preparing for Incomplete Recovery

In this phase, you examine the source of the media problem.

To prepare for incomplete recovery:

1. If you are uncertain about performing incomplete recovery, then back up the whole database—all datafiles, a control file, and the parameter files—as a precautionary measure in case an error occurs during the recovery procedure.
2. If the database is still open and incomplete media recovery is necessary, then terminate the instance:

```
SHUTDOWN ABORT
```

3. If a media failure occurred, correct the hardware problem that caused the failure. If the hardware problem cannot be repaired quickly, then proceed with database recovery by restoring damaged files to an alternative storage device.

Restoring Datafiles Before Performing Incomplete Recovery

In this phase, you restore a whole database backup.

To restore the files necessary for cancel-based recovery and bring them online:

1. If the current control files do not match the physical structure of the database at the intended time of recovery, then restore a backup control file as described in ["Restore Control File from Backup After Loss of All Current Control Files"](#) on page 18-7. The restored control file should reflect the database's physical file structure at the point at which incomplete media recovery should finish. To determine which control file backup to use:
 - Review the list of files that corresponds to the current control file and each control file backup to determine the correct control file to use.
 - If necessary, replace all current control files of the database with the correct control file backup.
 - Alternatively, create a new control file to replace the missing one.

Note: If you are unable to restore a control file backup to one of the CONTROL_FILES locations, then edit the initialization parameter file so that this CONTROL_FILES location is removed.

2. Restore backups of *all* datafiles of the database. All backups used to replace existing datafiles must have been taken before the intended time of recovery. For example, if you intend to recover to January 2 at 2:00 p.m., then restore all datafiles with backups completed before this time. Follow these guidelines:

If . . .	Then . . .
You do not have a backup of a datafile	Create an empty replacement file that can be recovered as described in "Restoring Backups of the Damaged or Missing Files" on page 18-16.
A datafile was added after the intended time of recovery	Do not restore a backup of this file because it will no longer be used for the database after recovery completes.
The hardware problem causing the failure has been solved and all datafiles can be restored to their default locations	Restore the files as described in "Restoring Datafiles and Archived Redo Logs" on page 18-4 and skip step 4 of this procedure.
A hardware problem persists	Restore damaged datafiles to an alternative storage device.

Note: Files in read-only tablespaces should be offline if you are using a control file backup. Otherwise, the recovery will try to update the headers of the read-only files.

1. Start SQL*Plus and connect to the database with administrator privilege, then start a new instance and mount the database:

```
STARTUP MOUNT
```

2. If one or more damaged datafiles were restored to alternative locations in step 2, then indicate the new locations of these files to the control file of the associated database. For example, enter:

```
ALTER DATABASE RENAME FILE '?/oradata/trgt/users01.dbf' TO
'/disk2/users01.dbf';
```

3. Obtain the datafile names and statuses of all datafiles by checking the list of datafiles that normally accompanies the current control file or querying the V\$DATAFILE view. For example, enter:

```
SELECT NAME,STATUS FROM V$DATAFILE;
```

4. Ensure that all datafiles of the database are online. All datafiles of the database requiring recovery must be online unless a tablespace was taken offline with the NORMAL option or is a read-only tablespace. For example, to guarantee that a datafile named ?/oradata/trgt/users01.dbf is online, enter the following:

```
ALTER DATABASE DATAFILE '?/oradata/trgt/users01.dbf' ONLINE;
```

If a specified datafile is already online, the statement has no effect. If you prefer, create a script to bring all datafiles online at once as in the following:

```
SPOOL onlineall.sql
SELECT 'ALTER DATABASE DATAFILE '''||name||''' ONLINE;' FROM V$DATAFILE;
SPOOL OFF
SQL> @onlineall
```

Performing Cancel-Based Incomplete Recovery

In cancel-based recovery, recovery proceeds by prompting you with the suggested filenames of archived redo log files. Recovery stops when you specify CANCEL instead of a filename or when all redo has been applied to the datafiles.

Cancel-based recovery is better than change-based or time-based recovery if you want to control which archived log terminates recovery. For example, you may know that you have lost all logs past sequence 1234, so you want to cancel recovery after log 1233 is applied.

You should perform cancel-based media recovery in these stages:

1. Prepare for recovery by backing up the database and correct any media failures as described in ["Preparing for Incomplete Recovery"](#) on page 18-20.
2. Restore backup datafiles as described in ["Restoring Datafiles Before Performing Incomplete Recovery"](#) on page 18-21. If you have a current control file, then do not restore a backup control file.
3. Perform media recovery on the restored database backup as described in the following procedure.

To perform cancel-based recovery:

1. Start SQL*Plus and connect to the database with administrator privileges, then start a new instance and mount the database:

```
STARTUP MOUNT
```

2. Begin cancel-based recovery by issuing the following command:

```
RECOVER DATABASE UNTIL CANCEL
```

If you are using a backup control file with this incomplete recovery, then specify the `USING BACKUP CONTROLFILE` option in the `RECOVER` command.

```
RECOVER DATABASE UNTIL CANCEL USING BACKUP CONTROLFILE
```

Note: If you fail to specify the `UNTIL` clause on the `RECOVER` command, then the database assumes a complete recovery and will not open until all redo is applied.

3. The database applies the necessary redo log files to reconstruct the restored datafiles. The database supplies the name it expects to find from `LOG_ARCHIVE_DEST_1` and requests you to stop or proceed with applying the log file. Note that if the control file is a backup, then you must supply the names of the online logs if you want to apply the changes in these logs.

Note: If you use a Real Application Clusters (RAC) configuration, and you are performing incomplete recovery or using a backup control file, then the database can only compute the name of the first archived redo log file from the *first* thread. The first redo log file from the other threads must be supplied by the user. After the first log file in a given thread has been supplied, the database can suggest the names of the subsequent log files in this thread.

4. Continue applying redo log files until the last log has been applied to the restored datafiles, then cancel recovery by executing the following command:

```
CANCEL
```

The database indicates whether recovery is successful. If you cancel before all the datafiles have been recovered to a consistent SCN and then try to open the database, you will get an `ORA-1113` error if more recovery is necessary. As explained in "[Determining Which Datafiles Require Recovery](#)" on page 18-3, you can query `V$RECOVER_FILE` to determine whether more recovery is needed, or if a backup of a datafile was not restored prior to starting incomplete recovery.

5. Open the database with the `RESETLOGS` option. You must always reset the logs after incomplete recovery or recovery with a backup control file. For example:

```
ALTER DATABASE OPEN RESETLOGS;
```

See Also: "[Opening the Database with the RESETLOGS Option](#)" on page 18-24

Performing Time-Based or Change-Based Incomplete Recovery

This section describes how to perform the time-based media recovery procedure in the following stages:

1. Prepare for recovery by backing up the database and correct any media failures as described in "[Preparing for Incomplete Recovery](#)" on page 18-20.

2. Restore backup datafiles as described in ["Restoring Datafiles Before Performing Incomplete Recovery"](#) on page 18-21. If you have a current control file, then do not restore a backup control file.
3. Perform media recovery with the following procedure.

To perform change-based or time-based recovery:

1. Issue the RECOVER DATABASE UNTIL statement to begin recovery. If recovering to an SCN, specify as a decimal number without quotation marks. For example, to recover through SCN 10034 issue:

```
RECOVER DATABASE UNTIL CHANGE 10034;
```

If recovering to a time, the time is always specified using the following format, delimited by single quotation marks: 'YYYY-MM-DD:HH24:MI:SS'. The following statement recovers the database up to a specified time:

```
RECOVER DATABASE UNTIL TIME '2000-12-31:12:47:30'
```

2. Apply the necessary redo log files to recover the restored datafiles. The database automatically terminates the recovery when it reaches the correct time, and returns a message indicating whether recovery is successful.

Note: Unless recovery is automated, the database supplies the name from LOG_ARCHIVE_DEST_1 and asks you to stop or proceed with after each log. If the control file is a backup, then after the archived logs are applied you must supply the names of the online logs.

3. Open the database in RESETLOGS mode. You must always reset the online logs after incomplete recovery or recovery with a backup control file. For example:

```
ALTER DATABASE OPEN RESETLOGS;
```

See Also: ["Opening the Database with the RESETLOGS Option"](#) on page 18-24

Opening the Database with the RESETLOGS Option

Whenever you perform incomplete recovery or recovery with a backup control file, you must reset the online logs when you open the database. The new version of the reset database is called a new **incarnation**.

This section contains the following topics:

- [About Opening with the RESETLOGS Option](#)
- [Executing the ALTER DATABASE OPEN Statements](#)
- [Checking the Alert Log After a RESETLOGS Operation](#)

About Opening with the RESETLOGS Option

The RESETLOGS option is always *required* after incomplete media recovery or recovery using a backup control file. Resetting the redo log does the following:

- Archives the current online redo logs (if they are accessible) and then erases the contents of the online redo logs and resets the log sequence number to 1. For example, if the current online redo logs are sequence 1000 and 1001 when you

open RESETLOGS, then the database archives logs 1000 and 1001 and then resets the online logs to sequence 1 and 2.

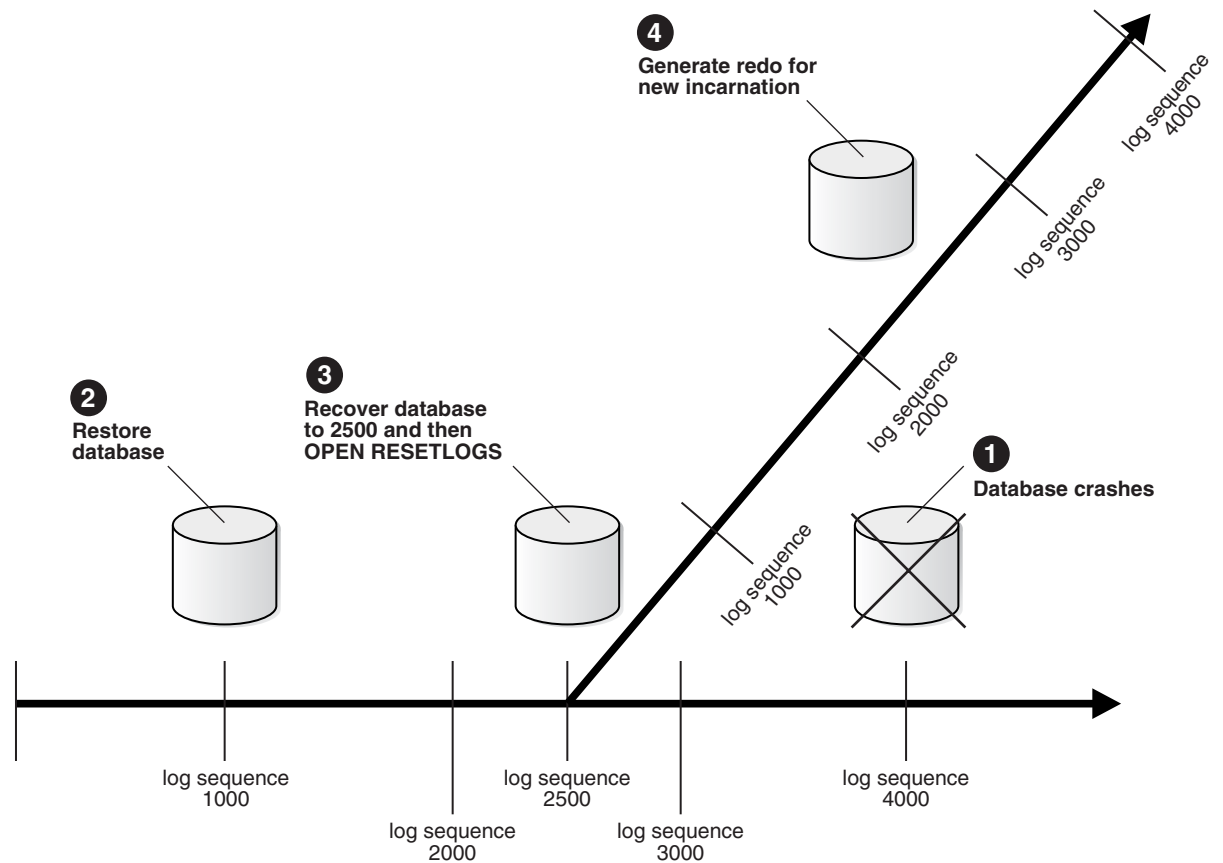
- Creates the online redo log files if they do not currently exist.
- Reinitializes the control file metadata about online redo logs and redo threads.
- Updates all current datafiles and online redo logs and all subsequent archived redo logs with a new RESETLOGS SCN and time stamp.

Because the database will not apply an archived log to a datafile unless the RESETLOGS SCN and time stamps match, the RESETLOGS prevents you from corrupting datafiles with archived logs that are not from direct parent incarnations of the current incarnation.

In prior releases, it was recommended that you back up the database immediately after the RESETLOGS. Because you can now easily recover a pre-RESETLOGS backup like any other backup, making a new database backup is optional. In order to perform recovery through resetlogs you must have all archived logs generated since the last backup and at least one control file (current, backup, or created).

Figure 18–1 shows the case of a database that can only be recovered to log sequence 2500 because an archived redo log is missing. When the online redo log is at sequence 4000, the database crashes. You restore the sequence 1000 backup and prepare for complete recovery. Unfortunately, one of your archived logs is corrupted. The log before the missing log contains sequence 2500, so you recover to this log sequence and open RESETLOGS. As part of the RESETLOGS, the database archives the current online logs (sequence 4000 and 4001) and resets the log sequence to 1.

You generate changes in the new incarnation of the database, eventually reaching log sequence 4000. The changes between sequence 2500 and sequence 4000 for the new incarnation of the database are different from the changes between sequence 2500 and sequence 4000 for the old incarnation. You cannot apply logs generated *after* 2500 in the old incarnation to the new incarnation, but you can apply the logs generated *before* sequence 2500 in the old incarnation to the new incarnation. The logs from after sequence 2500 are said to be **orphaned** in the new incarnation because they are unusable for recovery in that incarnation.

Figure 18–1 Creating a New Database Incarnation

Executing the ALTER DATABASE OPEN Statements

To preserve the log sequence number when opening a database after media recovery, execute either of the following statements:

```
ALTER DATABASE OPEN NORESETLOGS;
ALTER DATABASE OPEN;
```

To reset the log sequence number when opening a database after recovery and thereby create a new incarnation of the database, execute the following statement:

```
ALTER DATABASE OPEN RESETLOGS;
```

If you open with the `RESETLOGS` option, the database returns different messages depending on whether recovery was complete or incomplete. If the recovery was complete, then the following message appears in the `alert_SID.log` file:

```
RESETLOGS after complete recovery through change scn
```

If the recovery was incomplete, then this message is reported in the `alert_SID.log` file, where `scn` refers to the end point of incomplete recovery:

```
RESETLOGS after incomplete recovery UNTIL CHANGE scn
```

If you attempt to `OPEN RESETLOGS` when you should not, or if you neglect to reset the log when you should, then the database returns an error and does not open the database. Correct the problem and try again.

See Also: ["About User-Managed Media Recovery Problems"](#) on page 21-1 for descriptions of situations that can cause ALTER DATABASE OPEN RESETLOGS to fail

Checking the Alert Log After a RESETLOGS Operation

After opening the database with the RESETLOGS option, check the alert_<SID>.log to see whether the database detected inconsistencies between the data dictionary and the control file, for example, a datafile that the data dictionary includes but which is not listed in the new control file. The following table describes two possible scenarios.

Control File	Data Dictionary	Result
Datafile is listed	Datafile is not listed	References to the unlisted datafile are removed from the control file. A message in the alert log indicates what was found.
Datafile is not listed	Datafile is listed	The database creates a placeholder entry in the control file under MISSING<nnnnn> (where <nnnnn> is the file number in decimal). MISSING<nnnnn> is flagged in the control file as offline and requiring media recovery. You can make the datafile corresponding to MISSING<nnnnn> accessible by using ALTER DATABASE RENAME FILE for MISSING<nnnnn> so that it points to the datafile. If you do not have a backup of this datafile, then drop the tablespace.

Recovering a Database in NOARCHIVELOG Mode

If a media failure damages datafiles in a NOARCHIVELOG database, then the only option for recovery is usually to restore a consistent whole database backup. If you are using logical backups created by an Oracle export utility to supplement regular physical backups, then you can also attempt to restore the database by importing an exported backup of the database into a re-created database or a database restored from an old backup.

Restoring a NOARCHIVELOG Database to its Default Location

In this scenario, the media failure is repaired so that you are able to restore all database files to their original location.

To restore the most recent whole database backup to the default location:

1. If the database is open, then shut down the database. For example, enter:
SHUTDOWN IMMEDIATE
2. If possible, correct the media problem so that the backup database files can be restored to their original locations.
3. Restore the most recent whole database backup with operating system commands as described in ["Restoring Datafiles and Archived Redo Logs"](#) on page 18-4. Restore all of the datafiles and control files of the whole database backup, not just the damaged files. The following example restores a whole database backup to its default location:

```
% cp /backup/*.dbf $ORACLE_HOME/oradata/trgt/
```

4. Because online redo logs are not backed up, you cannot restore them with the datafiles and control files. In order to allow the database to reset the online redo logs, you must first mimic incomplete recovery:

```
RECOVER DATABASE UNTIL CANCEL
CANCEL
```

5. Open the database in RESETLOGS mode:

```
ALTER DATABASE OPEN RESETLOGS;
```

Restoring a NOARCHIVELOG Database to a New Location

In this scenario, you restore the database files to an alternative location because the original location is damaged by a media failure.

To restore the most recent whole database backup to a new location:

1. If the database is open, then shut it down. For example, enter:

```
SHUTDOWN IMMEDIATE
```

2. Restore all of the datafiles and control files of the whole database backup, not just the damaged files. If the hardware problem has not been corrected and some or all of the database files must be restored to alternative locations, then restore the whole database backup to a new location. For example, enter:

```
% cp /backup/*.dbf /new_disk/oradata/trgt/
```

3. If necessary, edit the restored parameter file to indicate the new location of the control files. For example:

```
CONTROL_FILES = "/new_disk/oradata/trgt/control01.dbf"
```

4. Start an instance using the restored and edited parameter file and mount, but do not open, the database. For example:

```
STARTUP MOUNT
```

5. If the restored datafile filenames will be different (as will be the case when you restore to a different file system or directory, on the same node or a different node), then update the control file to reflect the new datafile locations. For example, to rename datafile 1 you might enter:

```
ALTER DATABASE RENAME FILE '?/oradata/trgt/system01.dbf' TO
'/new_disk/oradata/system01.dbf';
```

6. If the online redo logs were located on a damaged disk, and the hardware problem is not corrected, then specify a new location for each affected online log. For example, enter:

```
ALTER DATABASE RENAME FILE '?/oradata/trgt/redo01.log' TO
'/new_disk/oradata/redo_01.log';
ALTER DATABASE RENAME FILE '?/oradata/trgt/redo02.log' TO
'/new_disk/oradata/redo_02.log';
```

7. Because online redo logs are not backed up, you cannot restore them with the datafiles and control files. In order to allow the database to reset the online redo logs, you must first mimic incomplete recovery:

```
RECOVER DATABASE UNTIL CANCEL;
CANCEL;
```

8. Open the database in RESETLOGS mode. This command clears the online redo logs and resets the log sequence to 1:

```
ALTER DATABASE OPEN RESETLOGS;
```

Note that restoring a NOARCHIVELOG database backup and then resetting the log discards all changes to the database made from the time the backup was taken to the time of the failure.

See Also: *Oracle Database Administrator's Guide* for more information about renaming and relocating datafiles, and *Oracle Database SQL Reference* to learn about ALTER DATABASE RENAME FILE

Controlling Parallel Media Recovery

By default, Oracle uses **parallel media recovery** to improve performance of the roll forward phase of media recovery. In parallel media recovery, the database uses a "division of labor" approach to allocate different processes to different data blocks while rolling forward, thereby making the procedure more efficient. The number of processes used is derived from the CPU_COUNT initialization parameter, which by default is equal to the number of CPUs on the system. For example, if parallel recovery is performed on a system where CPU_COUNT is 4, and only one datafile is recovered, then four spawned processes read blocks from the datafile and apply redo.

Typically, recovery is I/O-bound on reads from and writes to data blocks. Parallelism at the block level may only help recovery performance if it increases total I/Os, for example, by bypassing operating system restrictions on asynchronous I/Os. Systems with efficient asynchronous I/O see little benefit from parallel media recovery.

To override the default behavior of performing parallel recovery, use the RECOVER with the NOPARALLEL option, or RECOVER PARALLEL 0.

Note: The RECOVERY_PARALLELISM initialization parameter controls instance or crash recovery *only*. Media recovery is not affected by the value used for RECOVERY_PARALLELISM.

See Also:

- *Oracle Database Performance Tuning Guide* for more information on parallel recovery
- *SQL*Plus User's Guide and Reference* for more information about the SQL*Plus RECOVER . . . PARALLEL and NOPARALLEL statements

Advanced User-Managed Recovery Scenarios

This chapter describes several common media failure scenarios. It shows how to recover from each failure when using a user-managed backup and recovery strategy, that is, a strategy that does not depend upon Recovery Manager. This chapter includes the following topics:

- [Recovering After the Loss of Datafiles: Scenarios](#)
- [Recovering Through an Added Datafile with a Backup Control File: Scenario](#)
- [Re-Creating Datafiles When Backups Are Unavailable: Scenario](#)
- [Recovering Through RESETLOGS with Created Control File: Scenario](#)
- [Recovering NOLOGGING Tables and Indexes: Scenario](#)
- [Recovering Read-Only Tablespaces with a Backup Control File: Scenario](#)
- [Media Recovery of Transportable Tablespaces: Scenario](#)
- [Recovering After the Loss of Online Redo Log Files: Scenarios](#)
- [Recovering After the Loss of Archived Redo Log Files: Scenario](#)
- [Recovering from a Dropped Table: Scenario](#)
- [Performing Media Recovery in a Distributed Environment: Scenario](#)
- [Dropping a Database with SQL*Plus](#)

Recovering After the Loss of Datafiles: Scenarios

If a media failure affects datafiles, then the recovery procedure depends on:

- The archiving mode of the database: ARCHIVELOG or NOARCHIVELOG
- The type of media failure
- The files affected by the media failure

Losing Datafiles in NOARCHIVELOG Mode

If either a permanent or temporary media failure affects any datafiles of a database operating in NOARCHIVELOG mode, then the database automatically shuts down.

If the media failure is temporary, correct the underlying problem and restart the database. Usually, crash recovery will recover all committed transactions from the

online redo log. If the media failure is permanent, then restore the database as described in ["Recovering a Database in NOARCHIVELOG Mode"](#) on page 18-27.

Losing Datafiles in ARCHIVELOG Mode

If either a permanent or temporary media failure affects the datafiles of a database operating in ARCHIVELOG mode, then the following scenarios can occur.

Damaged Datafiles	Database Status	Solution
Datafiles in the <code>SYSTEM</code> tablespace or datafiles with active undo segments.	Database shuts down.	If the hardware problem is temporary, then fix it and restart the database. Usually, crash recovery recovers lost transactions. If the hardware problem is permanent, then recover the database as described in "Performing Closed Database Recovery" on page 18-16.
Datafiles not in the <code>SYSTEM</code> tablespace or datafiles that do not contain active rollback or undo segments.	Affected datafiles are taken offline, but the database stays open.	If the unaffected portions of the database must remain available, then do not shut down the database. Take tablespaces containing problem datafiles offline using the temporary option, then recover them as described in "Performing Datafile Recovery in an Open Database" on page 18-18.

Recovering Through an Added Datafile with a Backup Control File: Scenario

If database recovery with a backup control file rolls forward through a `CREATE TABLESPACE` or an `ALTER TABLESPACE ADD DATAFILE` operation, then the database stops recovery when applying the redo record for the added files and lets you confirm the filenames.

For example, suppose the following sequence of events occurs:

1. You back up the database
2. You create a new tablespace containing two datafiles:
`/oracle/oradata/trgt/test01.dbf` and
`/oracle/oradata/trgt/test02.dbf`.
3. You later restore a backup control file and perform media recovery through the `CREATE TABLESPACE` operation.

You may see the following error when applying the `CREATE TABLESPACE` redo data:

```
ORA-00283: recovery session canceled due to errors
ORA-01244: unnamed datafile(s) added to control file by media recovery
ORA-01110: data file 11: '/oracle/oradata/trgt/test02.dbf'
ORA-01110: data file 10: '/oracle/oradata/trgt/test01.dbf'
```

To recover through an `ADD DATAFILE` operation, use the following procedure:

1. View the files added by selecting from `V$DATAFILE`. For example:

```
SELECT FILE#,NAME
FROM V$DATAFILE;

FILE#          NAME
-----
1              /oracle/oradata/trgt/system01.dbf
.
```

```

10          /oracle/oradata/trgt/UNNAMED00001
11          /oracle/oradata/trgt/UNNAMED00002

```

2. If multiple unnamed files exist, then determine which unnamed file corresponds to which datafile by using one of these methods:

- Open the `alert_SID.log`, which contains messages about the original file location for each unnamed file.
- Derive the original file location of each unnamed file from the error message and `V$DATAFILE`: each unnamed file corresponds to the file in the error message with the same file number.

3. Issue the `ALTER DATABASE RENAME FILE` statement to rename the datafiles. For example, enter:

```

ALTER DATABASE RENAME FILE '/db/UNNAMED00001' TO
                          '/oracle/oradata/trgt/test01.dbf';
ALTER DATABASE RENAME FILE '/db/UNNAMED00002' TO
                          '/oracle/oradata/trgt/test02.dbf';

```

4. Continue recovery by issuing the previous recovery statement. For example:

```
RECOVER AUTOMATIC DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL
```

Re-Creating Datafiles When Backups Are Unavailable: Scenario

If a datafile is damaged and no backup of the file is available, then you can still recover the datafile if:

- All archived log files written after the creation of the original datafile are available
- The control file contains the name of the damaged file (that is, the control file is current, or is a backup taken after the damaged datafile was added to the database)

Note: You cannot re-create any of the datafiles for the `SYSTEM` tablespace by using the `CREATE DATAFILE` clause of the `ALTER DATABASE` statement because the necessary redo is not available.

To re-create a datafile for recovery:

1. Create a new, empty datafile to replace a damaged datafile that has no corresponding backup. For example, assume that the datafile `/?/oradata/trgt/users01.dbf` has been damaged, and no backup is available. The following statement re-creates the original datafile (same size) on `disk2`:

```

ALTER DATABASE CREATE DATAFILE '/oradata/trgt/users01.dbf' AS
                          '/disk2/users01.dbf';

```

This statement creates an empty file that is the same size as the lost file. The database looks at information in the control file and the data dictionary to obtain size information. The old datafile is renamed as the new datafile.

2. Perform media recovery on the empty datafile. For example, enter:

```
RECOVER DATAFILE '/disk2/users01.dbf'
```

3. All archived logs written after the original datafile was created must be applied to the new, empty version of the lost datafile during recovery.

Recovering Through RESETLOGS with Created Control File: Scenario

You can recover backups through an OPEN RESETLOGS so long as:

- You have a current, backup, or created control file that knows about the prior incarnations
- You have all available archived redo logs

If you need to re-create the control file, the trace file generated by ALTER DATABASE BACKUP CONTROLFILE TO TRACE will contain the necessary commands to re-construct the complete incarnation history. The V\$DATABASE_INCARNATION view displays the RESETLOGS history known to the control file, while the V\$LOG_HISTORY view displays the archived log history.

It is possible for the incarnation history to be incomplete in the in re-created control file. For example, archived logs necessary for recovery may be missing. In this case, it is possible to create incarnation records explicitly with the ALTER DATABASE REGISTER LOGFILE statement.

In the following example, you register four logs that are necessary for recovery but are not recorded in the re-created control file, and then recover the database:

```
ALTER DATABASE REGISTER LOGFILE '?/oradata/trgt/arch/arcr_1_1_42343523.arc';
ALTER DATABASE REGISTER LOGFILE '?/oradata/trgt/arch/arcr_1_1_34546466.arc';
ALTER DATABASE REGISTER LOGFILE '?/oradata/trgt/arch/arcr_1_1_23435466.arc';
ALTER DATABASE REGISTER LOGFILE '?/oradata/trgt/arch/arcr_1_1_12343533.arc';
RECOVER AUTOMATIC DATABASE;
```

Recovering NOLOGGING Tables and Indexes: Scenario

You can create tables and indexes with the CREATE TABLE AS SELECT statement. You can also specify that the database create them with the NOLOGGING option. When you create a table or index as NOLOGGING, the database does not generate redo log records for the operation. Thus, you cannot recover objects created with NOLOGGING, even if you are running in ARCHIVELOG mode.

Note: If you cannot afford to lose tables or indexes created with NOLOGGING, then make a backup after the unrecoverable table or index is created.

Be aware that when you perform media recovery, and some tables or indexes are created normally whereas others are created with the NOLOGGING option, the NOLOGGING objects are marked logically corrupt by the RECOVER operation. Any attempt to access the unrecoverable objects returns an ORA-01578 error message. Drop the NOLOGGING objects and re-create them if needed.

Because it is possible to create a table with the NOLOGGING option and then create an index with the LOGGING option on that table, the index is not marked as logically corrupt after you perform media recovery. The table was unrecoverable (and thus marked as corrupt after recovery), however, so the index points to corrupt blocks. The index must be dropped, and the table and index must be re-created if necessary.

See Also: *Oracle Data Guard Concepts and Administration* for information about the impact of NOLOGGING on a database

Recovering Read-Only Tablespaces with a Backup Control File: Scenario

If you have a read-only tablespace on read-only or slow media, then you may encounter errors or poor performance when recovering with the `USING BACKUP CONTROLFILE` option. This situation occurs when the backup control file indicates that a tablespace was read/write when the control file was backed up. In this case, media recovery may attempt to write to the files. For read-only media, the database issues an error saying that it cannot write to the files. For slow media, such as a hierarchical storage system backed up by tapes, performance may suffer.

To avoid these recovery problems, use current control files rather than backups to recover the database. If you need to use a backup control file, then you can also avoid this problem if the read-only tablespace has not suffered a media failure.

Recovery of Read-Only or Slow Media with a Backup Control File

You have these alternatives for recovering read-only and slow media when using a backup control file:

- Take datafiles from read-only tablespaces offline before doing recovery with a backup control file, and then bring the files online at the end of media recovery.
- Use the correct version of the control file for the recovery. If the tablespace will be read-only when recovery completes, then the control file backup must be from a time when the tablespace was read-only. Similarly, if the tablespace will be read/write at the end of recovery, then the control file must be from a time when the tablespace was read/write.

Recovery of Read-Only Files with a Re-Created Control File

If a current or backup control file is unavailable for recovery, then you can execute a `CREATE CONTROLFILE` statement as described in "[Create New Control File After Losing All Current and Backup Control Files](#)" on page 18-9. Read-only files should not be listed in the `CREATE CONTROLFILE` statement so that recovery can skip these files. No recovery is required for read-only datafiles unless you restored backups of these files from a time when the datafiles were read/write.

After you create a new control file and attempt to mount and open the database, the database performs a data dictionary check against the files listed in the control file. For each file that is not listed in the `CREATE CONTROLFILE` statement but is present in the data dictionary, an entry is created for them in the control file. These files are named as `MISSINGnnnnn`, where `nnnnn` is a five digit number starting with 0.

After the database is open, rename the read-only files to their correct filenames by executing the `ALTER DATABASE RENAME FILE` statement for all the files whose name is prefixed with `MISSING`.

To prepare for a scenario in which you might have to re-create the control file, run the following statement when the database is mounted or open to obtain the `CREATE CONTROLFILE` syntax:

```
ALTER DATABASE BACKUP CONTROLFILE TO TRACE;
```

This SQL statement produces a trace file that you can edit and use as a script to re-create the control file. You can specify either the `RESETLOGS` or `NORESETLOGS`

(default) keywords to generate `CREATE CONTROLFILE . . . RESETLOGS` or `CREATE CONTROLFILE . . . NORESETLOGS` versions of the script.

All the restrictions related to read-only files in `CREATE CONTROLFILE` statements also apply to offline normal tablespaces, except that you need to bring the tablespace online after the database is open. You should leave out tempfiles from the `CREATE CONTROLFILE` statement and add them after database open.

See Also: *Oracle Database Backup and Recovery Basics* to learn how to make trace backups of the control file

Media Recovery of Transportable Tablespaces: Scenario

The **transportable tablespace** feature of Oracle allows a user to transport a set of tablespaces from one database to another. Transporting a tablespace into a database is like creating a tablespace with preloaded data. Using this feature is often an advantage because:

- It is faster than using the Export or SQL*Loader utilities because it involves only copying datafiles and integrating metadata
- You can use it to move index data, hence avoiding the necessity of rebuilding indexes

See Also: *Oracle Database Administrator's Guide* for detailed information about using the transportable tablespace feature

Like normal tablespaces, transportable tablespaces are recoverable. However, while you can recover normal tablespaces without a backup, you must have a version of the transported datafiles in order to recover a transported tablespace.

To recover a transportable tablespace, use the following procedure:

1. If the database is open, then take the transported tablespace offline. For example, if you want to recover the `users` tablespace, then issue:

```
ALTER TABLESPACE users OFFLINE IMMEDIATE;
```

2. Restore a backup of the transported datafiles with an operating system utility. The backup can be the initial version of the transported datafiles or any backup taken after the tablespace is transported. For example, enter:

```
% cp /backup/users.dbf $ORACLE_HOME/oradata/trgt/users01.dbf
```

3. Recover the tablespace as normal. For example, enter:

```
RECOVER TABLESPACE users
```

You may see the error `ORA-01244` when recovering through a transportable tablespace operation just as when recovering through a `CREATE TABLESPACE` operation. In this case, rename the unnamed files to the correct locations using the procedure in ["Recovering Through an Added Datafile with a Backup Control File: Scenario"](#) on page 19-2.

Recovering After the Loss of Online Redo Log Files: Scenarios

If a media failure has affected the online redo logs of a database, then the appropriate recovery procedure depends on the following:

- The configuration of the online redo log: mirrored or non-mirrored

- The type of media failure: temporary or permanent
- The types of online redo log files affected by the media failure: current, active, unarchived, or inactive

Table 19–1 displays V\$LOG status information that can be crucial in a recovery situation involving online redo logs.

Table 19–1 STATUS Column of V\$LOG

Status	Description
UNUSED	The online redo log has never been written to.
CURRENT	The online redo log is active, that is, needed for instance recovery, and it is the log to which the database is currently writing. The redo log can be open or closed.
ACTIVE	The online redo log is active, that is, needed for instance recovery, but is not the log to which the database is currently writing. It may be in use for block recovery, and may or may not be archived.
CLEARING	The log is being re-created as an empty log after an ALTER DATABASE CLEAR LOGFILE statement. After the log is cleared, then the status changes to UNUSED.
CLEARING_CURRENT	The current log is being cleared of a closed thread. The log can stay in this status if there is some failure in the switch such as an I/O error writing the new log header.
INACTIVE	The log is no longer needed for instance recovery. It may be in use for media recovery, and may or may not be archived.

Recovering After Losing a Member of a Multiplexed Online Redo Log Group

If the online redo log of a database is multiplexed, and if at least one member of each online redo log group is not affected by the media failure, then the database continues functioning as normal, but error messages are written to the log writer trace file and the `alert_SID.log` of the database.

Solve the problem by taking one of the following actions:

- If the hardware problem is temporary, then correct it. The log writer process accesses the previously unavailable online redo log files as if the problem never existed.
- If the hardware problem is permanent, then drop the damaged member and add a new member by using the following procedure.

Note: The newly added member provides no redundancy until the log group is reused.

1. Locate the filename of the damaged member in V\$LOGFILE. The status is INVALID if the file is inaccessible:

```
SELECT GROUP#, STATUS, MEMBER
FROM V$LOGFILE
WHERE STATUS='INVALID';
```

```
GROUP#      STATUS      MEMBER
-----
0002      INVALID      /oracle/oradata/trgt/redo02.log
```

- Drop the damaged member. For example, to drop member redo01.log from group 2, issue:

```
ALTER DATABASE DROP LOGFILE MEMBER '/oracle/oradata/trgt/redo02.log';
```

- Add a new member to the group. For example, to add redo02.log to group 2, issue:

```
ALTER DATABASE ADD LOGFILE MEMBER '/oracle/oradata/trgt/redo02b.log'
TO GROUP 2;
```

If the file you want to add already exists, then it must be the same size as the other group members, and you must specify REUSE. For example:

```
ALTER DATABASE ADD LOGFILE MEMBER '/oracle/oradata/trgt/redo02b.log'
REUSE TO GROUP 2;
```

Recovering After the Loss of All Members of an Online Redo Log Group

If a media failure damages all members of an online redo log group, then different scenarios can occur depending on the type of online redo log group affected by the failure and the archiving mode of the database.

If the damaged log group is active, then it is needed for crash recovery; otherwise, it is not.

If the group is . . .	Then . . .	And you should . . .
Inactive	It is not needed for crash recovery	Clear the archived or unarchived group.
Active	It is needed for crash recovery	Attempt to issue a checkpoint and clear the log; if impossible, then you must restore a backup and perform incomplete recovery up to the most recent available redo log.
Current	It is the log that the database is currently writing to	Attempt to clear the log; if impossible, then you must restore a backup and perform incomplete recovery up to the most recent available redo log.

Your first task is to determine whether the damaged group is active or inactive.

- Locate the filename of the lost redo log in V\$LOGFILE and then look for the group number corresponding to it. For example, enter:

```
SELECT GROUP#, STATUS, MEMBER FROM V$LOGFILE;
```

```
GROUP#    STATUS    MEMBER
-----
0001
0001
0002      INVALID  /oracle/dbs/log2a.f
0002      INVALID  /oracle/dbs/log2b.f
0003
0003
```

- Determine which groups are active. For example, enter:

```
SELECT GROUP#, MEMBERS, STATUS, ARCHIVED
FROM V$LOG;

GROUP# MEMBERS          STATUS  ARCHIVED
```

```

-----
0001  2          INACTIVE  YES
0002  2          ACTIVE    NO
0003  2          CURRENT   NO

```

- If the affected group is inactive, follow the procedure in [Losing an Inactive Online Redo Log Group](#) on page 19-9. If the affected group is active (as in the preceding example), then follow the procedure in ["Losing an Active Online Redo Log Group"](#) on page 19-10.

Losing an Inactive Online Redo Log Group

If all members of an online redo log group with `INACTIVE` status are damaged, then the procedure depends on whether you can fix the media problem that damaged the inactive redo log group.

If the failure is . . .	Then . . .
Temporary	Fix the problem. LGWR can reuse the redo log group when required.
Permanent	The damaged inactive online redo log group eventually halts normal database operation. Reinitialize the damaged group manually by issuing the <code>ALTER DATABASE CLEAR LOGFILE</code> statement as described in this section.

Clearing Inactive, Archived Redo You can clear an inactive redo log group when the database is open or closed. The procedure depends on whether the damaged group has been archived.

To clear an inactive, online redo log group that has been archived, use the following procedure:

- If the database is shut down, then start a new instance and mount the database:

```
STARTUP MOUNT
```
- Reinitialize the damaged log group. For example, to clear redo log group 2, issue the following statement:

```
ALTER DATABASE CLEAR LOGFILE GROUP 2;
```

Clearing Inactive, Not-Yet-Archived Redo Clearing a not-yet-archived redo log allows it to be reused without archiving it. This action makes backups unusable if they were started before the last change in the log, unless the file was taken offline prior to the first change in the log. Hence, if you need the cleared log file for recovery of a backup, then you cannot recover that backup. Also, it prevents complete recovery from backups due to the missing log.

To clear an inactive, online redo log group that has not been archived, use the following procedure:

- If the database is shut down, then start a new instance and mount the database:

```
STARTUP MOUNT
```
- Clear the log using the `UNARCHIVED` keyword. For example, to clear log group 2, issue:

```
ALTER DATABASE CLEAR LOGFILE UNARCHIVED GROUP 2;
```

If there is an offline datafile that requires the cleared log to bring it online, then the keywords `UNRECOVERABLE DATAFILE` are required. The datafile and its entire tablespace have to be dropped because the redo necessary to bring it online is being cleared, and there is no copy of it. For example, enter:

```
ALTER DATABASE CLEAR LOGFILE UNARCHIVED GROUP 2 UNRECOVERABLE DATAFILE;
```

3. Immediately back up the whole database with an operating system utility, so that you have a backup you can use for complete recovery without relying on the cleared log group. For example, enter:

```
% cp /disk1/oracle/dbs/*.f /disk2/backup
```

4. Back up the database's control file with the `ALTER DATABASE` statement. For example, enter:

```
ALTER DATABASE BACKUP CONTROLFILE TO '/oracle/dbs/cf_backup.f';
```

Failure of CLEAR LOGFILE Operation The `ALTER DATABASE CLEAR LOGFILE` statement can fail with an I/O error due to media failure when it is not possible to:

- Relocate the redo log file onto alternative media by re-creating it under the currently configured redo log filename
- Reuse the currently configured log filename to re-create the redo log file because the name itself is invalid or unusable (for example, due to media failure)

In these cases, the `ALTER DATABASE CLEAR LOGFILE` statement (before receiving the I/O error) would have successfully informed the control file that the log was being cleared and did not require archiving. The I/O error occurred at the step in which the `CLEAR LOGFILE` statement attempts to create the new redo log file and write zeros to it. This fact is reflected in `V$LOG.CLEARING_CURRENT`.

Losing an Active Online Redo Log Group

If the database is still running and the lost active redo log is *not* the current log, then issue the `ALTER SYSTEM CHECKPOINT` statement. If successful, then the active redo log becomes inactive, and you can follow the procedure in "[Losing an Inactive Online Redo Log Group](#)" on page 19-9. If unsuccessful, or if your database has halted, then perform one of procedures in this section, depending on the archiving mode.

The current log is the one LGWR is currently writing to. If a LGWR I/O fails, then LGWR terminates and the instance crashes. In this case, you must restore a backup, perform incomplete recovery, and open the database with the `RESETLOGS` option.

To recover from loss of an active online log group in `NOARCHIVELOG` mode:

1. If the media failure is temporary, then correct the problem so that the database can reuse the group when required.
2. Restore the database from a consistent, whole database backup (datafiles and control files) as described in "[Restoring Datafiles Before Performing Incomplete Recovery](#)" on page 18-21. For example, enter:

```
% cp /disk2/backup/*.dbf $ORACLE_HOME/oradata/trgt/
```

3. Mount the database:

```
STARTUP MOUNT
```

4. Because online redo logs are not backed up, you cannot restore them with the datafiles and control files. In order to allow the database to reset the online redo logs, you must first mimic incomplete recovery:

```
RECOVER DATABASE UNTIL CANCEL
CANCEL
```

5. Open the database using the RESETLOGS option:

```
ALTER DATABASE OPEN RESETLOGS;
```

6. Shut down the database consistently. For example, enter:

```
SHUTDOWN IMMEDIATE
```

7. Make a whole database backup.

To recover from loss of an active online redo log group in ARCHIVELOG mode:

If the media failure is temporary, then correct the problem so that the database can reuse the group when required. If the media failure is not temporary, then use the following procedure.

1. Begin incomplete media recovery, recovering up through the log before the damaged log.
2. Ensure that the current name of the lost redo log can be used for a newly created file. If not, then rename the members of the damaged online redo log group to a new location. For example, enter:

```
ALTER DATABASE RENAME FILE "?/oradata/trgt/redo01.log" TO "/tmp/redo01.log";
ALTER DATABASE RENAME FILE "?/oradata/trgt/redo01.log" TO "/tmp/redo02.log";
```

3. Open the database using the RESETLOGS option:

```
ALTER DATABASE OPEN RESETLOGS;
```

Note: All updates executed from the endpoint of the incomplete recovery to the present must be re-executed.

Loss of Multiple Redo Log Groups

If you have lost multiple groups of the online redo log, then use the recovery method for the most difficult log to recover. The order of difficulty, from most difficult to least difficult, follows:

1. The current online redo log
2. An active online redo log
3. An unarchived online redo log
4. An inactive online redo log

Recovering After the Loss of Archived Redo Log Files: Scenario

If the database is operating in ARCHIVELOG mode, and if the only copy of an archived redo log file is damaged, then the damaged file does not affect the present operation of the database. The following situations can arise, however, depending on when the redo log was written and when you backed up the datafile.

If you backed up . . .	Then . . .
All datafiles after the filled online redo log group (which is now archived) was written	The archived version of the filled online redo log group is not required for complete media recovery operation.
A specific datafile before the filled online redo log group was written	If the corresponding datafile is damaged by a permanent media failure, use the most recent backup of the damaged datafile and perform incomplete recovery of the tablespace containing the damaged datafile, up to the damaged log.

Caution: If you know that an archived redo log group has been damaged, immediately back up all datafiles so that you will have a whole database backup that does not require the damaged archived redo log.

Recovering from a Dropped Table: Scenario

One not-uncommon error is the accidental dropping of a table from your database. In general, the fastest and simplest solution is to use the flashback drop feature, described in *Oracle Database Backup and Recovery Basics*, to reverse the dropping of the table. However, if for some reason, such as flashback drop being disabled or the table having been dropped with the PURGE option, you cannot use flashback table, you can create a copy of the database, perform point-in-time recovery of that copy to a time before the table was dropped, export the dropped table using an Oracle export utility, and re-import it into your primary database using an Oracle import utility.

In this scenario, assume that you do not have the flashback database functionality enabled, so FLASHBACK DATABASE is not an option, but you do have physical backups of the database.

Note: Grant powerful privileges (such as DROP ANY TABLE) only to only selected, appropriate users, to minimize user errors that require database recovery.

To recover a table that has been accidentally dropped, use the following procedure:

1. If possible, keep the database that experienced the user error online and available for use. Back up all datafiles of the existing database in case an error is made during the remaining steps of this procedure.
2. Restore a database backup to an alternate location, then perform incomplete recovery of this backup using a restored backup control file, to the point just before the table was dropped.
3. Export the lost data from the temporary, restored version of the database using an Oracle export utility. In this case, export the accidentally dropped table.

Note: System audit options are exported.

4. Use an Oracle import utility to import the data back into the production database.
5. Delete the files of the temporary copy of the database to conserve space.

See Also: *Oracle Database Utilities* for more information about the Oracle export and import utilities

Performing Media Recovery in a Distributed Environment: Scenario

How you perform media recovery depends on whether your database participates in a distributed database system. The Oracle distributed database architecture is autonomous. Therefore, depending on the type of recovery operation selected for a single damaged database, you may have to coordinate recovery operations globally among all databases in the distributed system.

Table 19–2, "Recovery Operations in a Distributed Database Environment" summarizes different types of recovery operations and whether coordination among nodes of a distributed database system is required.

Table 19–2 Recovery Operations in a Distributed Database Environment

If you are . . .	Then . . .
Restoring a whole backup for a database that was never accessed from a remote node	Use non-coordinated, autonomous database recovery.
Restoring a whole backup for a database that was accessed by a remote node for a database in NOARCHIVELOG mode	Shut down all databases and restore them using the same coordinated full backup.
Performing complete media recovery of one or more databases in a distributed database	Use non-coordinated, autonomous database recovery.
Performing incomplete media recovery of a database that was never accessed by a remote node	Use non-coordinated, autonomous database recovery.
Performing incomplete media recovery of a database that was accessed by a remote node	Use coordinated, incomplete recovery to the same global point in time for all databases in the distributed system.

Coordinating Time-Based and Change-Based Distributed Database Recovery

If one node in a distributed database requires recovery to a past time, it is often necessary to recover all other nodes in the system to the same point in time to preserve global data consistency. This operation is called **coordinated, time-based, distributed database recovery**. The following tasks should be performed with the standard procedures of time-based and change-based recovery described in this chapter.

1. Recover the database that requires the recovery operation using time-based recovery. For example, if a database needs to be recovered because of a media failure, then recover this database first using time-based recovery. Do not recover the other databases at this point.
2. After you have recovered the database and opened it with the RESETLOGS option, search the `alert_SID.log` of the database for the RESETLOGS message.

If the message is, "RESETLOGS after complete recovery through change xxx", then you have applied all the changes in the database and performed complete recovery. Do not recover any of the other databases in the distributed system, or you will unnecessarily remove changes in them. Recovery is complete.

If the message is, "RESETLOGS after incomplete recovery UNTIL CHANGE xxx", then you have successfully performed an incomplete recovery. Record the change number from the message and proceed to the next step.

3. Recover all other databases in the distributed database system using change-based recovery, specifying the change number (SCN) from Step 2.

Dropping a Database with SQL*Plus

You may need to remove a database, that is, the database files that form the database, from the operating system. For example, this scenario can occur when you create a test database and then no longer have a use for it. The SQL*Plus command `DROP DATABASE` can perform this function.

See Also: *Oracle Database Backup and Recovery Basics* to learn how to use the equivalent RMAN command `DROP DATABASE`

1. Start SQL*Plus and connect to the target database with administrator privileges, then ensure that the database is either mounted or open with no users connected. For example:

```
SQL> STARTUP FORCE MOUNT
```

2. Remove the datafiles and control files listed in the control file from the operating system. For example:

```
SQL> DROP DATABASE; # deletes all database files, both ASM and non-ASM
```

If the database is on raw disk, the command does not delete the actual raw disk special files.

3. Use an operating system utility to delete all backups and archived logs associated with the database because these are not automatically deleted by the SQL*Plus command. For example:

```
% rm /backup/* ?/oradata/trgt/arch/*
```

Performing User-Managed TSPITR

This chapter describes how to perform user-managed tablespace point-in-time recovery (TSPITR) with the transportable tablespace feature. The process for performing TSPITR described in this chapter does not depend upon the use of Recovery Manager (RMAN).

This chapter includes the following topics:

- [Introduction to User-Managed Tablespace Point-in-Time Recovery](#)
- [Preparing for User-Managed Tablespace Point-in-Time Recovery: Basic Steps](#)
- [Restoring and Recovering the Auxiliary Database in User-Managed TSPITR: Basic Steps](#)
- [Performing User-Managed TSPITR with Transportable Tablespaces](#)
- [Performing Partial TSPITR of Partitioned Tables](#)
- [Performing User-Managed TSPITR of Partitioned Tables With a Dropped Partition](#)
- [Performing User-Managed TSPITR of Partitioned Tables When a Partition Has Split](#)

See Also: [Chapter 8, "RMAN Tablespace Point-in-Time Recovery \(TSPITR\)"](#)

Introduction to User-Managed Tablespace Point-in-Time Recovery

Tablespace point-in-time recovery (TSPITR) with the transportable tablespace feature enables you to quickly recover one or more tablespaces (other than the `SYSTEM` tablespace) to a time that is prior to the rest of the database.

User-managed TSPITR is most useful for recovering the following:

- An erroneous `DROP TABLESPACE` operation
- An incorrect batch job or other DML statement that has affected only a subset of the database
- A logical schema to a point different from the rest of the physical database when multiple schemas exist in separate tablespaces of one physical database
- A tablespace in a VLDB (very large database) when TSPITR is more efficient than restoring the whole database from a backup and rolling it forward

Refer to ["Preparing for User-Managed Tablespace Point-in-Time Recovery: Basic Steps"](#) on page 20-3 before deciding to perform TSPITR.

TSPITR Terminology

Familiarize yourself with the following terms and abbreviations, which are used throughout this chapter:

TSPITR

Tablespace point-in-time recovery

Primary Database

The database containing the tablespace or tablespaces that you want to recover to a prior point in time.

Auxiliary Database

A copy of the current database that is restored from a backup. It includes restored backups on the auxiliary host of the following files:

- Datafiles belonging to the `SYSTEM` tablespace
- Datafiles in the set of tablespaces to be recovered
- Datafiles belonging to an undo tablespace or tablespace that contains rollback segments

All backups must be from a point in time prior to the desired recovery time.

Recovery Set

All the tablespaces on the primary database that require point-in-time recovery to be performed on them.

Recovery Set Self-Containment Check

All objects that are part of the recovery set must be self-contained: there can be no dependencies on objects outside the recovery set. For example, if a table is part of the recovery set and its indexes are in a separate tablespace, then the recovery set must include the tablespace containing the index. Alternatively, the index can be dropped. You can check the recovery set tablespaces for self-containment with the procedure `DBMS_TTS.TRANSPORT_SET_CHECK`.

Auxiliary Set

Any other files required for restoring the auxiliary database, including:

- Backup control file
- Datafiles from the `SYSTEM` tablespace
- Datafiles in an undo tablespace or datafiles containing rollback segments

Transportable Tablespace

A rapid method of transporting tablespaces across databases by unplugging them from a source database and plugging them into a target database. The databases can even be on different platforms, for example, Solaris and Windows 2000. The unplugging and plugging is done with the Export and Import utilities. Note that there is no actual export and import of the table data, but simply an export and import of internal metadata. During the procedure, the datafiles of the transported tablespaces are made part of the target database.

TSPITR Methods

In releases prior to Oracle9i, you had the following two methods for performing user-managed TSPITR:

- Traditional user-managed TSPITR, which required you to create a special type of database called a **clone database**
- User-managed TSPITR with the transportable tablespace feature

As of Oracle Database Release 10g, TSPITR should be performed by using the transportable tablespace feature. This procedure is relatively easy to use and is less error prone than the traditional method, which is currently deprecated (although not yet unsupported).

TSPITR is performed by dropping the tablespaces to be recovered from the primary database, restoring a copy of the database called an **auxiliary database** and recovering it to the desired point in time, then transporting the relevant tablespaces from the auxiliary database to the current version of the primary database.

For ease of use, it is highly recommended that you place the auxiliary and primary databases on different hosts. Nevertheless, you can also perform TSPITR when the databases are located on the same host.

The basic procedure for performing user-managed TSPITR is as follows:

1. Take the tablespaces requiring TSPITR offline.
2. Plan the setup of the auxiliary database.
3. Create the auxiliary database and recover it to the desired point in time.
4. Drop the tablespaces requiring TSPITR from the primary database.
5. Use the transportable tablespace feature to transport the set of tablespaces from the auxiliary database to the primary database.

See Also: *Oracle Database Administrator's Guide* for a complete account of how to use the transportable tablespace feature

Preparing for User-Managed Tablespace Point-in-Time Recovery: Basic Steps

User-managed tablespace point-in-time recovery (TSPITR) requires careful planning. Before proceeding you should read this chapter thoroughly.

This section contains the following topics:

- [Step 1: Review TSPITR Requirements](#)
- [Step 2: Identify All of the Files in the Recovery and Auxiliary Set Tablespaces](#)
- [Step 3: Determine Whether Objects Will Be Lost](#)
- [Step 4: Choose a Method for Connecting to the Auxiliary Instance](#)
- [Step 5: Create an Oracle Password File for the Auxiliary Instance](#)
- [Step 6: Create the Initialization Parameter File for the Auxiliary Instance](#)

Caution: You should not perform TSPITR for the first time on a production system, or when there is a time constraint.

Step 1: Review TSPITR Requirements

Satisfy the following requirements before performing TSPITR:

- Ensure that you have backups of all datafiles in the recovery and auxiliary set tablespaces. The datafile backups must have been created *before* the desired TSPITR time.
- Ensure that you have a control file backup that is usable on the auxiliary database. To be usable, the control file must meet these requirements:
 - The control file must have been backed up *before* the desired TSPITR time.
 - The control file must have been backed up with the following SQL statement, where *cf_name* refers to the fully specified filename:


```
ALTER DATABASE BACKUP CONTROLFILE TO 'cf_name' ;
```
- Ensure that all files constituting the recovery set tablespaces are in the recovery set on the auxiliary database; otherwise, the export phase during tablespace transport fails.
- Allocate enough disk space on the auxiliary host to accommodate the auxiliary database.
- Provide enough real memory to start the auxiliary instance.
- If the tablespace to be recovered has been renamed, ensure that the target SCN for TSPITR is after the time when the file was renamed. You cannot TSPITR a renamed tablespace to a point in time earlier than the rename. However, you can perform a DBPITR to an SCN before the rename. In this case, the tablespace reverts to its name as of the target SCN.

See Also: ["Step 6: Create the Initialization Parameter File for the Auxiliary Instance"](#) on page 20-6.

Step 2: Identify All of the Files in the Recovery and Auxiliary Set Tablespaces

Before you create the auxiliary database, make sure that you connect to the primary database with administrator privileges and obtain all of the following information about the primary database:

- The filenames of the datafiles in the recovery set tablespaces
- The filenames of the datafiles in the *SYSTEM* tablespace
- The filenames of the datafiles in an undo tablespace or datafiles containing rollback segments
- The filenames of the control files

The following useful query displays the filenames of all datafiles and control files in the database:

```
SELECT NAME FROM V$DATAFILE
UNION ALL
SELECT NAME FROM V$CONTROLFILE;
```

To determine the filenames of the datafiles in the *SYSTEM* and recovery set tablespaces, execute the following query and replace *RECO_TBS_1*, *RECO_TBS_2*, and so forth with the names of the recovery set tablespaces:

```
SELECT t.NAME AS "reco_tbs", d.NAME AS "dbf_name"
FROM V$DATAFILE d, V$TABLESPACE t
WHERE t.TS# = d.TS#
AND t.NAME IN ('SYSTEM', 'RECO_TBS_1', 'RECO_TBS_2');
```

If you run the database in manual undo management mode (which is deprecated), then the following query displays the names of the tablespaces containing rollback segments as well as the names of the datafiles in the tablespaces:

```
SELECT DISTINCT r.TABLESPACE_NAME AS "rbs_tbs", d.FILE_NAME AS "dbf_name"
       FROM DBA_ROLLBACK_SEGS r, DBA_DATA_FILES d
WHERE r.TABLESPACE_NAME=d.TABLESPACE_NAME;
```

If you run the database in automatic undo management mode, then the following query displays the names of the undo tablespaces as well as the names of the datafiles in the tablespaces:

```
SELECT DISTINCT u.TABLESPACE_NAME AS "undo_tbs", d.FILE_NAME AS "dbf_name"
       FROM DBA_UNDO_EXTENTS u, DBA_DATA_FILES d
WHERE u.TABLESPACE_NAME=d.TABLESPACE_NAME;
```

Step 3: Determine Whether Objects Will Be Lost

When TSPITR is performed on a tablespace, any objects created after the recovery time are lost. To determine which objects will be lost, query the `TS_PITR_OBJECTS_TO_BE_DROPPED` view on the primary database. The contents of the view are described in [Table 20-1](#).

Table 20-1 *TS_PITR_OBJECTS_TO_BE_DROPPED* View

Column Name	Meaning
OWNER	Owner of the object to be dropped.
NAME	The name of the object that will be lost as a result of TSPITR
CREATION_TIME	Creation time stamp for the object.
TABLESPACE_NAME	Name of the tablespace containing the object.

When querying this view, supply all the elements of the date field, otherwise the default setting is used. Also, use the `TO_CHAR` and `TO_DATE` functions. For example, with a recovery set consisting of `users` and `tools`, and a recovery point in time of 19 October 2002, 15:34:11, execute the following SQL script:

```
SELECT OWNER, NAME, TABLESPACE_NAME,
       TO_CHAR(CREATION_TIME, 'YYYY-MM-DD:HH24:MI:SS')
       FROM SYS.TS_PITR_OBJECTS_TO_BE_DROPPED
WHERE TABLESPACE_NAME IN ('users', 'tools')
AND CREATION_TIME > TO_DATE('02-OCT-19:15:34:11', 'YY-MON-DD:HH24:MI:SS')
ORDER BY TABLESPACE_NAME, CREATION_TIME;
```

See Also: *Oracle Database Reference* for more information about the `TS_PITR_OBJECTS_TO_BE_DROPPED` view

Step 4: Choose a Method for Connecting to the Auxiliary Instance

You must be able to connect to the auxiliary instance. You can either use Oracle Net or operating system authentication. To learn how to configure networking files, refer to *Oracle Database Net Services Administrator's Guide*.

Step 5: Create an Oracle Password File for the Auxiliary Instance

For information about creating and maintaining Oracle password files, refer to the *Oracle Database Administrator's Guide*. If you do not use a password file, then you can skip this step.

Step 6: Create the Initialization Parameter File for the Auxiliary Instance

Create a new initialization parameter file rather than copying and then editing the production database initialization parameter file. Save memory by using low settings for parameters such as the following:

- DB_CACHE_SIZE
- SHARED_POOL_SIZE
- LARGE_POOL_SIZE

Reducing the preceding parameter settings can prevent the auxiliary database from starting when other dependent parameters are set too high—for example, the initialization parameter ENQUEUE_RESOURCES, which allocates memory from within the shared pool.

The auxiliary database can be either on the same host as the primary database or on a different host. Because the auxiliary database filenames are identical to the primary database filenames in the auxiliary control file, you must update the auxiliary control file to point to the locations to which the files were restored for the auxiliary database. If the auxiliary database is on the same machine as the primary database, or if the auxiliary database is on a different machine that uses different path names, then you must rename the control files, datafiles, and online redo logs. If the auxiliary database is on a different machine with the same path names, then you can rename just the online redo logs. To view the names of the online redo log files of the primary database so that you can be sure to use unique names when creating the auxiliary, use this query on the primary database:

```
SELECT NAME FROM V$LOGFILE;
```

Caution: If the auxiliary and primary database are on the same machine, then failing to rename the online redo log files may cause primary database corruption.

Set the parameters shown in [Table 20–2](#) in the auxiliary initialization parameter file.

Table 20–2 Auxiliary Initialization Parameters

Parameter	Purpose
DB_NAME	Names the auxiliary database. Leave the name of the auxiliary database <i>the same as</i> the primary database.
CONTROL_FILES	Identifies auxiliary control files. Set to the filename of the auxiliary control file. If the auxiliary database is on the same host as the primary database, make sure that the control file name is <i>different from</i> the primary database control file name.
DB_UNIQUE_NAME	Allows the auxiliary database to start even though it has the same name as the primary database. Set to any unique value, for example, = AUX. This parameter is only needed if the auxiliary and primary database are on the same host.
DB_FILE_NAME_CONVERT	Uses patterns to convert filenames for the datafiles of the auxiliary database. This parameter is only necessary if you are either restoring the auxiliary database on the same host as the primary host, or on a different host that uses different path names from the primary host.
LOG_FILE_NAME_CONVERT	Uses patterns to convert filenames for the online redo logs of the auxiliary database. This parameter is mandatory.

Table 20–2 (Cont.) Auxiliary Initialization Parameters

Parameter	Purpose
LOG_ARCHIVE_DEST_1	Specifies the default directory containing the archived redo logs required for recovery. This parameter specifies the location on the auxiliary host in which the archived logs will be located.
LOG_ARCHIVE_FORMAT	Specifies the format of the archived logs. You should use the same format setting used in the primary initialization parameter file.

Set other parameters as needed, including the parameters that allow you to connect as SYSDBA through Oracle Net.

For example, the auxiliary parameter file for a database on the same host as the primary could look like the following:

```
DB_NAME = prod1
CONTROL_FILES = /oracle/aux/control01.dbf
DB_UNIQUE_NAME = aux
DB_FILE_NAME_CONVERT=("/oracle/oradata/", "/aux/")
LOG_FILE_NAME_CONVERT=("/oracle/oradata/", "/aux/")
LOG_ARCHIVE_DEST_1 = 'LOCATION=/oracle/oradata/arch/'
LOG_ARCHIVE_FORMAT = arcr_%t_%s_%r.arc
```

The auxiliary parameter file for a database on a different host with the same path names as the primary could look like the following:

```
DB_NAME = prod1
# you do not need to set CONTROL_FILES or DB_FILE_NAME_CONVERT because the file
# system structure on both hosts is identical
LOG_FILE_NAME_CONVERT=("/oracle/oradata/", "/tmp/oradata/")
LOG_ARCHIVE_DEST_1 = 'LOCATION=/tmp/arch/'
LOG_ARCHIVE_FORMAT = arcr_%t_%s_%r.arc
```

Restoring and Recovering the Auxiliary Database in User-Managed TSPITR: Basic Steps

The procedure for restore and recovery of the auxiliary database differs depending on whether the auxiliary database is on the same host as the primary database. The examples in this section assume:

- You are performing TSPITR on production database called `prod1` located on host `prim_host`.
- The recovery set tablespaces are `users` and `tools`. Tablespace `users` contains datafile `/oracle/oradata/users01.dbf` and tablespace `tools` contains datafile `/fs2/tools01.dbf`.
- The auxiliary set contains the `SYSTEM` tablespace datafile `/oracle/oradata/system.dbf`, the undo tablespace datafile `/oracle/oradata/undo01.dbf`, and the control file `/oracle/oradata/control01.dbf`.
- The online redo logs are named `/oracle/oradata/redo01.log` and `/oracle/oradata/redo02.log`.
- All the primary database files are contained in `/oracle/oradata`

The different cases are described in the following sections:

- [Restoring and Recovering the Auxiliary Database on the Same Host](#)
- [Restoring the Auxiliary Database on a Different Host with the Same Path Names](#)
- [Restoring the Auxiliary Database on a Different Host with Different Path Names](#)

Restoring and Recovering the Auxiliary Database on the Same Host

The following examples assume the case in which you restore the auxiliary database to the same host as the primary database. In this scenario, all of the primary database files are contained in `/oracle/oradata`, and you want to restore the auxiliary database to `/oracle/oradata/aux`. So, you set `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT` to convert the filenames from `/oracle/oradata` to `/oracle/oradata/aux`.

Perform the following tasks to restore and recover the auxiliary database:

1. Restore the auxiliary set and the recovery set to a location different from that of the primary database. For example, assume that the auxiliary set consists of the following files:

```
/oracle/oradata/control01.dbf    # control file
/oracle/oradata/undo01.dbf      # datafile in undo tablespace
/oracle/oradata/system.dbf      # datafile in SYSTEM tablespace
```

And the recovery set consists of the following datafiles:

```
/oracle/oradata/users01.dbf     # datafile in users tablespace
/oracle/oradata/tools01.dbf     # datafile in tools tablespace
```

You can restore backups of the auxiliary set files and recovery set files to a new location as follows:

```
cp /backup/control01.dbf /oracle/oradata/aux/control01.dbf
cp /backup/undo01.dbf /oracle/oradata/aux/undo01.dbf
cp /backup/system.dbf /oracle/oradata/aux/system.dbf
cp /backup/users01.dbf /oracle/oradata/aux/users01.dbf
cp /backup/tools01.dbf /oracle/oradata/aux/tools01.dbf
```

2. Start the auxiliary database without mounting it, specifying the initialization parameter file if necessary. For example, enter:

```
STARTUP NOMOUNT PFILE=/aux/initAUX.ora
```

3. Mount the auxiliary database, specifying the `CLONE` keyword:

```
ALTER DATABASE MOUNT CLONE DATABASE;
```

The `CLONE` keyword causes Oracle to take all datafiles offline automatically.

4. Manually rename all auxiliary database files to reflect their new locations *only if* these files are not renamed by `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT`. In our scenario, all datafiles and online redo logs are renamed by initialization parameters, so no manual renaming is necessary.
5. Run the following SQL script on the auxiliary database to ensure that all datafiles are named correctly:

```
SELECT NAME FROM V$DATAFILE
UNION ALL
SELECT MEMBER FROM V$LOGFILE
UNION ALL
SELECT NAME FROM V$CONTROLFILE
```

/

If not, then rename the files manually as in the previous step.

- Bring only the datafiles in the auxiliary and recovery set tablespaces online. For example, bring the four datafiles in the recovery and auxiliary sets online:

```
ALTER DATABASE DATAFILE /oracle/oradata/aux/system.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/aux/users01.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/aux/tools01.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/aux/undo01.dbf ONLINE;
```

Note: The export phase of TSPITR will not work if all the files of each recovery set tablespace are not online.

At this point, the auxiliary database is mounted and ready for media recovery.

- Recover the auxiliary database to the specified point in time with the `USING BACKUP CONTROLFILE` option. Use any form of incomplete recovery. The following example uses cancel-based incomplete recovery:

```
RECOVER DATABASE UNTIL CANCEL USING BACKUP CONTROLFILE
```

- Open the auxiliary database with the `RESETLOGS` option using the following statement:

```
ALTER DATABASE OPEN RESETLOGS;
```

Restoring the Auxiliary Database on a Different Host with the Same Path Names

The following example assumes that you create the auxiliary database on a different host called `aux_host`. The auxiliary host has the same path names as the primary host. Hence, you do not need to rename the auxiliary database datafiles. So, you do *not* need to set `DB_FILE_NAME_CONVERT`, although you should set `LOG_FILE_NAME_CONVERT`.

To restore and recover the auxiliary database:

- Restore the auxiliary set and the recovery set to the auxiliary host. For example, assume that the auxiliary set consists of the following files:

```
/oracle/oradata/control01.dbf      # control file
/oracle/oradata/undo01.dbf        # datafile in undo tablespace
/oracle/oradata/system.dbf       # datafile in SYSTEM tablespace
```

And the recovery set consists of the following datafiles:

```
/oracle/oradata/users01.dbf # 1st datafile in users tablespace
/oracle/oradata/tools01.dbf # 2nd datafile in tools tablespace
```

These files will occupy the same locations in the auxiliary host.

- Start the auxiliary database without mounting it, specifying the initialization parameter file if necessary. For example, enter:

```
STARTUP NOMOUNT PFILE=/aux/initAUX.ora
```

- Mount the auxiliary database, specifying the `CLONE` keyword:

```
ALTER DATABASE MOUNT CLONE DATABASE;
```

The `CLONE` keyword causes Oracle to take all datafiles offline automatically.

4. Rename all auxiliary database files to reflect their new locations *only if* these files are not renamed by `DB_FILE_NAME_CONVERT` and `LOG_FILE_NAME_CONVERT`. In our scenario, the datafiles do not require renaming, and the logs are converted with `LOG_FILE_NAME_CONVERT`. So, no manual renaming is necessary.
5. Run the following script in SQL*Plus on the auxiliary database to ensure that all datafiles are named correctly.

```
SELECT NAME FROM V$DATAFILE
UNION ALL
SELECT MEMBER FROM V$LOGFILE
UNION ALL
SELECT NAME FROM V$CONTROLFILE
;
```

If not, then rename them manually as in the previous step.

6. Bring all datafiles in the auxiliary and recovery set tablespaces online. For example, bring the four datafiles in the recovery and auxiliary sets online:

```
ALTER DATABASE DATAFILE /oracle/oradata/system.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/users01.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/tools01.dbf ONLINE;
ALTER DATABASE DATAFILE /oracle/oradata/undo01.dbf ONLINE;
```

Note: The export phase of TSPITR will not work if all the files of each recovery set tablespace are not online.

At this point, the auxiliary database is mounted and ready for media recovery.

7. Recover the auxiliary database to the specified point in time with the `USING BACKUP CONTROLFILE` option. Use any form of incomplete recovery. The following example uses cancel-based incomplete recovery:

```
RECOVER DATABASE UNTIL CANCEL USING BACKUP CONTROLFILE
```

8. Open the auxiliary database with the `RESETLOGS` option using the following statement:

```
ALTER DATABASE OPEN RESETLOGS;
```

Restoring the Auxiliary Database on a Different Host with Different Path Names

This case should be treated exactly like ["Restoring and Recovering the Auxiliary Database on the Same Host"](#) on page 20-8. The same guidelines for renaming files apply in both cases.

Performing User-Managed TSPITR with Transportable Tablespaces

After you have completed the preparation stage, begin the actual TSPITR procedure as described in *Oracle Database Administrator's Guide*. The procedure occurs in the following steps:

- [Step 1: Unplugging the Tablespaces from the Auxiliary Database](#)
- [Step 2: Transporting the Tablespaces into the Primary Database](#)

Step 1: Unplugging the Tablespaces from the Auxiliary Database

In this step, you recover the auxiliary database to the desired past time, then unplug the desired tablespaces.

To unplug the auxiliary database tablespaces:

1. Connect SQL*Plus to the auxiliary database with administrator privileges. For example:

```
% sqlplus 'SYS/oracle@aux AS SYSDBA'
```

2. Make the tablespaces in the recovery set read-only by running the ALTER TABLESPACE . . . READ ONLY statement. For example, make `users` and `tools` read-only as follows:

```
ALTER TABLESPACE users READ ONLY;
ALTER TABLESPACE tools READ ONLY;
```

3. Ensure that the recovery set is self-contained. For example:

```
EXECUTE SYS.DBMS_TTS.TRANSPORT_SET_CHECK('users,tools',TRUE,TRUE);
```

4. Query the transportable tablespace violations table to manage any dependencies. For example:

```
SELECT * FROM SYS.TRANSPORT_SET_VIOLATIONS;
```

This query should return no rows after all dependencies are managed. Refer to *Oracle Database Administrator's Guide* for more information about this table.

5. Generate the transportable set by running the Export utility as described in *Oracle Database Administrator's Guide*. Include all tablespaces in the recovery set, as in the following example:

```
% exp SYS/oracle TRANSPORT_TABLESPACE=y TABLESPACES=(users,tools) \
TTS_FULL_CHECK=y
```

This command generates an export file named `expdat.dmp`.

Step 2: Transporting the Tablespaces into the Primary Database

In this step, you transport the recovery set tablespaces into the primary database.

To plug the recovery set tablespaces into the primary database:

1. Connect SQL*Plus to the primary database (*not* the auxiliary database). For example:

```
% sqlplus 'SYS/oracle@primary AS SYSDBA'
```

2. Drop the tablespaces in the recovery set with the DROP TABLESPACE statement. For example:

```
DROP TABLESPACE users INCLUDING CONTENTS;
DROP TABLESPACE tools INCLUDING CONTENTS;
```

3. Restore the recovery set datafiles from the auxiliary database to the recovery set file locations in the primary database. For example:

```
% cp /net/aux_host/aux/users01.dbf \
> /net/primary_host/oracle/oradata/users01.dbf
% cp /net/aux_host/aux/tools01.dbf \
> /net/primary_host/oracle/oradata/tools01.dbf
```

4. Move the export file `expdat.dmp` to the primary host. For example, enter:

```
% cp /net/aux_host/aux/expdat.dmp \  
> /net/primary_host/oracle/oradata/expdat.dmp
```
5. Plug in the transportable set into the primary database by running Import as described in *Oracle Database Administrator's Guide*. For example:

```
% imp SYS/oracle TRANSPORT_TABLESPACE=y FILE=expat.dmp  
DATAFILES=('/oracle/oradata/users01.dbf','/oracle/oradata/tools01.dbf')
```
6. Make the recovered tablespaces read/write by executing the `ALTER TABLESPACE READ WRITE` statement. For example:

```
ALTER TABLESPACE users READ WRITE;  
ALTER TABLESPACE tools READ WRITE;
```
7. Back up the recovered tablespaces with an operating system utility.

Caution: You must back up the tablespace because otherwise you might lose it. For example, a media failure occurs, but the archived logs from the last backup of the database do not logically link to the recovered tablespaces. If you attempt to recover any recovery set tablespaces from a backup taken before TSPITR, then recovery fails.

Performing Partial TSPITR of Partitioned Tables

Partitioned tables can span multiple tablespaces. Follow this procedure only if the recovery set does not fully contain all of the partitions.

This section describes how to perform partial TSPITR of partitioned tables that have a range that has not changed or expanded, and includes the following steps:

- [Step 1: Create a Table on the Primary Database for Each Partition Being Recovered](#)
- [Step 2: Drop the Indexes on the Partition Being Recovered](#)
- [Step 3: Exchange Partitions with Standalone Tables](#)
- [Step 4: Drop the Recovery Set Tablespace](#)
- [Step 5: Create Tables at Auxiliary Database](#)
- [Step 6: Drop Indexes on Partitions Being Recovered](#)
- [Step 7: Exchange Partitions with Standalone Tables on the Auxiliary Database](#)
- [Step 8: Transport the Recovery Set Tablespaces](#)
- [Step 9: Exchange Partitions with Standalone Tables on the Primary Database](#)
- [Step 10: Back Up the Recovered Tablespaces in the Primary Database](#)

Note: Often you have to recover the dropped partition along with recovering a partition whose range has expanded. Refer to ["Performing User-Managed TSPITR of Partitioned Tables With a Dropped Partition"](#) on page 20-14.

Step 1: Create a Table on the Primary Database for Each Partition Being Recovered

This table should have the exact same column names and column datatypes as the partitioned table you are recovering. Create the table using the following template:

```
CREATE TABLE new_table AS
  SELECT * FROM partitioned_table
  WHERE 1=2;
```

These tables are used to swap each recovery set partition (see ["Step 3: Exchange Partitions with Standalone Tables"](#) on page 20-13).

Note: The table and the partition must belong to the same schema.

Step 2: Drop the Indexes on the Partition Being Recovered

Drop the indexes on the partition you wish to recover, or create identical, nonpartitioned indexes that exist on the partition you wish to recover. If you drop the indexes on the partition being recovered, then you need to drop them on the auxiliary database (see ["Step 6: Drop Indexes on Partitions Being Recovered"](#) on page 20-13). Rebuild the indexes after TSPITR is complete.

Step 3: Exchange Partitions with Standalone Tables

Exchange each partition in the recovery set with its associated standalone table (created in ["Step 1: Create a Table on the Primary Database for Each Partition Being Recovered"](#) on page 20-13) by issuing the following statement, replacing the variables with the names of the appropriate objects:

```
ALTER TABLE table_name EXCHANGE PARTITION partition_name WITH TABLE table_name;
```

Step 4: Drop the Recovery Set Tablespace

On the primary database, drop each tablespace in the recovery set. For example, enter the following, replacing *tablespace_name* with the name of the tablespace:

```
DROP TABLESPACE tablespace_name INCLUDING CONTENTS;
```

Step 5: Create Tables at Auxiliary Database

After recovering the auxiliary database and opening it with the RESETLOGS option, create a table in the SYSTEM tablespace that has the same column names and column data types as the partitioned table you are recovering. You must create the table in the SYSTEM tablespace: otherwise, Oracle issues the ORA-01552 error.

Create a table for each partition you wish to recover. These tables are used later to swap each recovery set partition.

Note: The table and the partition must belong to the same schema.

Step 6: Drop Indexes on Partitions Being Recovered

Drop the indexes on the partition you wish to recover, or create identical, non-partitioned indexes that exist on the partition you wish to recover (on the table created in ["Step 1: Create a Table on the Primary Database for Each Partition Being Recovered"](#) on page 20-13).

Step 7: Exchange Partitions with Standalone Tables on the Auxiliary Database

For each partition in the auxiliary database recovery set, exchange the partitions with the standalone tables (created in ["Step 5: Create Tables at Auxiliary Database"](#) on page 20-13) by executing the following SQL script, replacing the variables with the appropriate object names:

```
ALTER TABLE partitioned_table_name
EXCHANGE PARTITION partition_name
WITH TABLE table_name;
```

Step 8: Transport the Recovery Set Tablespaces

Export the recovery set tablespaces from the auxiliary database and then import them into the primary database as described in ["Performing User-Managed TSPITR with Transportable Tablespaces"](#) on page 20-10.

Step 9: Exchange Partitions with Standalone Tables on the Primary Database

For each recovered partition on the primary database, swap its associated standalone table with the following statement, replacing the variables with the appropriate object names:

```
ALTER TABLE table_name EXCHANGE PARTITION partition_name WITH TABLE table_name;
```

If the associated indexes have been dropped, then re-create them.

Step 10: Back Up the Recovered Tablespaces in the Primary Database

Back up the recovered tablespaces on the primary database. Failure to do so results in loss of data in the event of media failure.

Performing User-Managed TSPITR of Partitioned Tables With a Dropped Partition

This section describes how to perform TSPITR on partitioned tables when a partition has been dropped. It includes the following steps:

- [Step 1: Find the Low and High Range of the Partition that Was Dropped](#)
- [Step 2: Create a Temporary Table](#)
- [Step 3: Exchange Partitions with Standalone Tables](#)
- [Step 4: Drop the Recovery Set Tablespace](#)
- [Step 5: Create Tables at the Auxiliary Database](#)
- [Step 6: Drop Indexes on Partitions Being Recovered](#)
- [Step 7: Exchange Partitions with Standalone Tables](#)
- [Step 8: Transport the Recovery Set Tablespaces](#)
- [Step 9: Insert Standalone Tables into Partitioned Tables](#)
- [Step 10: Back Up the Recovered Tablespaces in the Primary Database](#)

Step 1: Find the Low and High Range of the Partition that Was Dropped

When a partition is dropped, the range of the partition preceding it expands downwards. Therefore, there may be records in the preceding partition that should actually be in the dropped partition after it has been recovered. To ascertain this, run the following SQL script at the primary database, replacing the variables with the appropriate values:

```
SELECT * FROM partitioned_table
WHERE relevant_key
BETWEEN low_range_of_partition_that_was_dropped
AND high_range_of_partition_that_was_dropped;
```

Step 2: Create a Temporary Table

If any records are returned, then create a temporary table in which to store these records so that if necessary they can be inserted into the recovered partition later.

Step 3: Delete Records From the Partitioned Table

Delete all the records stored in the temporary table from the partitioned table.

```
DELETE FROM partitioned_table
WHERE relevant_key
BETWEEN low_range_of_partition_that_was_dropped
AND high_range_of_partition_that_was_dropped;
```

Step 4: Drop the Recovery Set Tablespace

On the primary database, drop each tablespace in the recovery set. For example, enter the following, replacing *tablespace_name* with the name of the tablespace:

```
DROP TABLESPACE tablespace_name INCLUDING CONTENTS;
```

Step 5: Create Tables at the Auxiliary Database

After opening the auxiliary database with the RESETLOGS option, create a table in the SYSTEM tablespace that has the same column names and column data types as the partitioned table you are recovering. You must create the table in the SYSTEM tablespace: otherwise, Oracle issues the ORA-01552 error. Create a table for each partition that you want to recover. These tables will be used later to swap each recovery set partition.

Step 6: Drop Indexes on Partitions Being Recovered

Drop the indexes on the partition you wish to recover, or create identical, nonpartitioned indexes that exist on the partition you wish to recover.

Step 7: Exchange Partitions with Standalone Tables

For each partition in the auxiliary recovery set, exchange the partitions into the standalone tables created in "[Step 5: Create Tables at the Auxiliary Database](#)" on page 20-15 by issuing the following statement, replacing the variables with the appropriate values:

```
ALTER TABLE partitioned_table_name
EXCHANGE PARTITION partition_name
WITH TABLE table_name;
```

Step 8: Transport the Recovery Set Tablespaces

Export the recovery set tablespaces from the auxiliary database and then import them into the primary database as described in "[Performing User-Managed TSPITR with Transportable Tablespaces](#)" on page 20-10.

Step 9: Insert Standalone Tables into Partitioned Tables

At this point you must insert the standalone tables into the partitioned tables; you can do this by first issuing the following statement, replacing the variables with the appropriate values:

```
ALTER TABLE table_name
  SPLIT PARTITION partition_name AT (key_value)
  INTO
  (PARTITION partition_1_name TABLESPACE tablespace_name,
   PARTITION partition_2_name TABLESPACE tablespace_name);
```

At this point, partition 2 is empty because keys in that range have already been deleted from the table.

Issue the following statement to swap the standalone table into the partition, replacing the variables with the appropriate values:

```
ALTER TABLE EXCHANGE PARTITION partition_name WITH TABLE table_name;
```

Now insert the records saved in "[Step 2: Create a Temporary Table](#)" on page 20-15 into the recovered partition (if desired).

Note: If the partition that has been dropped is the last partition in the table, then add it with the ALTER TABLE ADD PARTITION statement.

Step 10: Back Up the Recovered Tablespaces in the Primary Database

Back up the recovered tablespaces in the primary database. Failure to do so results in loss of data in the event of media failure.

Performing User-Managed TSPITR of Partitioned Tables When a Partition Has Split

This section describes how to recover partitioned tables when a partition has been split, and includes the following sections:

- [Step 1: Drop the Lower of the Two Partitions at the Primary Database](#)
- [Steps 2: Follow Same Procedure as for Partial TSPITR of Partitioned Tablespaces](#)

Step 1: Drop the Lower of the Two Partitions at the Primary Database

For each partition you wish to recover whose range has been split, drop the lower of the two partitions so that the higher expands downwards. In other words, the higher partition has the same range as before the split. For example, if P1 was split into partitions P1A and P1B, then P1B must be dropped, meaning that partition P1A now has the same range as P1.

For each partition that you wish to recover whose range has split, create a table that has exactly the same column names and column datatypes as the partitioned table you

are recovering. For example, execute the following, replacing the variables with the appropriate values:

```
CREATE TABLE new_table
AS
(
  SELECT *
  FROM partitioned_table
  WHERE 1=2
);
```

These tables will be used to exchange each recovery set partition in "[Step 3: Exchange Partitions with Standalone Tables](#)" on page 20-13.

Steps 2: Follow Same Procedure as for Partial TSPITR of Partitioned Tablespaces

Follow the same procedure as for "[Performing Partial TSPITR of Partitioned Tables](#)" on page 20-12, but skip the first step of this procedure: "[Step 1: Create a Table on the Primary Database for Each Partition Being Recovered](#)" on page 20-13. In other words, start with "[Step 2: Drop the Indexes on the Partition Being Recovered](#)" on page 20-13 and follow all subsequent steps.

Troubleshooting User-Managed Media Recovery

This chapter describes how to troubleshoot user-managed media recovery, that is, media recovery performed without using Recovery Manager (RMAN).

This chapter includes the following topics:

- [About User-Managed Media Recovery Problems](#)
- [Investigating the Media Recovery Problem: Phase 1](#)
- [Trying to Fix the Recovery Problem Without Corrupting Blocks: Phase 2](#)
- [Deciding Whether to Allow Recovery to Corrupt Blocks: Phase 3](#)
- [Allowing Recovery to Corrupt Blocks: Phase 4](#)
- [Performing Trial Recovery](#)

About User-Managed Media Recovery Problems

[Table 21–1, "Media Recovery Problems"](#) describes potential problems that can occur during media recovery.

Table 21–1 Media Recovery Problems

Problem	Description
Missing or misnamed archived log	Recovery stops because the database cannot find the archived log recorded in the control file.
When you attempt to open the database, error ORA-1113 indicates that a datafile needs media recovery	This error commonly occurs because: <ul style="list-style-type: none"> ■ You are performing incomplete recovery but failed to restore all needed datafile backups. ■ Incomplete recovery stopped before datafiles reached a consistent SCN. ■ You are recovering datafiles from an online backup, but not enough redo was applied to make the datafiles consistent. ■ You are performing recovery with a backup control file, and did not specify the location of a needed online redo log. ■ A datafile is undergoing media recovery when you attempt to open the database. ■ Datafiles needing recovery were not brought online before executing RECOVER DATABASE, and so were not recovered.
Redo record problems	Two possible cases are as follows: <ul style="list-style-type: none"> ■ Recovery stops because of failed consistency checks, a problem called stuck recovery. Stuck recovery can occur when an underlying operating system or storage system loses a write issued by the database during normal operation. ■ The database signals an internal error when applying the redo. This problem can be caused by an Oracle bug. If checksums are not being used, it can also be caused by corruptions to the redo or data blocks.
Corrupted archived logs	Logs may be corrupted while they are stored on or copied between storage systems. If DB_BLOCK_CHECKSUM is enabled, then the database usually signals checksum errors. If checksumming is not on, then log corruption may appear as a problem with redo.
Archived logs with incompatible parallel redo format	If you enable the parallel redo feature, then the database generates redo logs in a new format. Prior releases of Oracle are unable to apply parallel redo logs. However, releases prior to Oracle9i Release 2 (9.2) can detect the parallel redo format and indicate the inconsistency with the following error message: External error 00303, 00000, "cannot process Parallel Redo". See Also: <i>Oracle Database Performance Tuning Guide</i> to learn about the parallel redo feature
Corrupted data blocks	A datafile backup may have contained a corrupted data block, or the data block may become corrupted either during recovery or when it was copied to the backup. If checksums are being used, then the database signals a checksum error. Otherwise, the problem may also appear as a redo corruption.
Random problems	Memory corruptions and other transient problems can occur during recovery.

The symptoms of media recovery problems are usually external or internal errors signaled during recovery. For example, an external error indicates that a redo block or a data block has failed checksum verification checks. Internal errors can be caused by either bugs in the database or errors arising from the underlying operating system and hardware.

If media recovery encounters a problem while recovering a database backup, whether it is a stuck recovery problem or a problem during redo application, the database always stops and leaves the datafiles undergoing recovery in a consistent state, that is, at a consistent SCN preceding the failure. You can then do one of the following:

- Open the database read-only to investigate the problem.

- Open the database with the `RESETLOGS` option, as long as the requirements for opening `RESETLOGS` have been met. Note that the `RESETLOGS` restrictions apply to opening the standby database as well, because a standby database is updated by a form of media recovery.

In general, opening the database read-only or opening with the `RESETLOGS` option require all online datafiles to be recovered to the same SCN. If this requirement is not met, then the database may signal `ORA-1113` or other errors when you attempt to open. Some common causes of `ORA-1113` are described in [Table 21-1, "Media Recovery Problems"](#).

The basic methodology for responding to media recovery problems occurs in the following phases:

1. Try to identify the cause of the problem. Run a trial recovery if needed.
2. If the problem is related to missing redo logs or you suspect there is a redo log, memory, or data block corruption, then try to resolve it using the methods described in [Table 21-2](#).
3. If you cannot resolve the problem using the methods described in [Table 21-2](#), then do one of the following:
 - Open the database with the `RESETLOGS` option if you are recovering a whole database backup. If you have performed serial media recovery, then the database contains all the changes up to but not including the changes at the SCN where the corruption occurred. No changes from this SCN onward are in the recovered part of the database. If you have restored online backups, then opening `RESETLOGS` succeeds only if you have recovered through all the `ALTER . . . END BACKUP` operations in the redo stream.
 - Proceed with recovery by allowing media recovery to corrupt data blocks. After media recovery completes, try performing block media recovery using `RMAN`.
 - Call Oracle Support Services as a last resort.

See Also: ["Performing Disaster Recovery"](#) on page 7-10 to learn about block media recovery

Investigating the Media Recovery Problem: Phase 1

If media recovery encounters a problem, then obtain as much information as possible after recovery halts. You do not want to waste time fixing the wrong problem, which may in fact make matters worse.

The goal of this initial investigation is to determine whether the problem is caused by incorrect setup, corrupted redo logs, corrupted data blocks, memory corruption, or other problems. If you see a checksum error on a data block, then the data block is corrupted. If you see a checksum error on a redo log block, then the redo log is corrupted.

Sometimes the cause of a recovery problem can be difficult to determine. Nevertheless, the methods in this chapter allow you to quickly recover a database even when you do not completely understand the cause of the problem.

To investigate media recovery problems:

1. Examine the `alert.log` to see whether the error messages give general information about the nature of the problem. For example, does the `alert_`

SID.log indicate any checksum failures? Does the *alert_SID*.log indicate that media recovery may have to corrupt data blocks in order to continue?

2. Check the trace file generated by the Oracle process during recovery. It may contain additional error information.

Trying to Fix the Recovery Problem Without Corrupting Blocks: Phase 2

Depending on the type of media recovery problem you suspect, you have different solutions at your disposal. You can try one or a combination of the methods described in [Table 21–2](#). Note that these methods are fairly safe: in almost all cases, they should not cause any damage to the database.

Table 21–2 Media Recovery Solutions

If you suspect . . .	Then . . .
Missing/misnamed archived logs	Determine whether you entered the correct filename. If you did, then check to see whether the log is missing from the operating system. If it is missing, and you have a backup, then restore the backup and apply the log. If you do not have a backup, then if possible perform incomplete recovery up to the point of the missing log.
ORA-1113 for ALTER DATABASE OPEN	Review the causes of this error in Table 21–1, "Media Recovery Problems" . Make sure that all read/write datafiles requiring recovery are online. If you use a backup control file for recovery, then the control file and datafiles must be at a consistent SCN for the database to be opened. If you do not have the necessary redo, then you must re-create the control file.
Corrupt archived logs	The log is corrupted if the checksum verification on the log redo block fails. If <code>DB_BLOCK_CHECKSUM</code> is not enabled either during the recovery session or when the database generated the redo, then recovery problems may be caused by corrupted logs. If the log is corrupt and an alternate copy of the corrupt log is available, then try to apply it and see whether this tactic fixes the problem. The <code>DB_BLOCK_CHECKSUM</code> initialization parameter determines whether checksums are computed for redo log and data blocks.
Archived logs with incompatible parallel redo format	If you are running an Oracle release prior to Oracle9i Release 2, and if you are attempting to apply redo logs created with the parallel redo format, then you must do the following steps: <ol style="list-style-type: none"> 1. Upgrade the database to a later release. 2. Perform media recovery. 3. Shut down the database consistently and back up the database. 4. Downgrade the database to the original release. <p>See Also: <i>Oracle Database Performance Tuning Guide</i> to learn about the parallel redo feature</p>
Memory corruption or transient problems	You may be able to fix the problem by shutting down the database and restarting recovery. The database should be left in a consistent state if the second attempt also fails.
Corrupt data blocks	Restore and recover the datafile again with user-managed methods, or restore and recover individual data blocks with the <code>RMAN BLOCKRECOVER</code> command. This tactic may fix the problem. A data block is corrupted if the checksum verification on the block fails. If <code>DB_BLOCK_CHECKING</code> is disabled, a corrupted data block problem may appear as a redo problem. If you must proceed with recovery, then you may want to corrupt the block now and continue recovery, and use <code>RMAN</code> to perform block media recovery later.

If you cannot fix the problem with the methods described in [Table 21–2](#), then there may be no easy way to fix the problem without losing data. You have these options:

- Open the database with the `RESETLOGS` option (for whole database recovery). This solution discards all changes after the point where the redo problem occurred, but guarantees a logically consistent database.
- Allow media recovery to corrupt one or more data blocks and proceed with media recovery. This option will only succeed if the `alert_SID.log` indicates that recovery can continue if it is allowed to corrupt a data block, which should be the case for most recovery problems. This option is best if it is important to bring up the database quickly and recover all changes. If you are contemplating this option as a last resort, then proceed to ["Deciding Whether to Allow Recovery to Corrupt Blocks: Phase 3"](#) on page 21-5.

See Also: ["Performing Disaster Recovery"](#) on page 7-10 to learn how to perform block media recovery with the `BLOCKRECOVER` command

Deciding Whether to Allow Recovery to Corrupt Blocks: Phase 3

When media recovery encounters a problem, the `alert_SID.log` may indicate that recovery can continue if it is allowed to corrupt the data block causing the problem. The `alert_SID.log` always contains information about the block: its block type, block address, the tablespace it belongs to, and so forth. For blocks containing user data, the alert log may also report the data object number.

In this case, the database can proceed with recovery if it is allowed to mark the problem block as corrupt. Nevertheless, this response is not always advisable. For example, if the block is an important block in the `SYSTEM` tablespace, marking the block as corrupt can eventually prevent you from opening the recovered database. Another consideration is whether the recovery problem is isolated. If this problem is followed immediately by many other problems in the redo stream, then you may want to open the database with the `RESETLOGS` option.

For a block containing user data, you can usually query the database to find out which object or table owns this block. If the database is not open, then you should be able to open the database read-only, even if you are recovering a whole database backup. The following example cancels recovery and opens read-only:

```
CANCEL
ALTER DATABASE OPEN READ ONLY;
```

Assume that the data object number reported in the `alert_SID.log` is 8031. You can determine the owner, object name, and object type by issuing this query:

```
SELECT OWNER, OBJECT_NAME, SUBOBJECT_NAME, OBJECT_TYPE
FROM DBA_OBJECTS
WHERE DATA_OBJECT_ID = 8031;
```

To determine whether a recovery problem is isolated, you can run a diagnostic **trial recovery**, which scans the redo stream for problems but does not actually make any changes to the recovered database. If a trial recovery discovers any recovery problems, it reports them in the `alert_SID.log`. You can use the `RECOVER . . . TEST` statement to invoke trial recovery.

See Also: ["Performing Trial Recovery"](#) on page 21-6

After you have done these investigations, you can follow the guidelines in [Table 21-3](#) to decide whether to allow recovery to corrupt blocks.

Table 21–3 Guidelines for Allowing Recovery to Permit Corruption

If the problem is . . .	and the block is . . .	Then . . .
not isolated	n/a	You should probably open the database with the <code>RESETLOGS</code> option. This response is important for stuck recovery problems, because stuck recovery can be caused by the operating system or a storage system losing writes. If an operating system or storage system suddenly fails, it can cause stuck recovery problems on several blocks.
isolated	in the <code>SYSTEM</code> tablespace	Do not corrupt the block, because it may eventually prevent you from opening the database. However, sometimes data in the <code>SYSTEM</code> tablespace is unimportant. If you must corrupt a <code>SYSTEM</code> block and recover all changes, contact Oracle Support.
isolated	index data	Consider corrupting index blocks because the index can be rebuilt later after the database has been recovered.
isolated	user data	Decide based on the importance of the data. If you continue with datafile recovery and corrupt a block, you lose data in the block. However, you can use <code>RMAN</code> to perform block media recovery later after datafile recovery completes. If you open <code>RESETLOGS</code> , then the database is consistent but loses any changes made after the point where recovery was stopped.
isolated	rollback or undo data	Consider corrupting the rollback or undo block because it does not harm the database if the transactions that generated the undo are never rolled back. However, if those transactions are rolled back, then corrupting the undo block can cause problems. If you are unsure, then call Oracle Support.

See Also: ["Performing Trial Recovery"](#) on page 21-6 to learn how to perform trial recovery, and ["Allowing Recovery to Corrupt Blocks: Phase 4"](#) on page 21-6 if you decide to corrupt blocks

Allowing Recovery to Corrupt Blocks: Phase 4

If you decide to allow recovery to proceed in spite of block corruptions, then run the `RECOVER` command with the `ALLOW n CORRUPTION` clause, where *n* is the number of allowable corrupt blocks.

To allow recovery to corrupt blocks:

1. Ensure that all normal recovery preconditions are met. For example, if the database is open, then take tablespaces offline before attempting recovery.
2. Run the `RECOVER` command, allowing a single corruption, repeating as necessary for each corruption to be made. The following statements shows a valid example:

```
RECOVER DATABASE ALLOW 1 CORRUPTION
```

Performing Trial Recovery

When problems such as stuck recovery occur, you have a difficult choice. If the block is relatively unimportant, and if the problem is isolated, then it is better to corrupt the block. But if the problem is not isolated, then it may be better to open the database with the `RESETLOGS` option.

Because of this situation, the Oracle database supports trial recovery. A trial recovery applies redo in a way similar to normal media recovery, but it never writes its changes to disk and it always rolls back its changes. Trial recovery occurs only in memory.

See Also: ["Allowing Recovery to Corrupt Blocks: Phase 4"](#) on page 21-6

How Trial Recovery Works

By default, if a trial recovery encounters a stuck recovery or similar problem, then it always marks the data block as corrupt in memory when this action can allow recovery to proceed. The database writes errors generated during trial recovery to alert files. These errors are clearly marked as test run errors.

Like normal media recovery, trial recovery can prompt you for archived log filenames and ask you to apply them. Trial recovery ends when:

- The database runs out of the maximum number of buffers in memory that trial recovery is permitted to use
- An unrecoverable error is signaled, that is, an error that cannot be resolved by corrupting a data block
- You cancel or interrupt the recovery session
- The next redo record in the redo stream changes the control file
- All requested redo has been applied

When trial recovery ends, the database removes all effects of the test run from the system—except the possible error messages in the alert files. If the instance fails during trial recovery, then the database removes all effects of trial recovery from the system because trial recovery never writes changes to disk.

Trial recovery lets you foresee what problems might occur if you were to continue with normal recovery. For problems caused by ongoing memory corruption, trial recovery and normal recovery can encounter different errors.

Executing the RECOVER ... TEST Statement

You can use the TEST option for any RECOVER command. For example, you can start SQL*Plus and then issue any of the following commands:

```
RECOVER DATABASE TEST
RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL TEST
RECOVER TABLESPACE users TEST
RECOVER DATABASE UNTIL CANCEL TEST
```

By default, trial recovery always attempts to corrupt blocks in memory if this action allows trial recovery to proceed. In other words, trial recovery by default can corrupt an unlimited number of data blocks. You can specify the ALLOW *n* CORRUPTION clause on the RECOVER . . . TEST statement to limit the number of data blocks trial recovery can corrupt in memory.

A trial recovery command is usable in any scenario in which a normal recovery command is usable. Nevertheless, you should only need to run trial recovery when recovery runs into problems.

A

ABORT option
 SHUTDOWN statement, 18-6, 18-7, 18-20, 18-27, 18-28

active online redo log
 loss of group, 19-10, 19-11

alert log, 19-13
 checking after RESETLOGS, 18-27
 useful for RMAN, 12-2

ALLOW ... CORRUPTION clause
 RECOVER command, 21-6

ALTER DATABASE statement
 BACKUP CONTROLFILE clause, 17-11
 TO, 17-11
 CLEAR LOGFILE clause, 19-10
 END BACKUP clause, 17-8
 NORESETLOGS option, 18-26
 OPEN RESETLOGS clause, 10-8
 RECOVER clause, 18-5, 18-14
 RESETLOGS option, 18-26, 18-28, 18-29

ALTER SYSTEM statement
 KILL SESSION clause, 12-9
 RESUME clause, 17-14
 SUSPEND clause, 17-14

ALTER TABLESPACE statement
 BEGIN BACKUP clause, 17-5, 17-7
 END BACKUP option, 17-7

archived redo logs
 applying during media recovery, 18-11
 automating application, 18-12
 backing up, 6-15
 cataloging, 10-6
 changing default location, 18-14
 corrupted, 21-2
 deleting after recovery, 18-6
 deletion after backup, 2-11, 9-7
 deletion after restore, 3-6
 errors during recovery, 18-15
 incompatible format, 21-2
 location during recovery, 18-11
 loss of, 19-11
 restoring, 18-5
 RMAN fails to delete, 12-19
 using for recovery
 in default location, 18-13

 in nondefault location, 18-14

ARCHIVELOG mode
 datafile loss in, 19-2

AS SELECT clause
 CREATE TABLE statement, 19-4

autobackups
 control file, 2-28, 2-29
 server parameter file, 2-28

automatic channels, 2-1
 allocation, 2-2
 configuring, 5-9, 5-12, 6-1
 generic
 configuring, 2-5, 5-9
 definition, 5-10
 naming conventions, 2-5
 overriding, 5-9
 parallelism, 5-9
 specific configurations, 2-6

AUTORECOVERY option
 SET statement, 18-11

auxiliary instance parameter file
 with TRANSPORT TABLESPACE, 14-10

B

BACKUP
 and corrupt datafile blocks and
 MAXCORRUPT, 2-42

AS COPY
 and DB_FILE_NAME_CONVERT, 2-19
 and output filenames, 2-19
 FORMAT, 2-19

image copy
 and output filenames, 2-19

INCREMENTAL
 FROM SCN, 13-24

INCREMENTAL FROM SCN, 13-24

output filenames
 image copies, 2-19

BACKUP command, 6-22

BACKUPSET option, 2-16, 6-5

DELETE INPUT option, 9-7

FORMAT parameter, 2-18

KEEP option, 2-34

NOT BACKED UP SINCE clause, 2-39, 6-10

PROXY ONLY option, 2-11

- PROXY option, 2-10
 - SKIP OFFLINE option, 6-13
 - VALIDATE option, 2-44
- BACKUP CONTROLFILE clause
 - of ALTER DATABASE, 17-2
- BACKUP CONTROLFILE TO TRACE clause
 - of ALTER DATABASE, 17-2
- BACKUP COPIES parameter
 - CONFIGURE command, 5-17
- backup encryption, 6-7
 - dual-mode, 6-8
 - overview, 6-7
 - password, 6-8
 - transparent, 6-7
- backup mode
 - ending with ALTER DATABASE END
 - BACKUP, 17-8
 - for online user-managed backups, 17-5
 - instance failure, 17-7
- backup optimization
 - configuring, 5-16
 - definition, 2-35, 6-10
 - disabling, 2-37, 5-16
 - enabling, 2-37, 5-16
 - recovery window and, 2-38
 - redundancy and, 2-38
 - retention policies and, 2-37
- BACKUP OPTIMIZATION option
 - of CONFIGURE, 6-10
- backup retention policy
 - definition, 2-30
- backup sets
 - backing up, 2-16, 6-5
 - configuring maximum size, 5-15
 - crosschecking, 9-4
 - duplexing, 6-2
 - errors during creation, 2-41
 - failover during backups, 2-17
 - how RMAN generates, 2-22
 - limiting size, 2-22
 - multiplexing, 2-12
 - naming, 2-18
 - specifying maximum size (in bytes), 2-21
 - specifying number of, 2-22
- Backup Solutions Program (BSP), 1-8
- backups
 - archived redo logs, 6-15
 - deletion after backing up, 9-7
 - availability
 - altering with CHANGE command, 9-9
 - backup sets, 6-5
 - backups of, 2-16
 - closed, 17-3
 - consistent, 17-3
 - control files, 17-10
 - binary, 17-10
 - trace files, 17-11
 - correlating RMAN channels with, 9-12
 - cumulative incremental, 2-26, 3-10
 - datafile
 - using RMAN, 6-5, 6-11
 - DBVERIFY utility, 17-18
 - deleting, 9-6
 - determining datafile status, 17-2
 - duplexing, 5-17, 6-2
 - excluding tablespaces from backups, 5-18
 - failed RMAN, 12-20
 - failover during BACKUP BACKUPSET, 2-17
 - hung, 12-13
 - image copies, 2-9
 - inconsistent, 17-3
 - incremental, 2-24
 - differential, 2-25
 - using RMAN, 6-2, 6-3
 - interrupted, 6-10
 - keeping, 6-19
 - keeping records, 17-20
 - limiting I/O rate, 2-23
 - listing files needed, 17-1
 - logical, 17-19
 - long-term, 2-34
 - changing status, 9-10
 - multiple copies, 5-17
 - NOARCHIVELOG mode, in, 6-18
 - obsolete
 - batch deletes, 2-33
 - offline datafiles, 17-4
 - offline tablespaces, 17-4
 - optimizing, 2-35
 - read-only tablespaces, 17-9
 - recovery catalog, 1-7, 10-17
 - restartable, 2-39, 6-10
 - restoring user-managed, 18-2
 - RMAN error handling, 6-21
 - specifying number of files in a backup set, 2-22
 - split mirror, 2-10
 - using RMAN, 6-3
 - stored scripts, 10-13
 - tablespace, 17-6
 - using RMAN, 6-5, 6-11
 - tags, 2-19
 - testing RMAN, 2-44, 6-11
 - using media manager, 5-7
 - troubleshooting failed RMAN, 12-12, 12-15, 12-18
 - types, 2-23
 - Updating standby databases with
 - incrementals, 13-24
 - user-managed, 17-1
 - restoring, 18-4
 - validating, 6-11
 - verifying, 17-18
 - whole database
 - preparing for, 17-3
- BEGIN BACKUP clause
 - ALTER TABLESPACE statement, 17-5
- block corruptions
 - stored in V\$DATABASE_BLOCK_
 - CORRUPTION, 6-12
- block media recovery, 7-12
 - guidelines, 3-8

overview, 3-7
BLOCKRECOVER command, 3-7, 7-12
BSP. *See* Backup Solutions Program (BSP)

C

cancel-based media recovery
 procedures, 18-18, 18-23
canceling RMAN commands, 12-9
CATALOG command, 10-6
cataloging
 archived redo logs, 10-6
 datafiles, 10-6
catalog.sql script, 10-3
catproc.sql script, 10-3
CHANGE command, 9-4
 AVAILABLE option, 9-9
 KEEP option, 9-10
change-based media recovery
 coordinated in dis, 19-13
channels
 allocating manually for backups, 6-1
 configuring automatic, 5-9
 configuring for backups, 6-1
 control options, 2-8
 definition, 2-1
 difference between manual and automatic, 2-2
 generic configurations, 2-5
 overriding automatic, 5-9
 parallelism for manual channels, 2-7
 preconfigured disk, 5-9
 Recovery Manager, 2-1
 RMAN naming conventions, 2-5
 specific configurations, 2-6
circular reuse records, 1-5
CLEAR LOGFILE clause
 of ALTER DATABASE, 19-10
clearing RMAN configuration, 2-6, 5-14
clone databases
 preparing for TSPITR, 20-8, 20-9
cold failover cluster
 definition, 17-8
command files
 Recovery Manager, 1-3
command interface
 RMAN, 1-2
commands, Recovery Manager
 BACKUP, 6-22
 PROXY ONLY option, 2-11
 PROXY option, 2-10
 SKIP OFFLINE option, 6-13
 batch execution, 1-3
 CATALOG, 10-6
 CHANGE, 9-4
 CONFIGURE, 2-5, 2-6
 DELETE, 9-6
 DROP CATALOG, 10-26
 DUPLICATE, 3-9
 EXECUTE SCRIPT, 10-13
 how RMAN interprets, 1-2
 interactive, 1-3
 LIST, 9-1
 INCARNATION option, 10-8
 piping, 4-5
 RECOVER, 3-4
 RESET DATABASE
 INCARNATION option, 10-8
 RESYNC CATALOG, 10-9
 FROM CONTROLFILECOPY option, 10-20
 SET
 MAXCORRUPT option, 6-22
 SHOW, 2-4
 standalone, 1-4
 terminating, 12-9
 UPGRADE CATALOG, 10-25
commands, SQL
 ALTER DATABASE, 18-5, 18-14
commands, SQL*Plus
 RECOVER
 UNTIL TIME option, 18-24
 SET, 18-5, 18-11, 18-14
compatibility
 recovery catalog, 1-7
compilation
 and execution of RMAN commands, 1-2
complete recovery
 procedures, 18-15
CONFIGURE command
 BACKUP OPTIMIZATION option, 5-16
 CHANNEL option, 2-5
 CLEAR option, 2-6, 5-14
 DEFAULT DEVICE TYPE clause, 2-4
 DEVICE TYPE clause, 2-3
 EXCLUDE option, 5-18
 RETENTION POLICY clause, 2-30, 2-31
configuring
 media manager
 installing, 5-4
 prerequisites, 5-4
 media managers for use with RMAN, 5-5
 Recovery Manager
 autobackups, 2-28
 automatic channels, 5-9
 backup optimization, 5-16
 backup set size, 5-15
 clearing, 2-6, 5-14
 default device types, 2-4
 device types, 2-3
 parallelism, 2-3
 shared server, 5-22
 snapshot control file location, 5-20
 specific channels, 5-12
 tablespace exclusion for backups, 5-18
consistent backups
 whole database, 17-3
control file autobackups
 after structural changes to database, 2-28
 configuring, 2-28
 default format, 2-29
 restoring, 2-28

- control files
 - automatic backups, 2-29
 - configuring, 2-28
 - backing up to trace file, 17-11
 - backups, 17-2, 17-10
 - binary, 17-10
 - recovery using, 7-7
 - trace files, 17-11
 - creating after loss of all copies, 18-10
 - duplicate database, 13-4
 - finding filenames, 17-2
 - multiplexed
 - loss of, 18-6
 - recreated, 18-9
 - restoring
 - to default location, 18-6
 - to nondefault location, 18-7
 - using SET DBID, 7-15
 - snapshot
 - specifying location of, 5-20
 - time-based recovery, 18-21
 - types of records, 1-5
 - user-managed restore after loss of all copies, 18-9
 - using instead of a recovery catalog, 1-5
- CONTROL_FILES initialization parameter, 8-20, 18-7
- CONVERT command
 - and DB_FILE_NAME_CONVERT, 15-3
 - moving data to/from ASM, 15-16
 - with databases, 15-7
 - with tablespaces and datafiles, 15-1
- CONVERT DATABASE, 15-7
- CONVERT DATAFILE
 - moving to/from ASM, 15-16
- CONVERT DATAFILE or TABLESPACE, 15-1
- CONVERT TABLESPACE
 - moving to/from ASM, 15-16
- coordinated time-based recovery
 - distributed databases, 19-13
- COPIES option
 - of BACKUP, 6-3
- corrupt datafile blocks, 2-43
 - detecting, 2-43
 - records in control file, 2-42
 - recovering, 7-14
 - RMAN and, 2-41
 - setting maximum for backup, 6-22
- corrupt datafile blocks during backup, 2-42
- corruption detection, 2-43
- CREATE DATAFILE clause
 - of ALTER DATABASE, 19-3
- CREATE TABLE statement
 - AS SELECT clause, 19-4
- CREATE TABLESPACE statement, 19-2
- creating
 - test databases, 3-9
- crosschecking
 - definition, 9-4
 - recovery catalog with the media manager, 9-4
- cross-platform transportable database, 15-7
 - transport script, 15-12
- cross-platform transportable tablespace, 15-1
- cumulative incremental backups, 2-26

D

- data blocks
 - corrupted, 21-2
- data dictionary views, 17-4, 17-5, 17-9
- Data Pump Export utility, 17-19
 - backups, 17-19
- Data Pump Import utility, 17-19
- database connections
 - Recovery Manager
 - auxiliary database, 4-2
 - hiding passwords, 4-4
- database incarnation, 18-24
- database point-in-time recovery (DBPITR)
 - user-managed, 18-20
- databases
 - listing for backups, 17-1
 - media recovery procedures, user-managed, 18-1
 - media recovery scenarios, 19-1
 - recovery
 - after control file damage, 18-6, 18-7
 - registering in recovery catalog, 10-4, 10-5
 - suspending, 17-13
 - unregistering in recovery catalog, 10-7
- datafile recovery
 - definition, 3-4
- datafiles
 - backing up
 - offline, 17-4
 - using Recovery Manager, 6-5, 6-11
 - cataloging, 10-6
 - determining status, 17-2
 - duplicate database, 13-5
 - listing
 - for backup, 17-1
 - losing, 19-1
 - in ARCHIVELOG mode, 19-2
 - in NOARCHIVELOG mode, 19-1
 - recovery
 - basic steps, 3-4
 - without backup, 19-3
 - re-creating, 19-3
 - renaming
 - after recovery, 19-3
 - restoring, 3-1, 18-4
 - to default location, 18-4
- db identifier
 - problems registering copied database, 10-6
 - setting during disaster recovery, 7-9
 - setting with DBNEWID, 10-6
- DB_FILE_NAME_CONVERT
 - and CONVERT command, 15-3
- DB_FILE_NAME_CONVERT initialization
 - parameter, 8-20, 20-6
 - using with RMAN DUPLICATE command, 13-6, 13-7

- DB_NAME initialization parameter, 8-19
- DBA_DATA_FILES view, 17-4, 17-5, 17-9
- DBMS_PIPE package, 4-5
 - using with RMAN, 4-5
- DBNEWID utility, 10-6
- DBVERIFY utility, 17-18
- DELETE command, 9-6
 - OBSOLETE option, 2-33
- deleting
 - expired backups, 9-6
 - files after backups, 9-7
 - obsolete backups, 9-6
 - using RMAN, 9-5
- device types
 - configuring in RMAN, 5-9
- differential incremental backups, 2-25
- disk API, 5-5
- disk channels
 - preconfigured, 5-9
- distributed databases
 - change-based recovery, 19-13
 - coordinated time-based recovery, 19-13
 - recovery, 19-13
- dropping the recovery catalog, 10-26
- dual mode backup encryption, 6-8
- dummy API, 5-5
- duplexing
 - backup sets, 2-15, 5-17, 6-2
- DUPLICATE command, 3-9
- duplicate database
 - synchronizing
 - using DUPLICATE DATABASE, 13-23
 - using incremental backups, 13-24
- duplicate databases
 - creating, 3-9
 - on local host, 13-16
 - on remote host with different file system, 13-12
 - on remote host with same file system, 13-12
 - past point-in-time, 13-22
 - using CONFIGURE AUXNAME, 13-15
 - using init.ora parameter and LOGFILE, 13-13
 - using SET NEWNAME, 13-14
 - datafiles, 13-5
 - excluding tablespaces, 3-10, 13-3
 - failed creation, 12-22
 - generating control files, 13-4
 - generating filenames, 13-4
 - how RMAN creates, 13-2
 - NOFILENAMECHECK option, 13-6
 - preparing for duplication, 13-8
 - skipping offline normal tablespaces, 13-7
 - skipping read-only tablespaces, 13-7
- duplicating a database, 3-9
 - troubleshooting, 12-22

E

- endian formats
 - and CONVERT DATAFILE/TABLESPACE, 15-1

- environment, Recovery Manager
 - definition, 1-1
- error codes
 - media manager, 12-3
 - RMAN, 12-1, 12-2
 - message numbers, 12-3
- error messages
 - Recovery Manager
 - interpreting, 12-5
- error stacks
 - interpreting, 12-5
- errors
 - during RMAN backups, 6-21
- EXCLUDE option
 - of CONFIGURE, 5-18
- expired backups
 - deleting, 9-6

F

- FAST_START_MTTR_TARGET
 - and tuning instance recovery, 11-12
- Fast-Start checkpointing architecture, 11-10
- Fast-Start Fault Recovery, 11-9, 11-10
- features, new, 0-xxv
- fileNameConversionSpec
 - and CONVERT command, 15-3
- filenames
 - listing for backup, 17-1
- FORCE option
 - DELETE command, 9-9
- fractured blocks
 - definition, 2-44
 - detection, 2-44

G

- generic channels
 - definition, 5-10
- groups
 - archived redo log, 19-7, 19-8
 - online redo log, 19-7, 19-8

H

- hot backup mode
 - for online user-managed backups, 17-6
- hot backups
 - failed, 17-7
 - ending with ALTER DATABASE END BACKUP, 17-8

I

- image copies, 2-9
 - and SWITCH commands, 2-10
- inactive online redo log
 - loss of, 19-9
- INCARNATION option
 - of LIST, 10-8
 - of RESET DATABASE, 10-8

- incomplete media recovery, 18-20
 - in Oracle Real Application Clusters
 - configuration, 18-12
 - time-based, 18-23
 - with backup control file, 18-12
- incomplete recovery
 - overview, 3-7
- incremental backups
 - differential, 2-25
 - how RMAN applies, 3-5
 - using RMAN, 6-2, 6-3
- Incremental Roll Forward of Database Copy, 13-24
- initialization parameter file, 3-4
- initialization parameters
 - CONTROL_FILES, 18-7
 - DB_FILE_NAME_CONVERT, 8-20
 - DB_NAME, 8-19
 - LARGE_POOL_SIZE, 11-7
 - LOCK_NAME_SPACE, 8-19
 - LOG_ARCHIVE_DEST_*n*, 18-13
 - LOG_ARCHIVE_FORMAT, 18-13
 - LOG_FILE_NAME_CONVERT, 8-19
 - RECOVERY_PARALLELISM, 18-29
- instance failures
 - in backup mode, 17-7
- instance recovery
 - Fast-Start Fault Recovery, 11-10
 - performance tuning, 11-9
- integrity checks, 2-42
- interpreting RMAN error stacks, 12-5
- interrupting media recovery, 18-15
- I/O errors
 - effect on backups, 2-41
 - ignoring during deletions, 9-9

J

- jobs
 - RMAN
 - monitoring performance, 9-16
 - monitoring progress, 9-13

K

- KEEP option
 - of BACKUP, 2-34
 - of CHANGE, 9-10

L

- level 0 incremental backups, 2-25
- LIST command, 9-1
 - INCARNATION option, 10-8
- LOCK_NAME_SPACE initialization parameter, 8-19
- log sequence numbers
 - requested during, 18-11
- LOG_ARCHIVE_DEST_*n* initialization
 - parameter, 18-13, 20-7
- LOG_ARCHIVE_FORMAT initialization
 - parameter, 18-13, 20-7
- LOG_FILE_NAME_CONVERT initialization

- parameter, 8-19, 20-6
- logical backups, 17-19
- LOGSOURCE variable
 - SET statement, 18-5, 18-14
- long waits
 - defined, 11-8
- long-term backups
 - changing status, 9-10
 - definition, 2-34
- loss of
 - inactive log group, 19-9

M

- managing RMAN metadata, 10-1, 13-1
- MAXCORRUPT setting, 2-42
- MAXPIECESIZE parameter
 - SET command, 5-6
- MAXSETSIZE parameter
 - BACKUP command, 5-15
 - CONFIGURE command, 5-15
- MAXSIZE parameter
 - RECOVER command, 3-6
- mean time to recovery (MTTR)
 - definition, 3-7
- media failures
 - archived redo log file loss, 19-11
 - complete recovery, 18-15
 - complete recovery, user-managed, 18-15
 - control file loss, 18-9
 - datafile loss, 19-1
 - NOARCHIVELOG mode, 18-27
 - online redo log group loss, 19-8
 - online redo log loss, 19-7
 - online redo log member loss, 19-7
 - recovery, 18-15
 - distributed databases, 19-13
 - recovery procedures
 - examples, 19-1
- media management
 - backing up files, 1-8
 - Backup Solutions Program, 1-8
 - crosschecking, 9-4
 - error codes, 12-3
 - linking to software, 5-4
 - sbttest program, 12-7
 - testing the API, 12-7
- media managers
 - configuring for use with RMAN, 5-5
 - installing, 5-4
 - linking
 - testing, 5-5
 - prerequisites for configuring, 5-4
 - testing, 5-5
 - testing backups, 5-7
 - troubleshooting, 5-7
- media recovery
 - ADD DATAFILE operation, 19-2
 - after control file damage, 18-6, 18-7
 - applying archived redo logs, 18-11

- cancel-based, 18-18, 18-20, 18-23
- ch, 18-20
- complete, 18-15
 - closed database, 18-16
- complete, user-managed, 18-15
- corruption
 - allowing to occur, 21-5
- datafiles
 - basic steps, 3-4
 - without backup, 19-3
- distributed databases, 19-13
- errors, 18-15, 21-2
- incomplete, 18-20
- interrupting, 18-15
- lost files
 - lost archived redo log files, 19-11
 - lost datafiles, 19-1
 - lost mirrored control files, 18-6
- NOARCHIVELOG mode, 18-27
- offline tablespaces in open database, 18-18
- online redo log files, 19-6
- opening database after, 18-24, 18-26
- parallel, 18-29
- problems, 21-1, 21-2
 - fixing, 21-4
 - investigating, 21-3
- restarting, 18-15
- restoring
 - archived redo log files, 18-5
 - whole database backups, 18-27
- resuming after interruption, 18-15
- roll forward phase, 18-11
- scenarios, 19-1
- time-based, 18-20
- transportable tablespaces, 19-6
- trial, 21-6
 - explanation, 21-7
 - overview, 21-6
- troubleshooting, 21-1
 - basic methodology, 21-3
- types
 - distributed databases, 19-13
- undamaged tablespaces online, 18-18
- unsuccessfully applied redo logs, 18-15
- using Recovery Manager, 3-4
- media recovery, user-managed, 18-1
- metadata
 - managing RMAN, 1-4, 10-1, 13-1
 - querying RMAN, 9-1
 - storing in control file, 1-5
- mirrored files
 - online redo log
 - loss of, 19-7
 - splitting, 17-13
 - suspend/resume mode, 17-13
 - using RMAN, 6-3
- mirroring
 - backups using, 6-3
- modes
 - NOARCHIVELOG

- recovery from failure, 18-27
- monitoring RMAN, 9-10
- MOUNT option
 - STARTUP statement, 18-22
- multiplexed files
 - control files
 - loss of, 18-6
- multiplexing
 - datafiles with Recovery Manager, 2-12

N

- naming backup sets, 2-18
- new features, 0-xxv
- NOARCHIVELOG mode
 - backing up, 6-18
 - datafile loss in, 19-1
 - disadvantages, 18-27
 - recovery, 18-27
- noncircular reuse records, 1-5
- NOT BACKED UP SINCE clause
 - BACKUP command, 6-10
- not feasible to test 1 oct 04 nocheck, 5-5

O

- obsolete backups
 - deleting, 2-33, 9-6
 - different from expired backups, 2-30
 - reporting, 9-2
- online redo logs, 19-9
 - active group, 19-7, 19-8
 - applying during media recovery, 18-11
 - archived group, 19-7, 19-8
 - clearing
 - failure, 19-10
 - clearing inactive logs
 - archived, 19-9
 - unarchived, 19-9
 - current group, 19-7, 19-8
 - determining active logs, 19-8
 - inactive group, 19-7, 19-8
 - listing log files for backup, 17-2
 - loss of
 - active group, 19-10, 19-11
 - all members, 19-8
 - group, 19-8
 - mirrored members, 19-7
 - recovery, 19-6
 - multiple group loss, 19-11
 - replacing damaged member, 19-7
 - status of members, 19-7, 19-8
- online tablespace transport, 14-1
- OPEN RESETLOGS clause
 - ALTER DATABASE statement, 10-8
- operating system copies
 - definition, 2-10
- ORA-01578 error message, 19-4
- Oracle Encryption Wallet
 - and backups, 6-7

P

- packages
 - DBMS_PIPE, 4-5
- parallel recovery, 18-29
- parallelism
 - backups, 2-13
 - configuring RMAN, 2-3, 5-9
 - manually allocated RMAN channels, 2-7
- partitioned tables
 - dropped partitions, 20-14
 - performing partial TSPITR, 20-12
 - split partitions, 20-16
- password backup encryption, 6-8
- password files
 - connecting to Recovery Manager with, 4-1
- passwords
 - connecting to RMAN, 4-4
- performance tuning
 - backup performance, 11-6
 - disk bandwidth and RATE channel parameter, 11-5
 - Fast-Start Fault Recovery, 11-9
 - instance recovery, 11-9
 - FAST_START_MTTR_TARGET, 11-10
 - setting FAST_START_MTTR_TARGET, 11-12
 - using V\$INSTANCE_RECOVERY, 11-11
 - LARGE_POOL_SIZE initialization parameter, 11-7
 - long waits
 - defined, 11-8
 - short waits
 - definition of, 11-8
- pipe interface, 4-5
- point of recoverability
 - recovery window, 2-31
- point-in-time recovery, 18-20
 - tablespace, 8-1 to ??, 8-3, 8-4 to ??, 20-12
 - user-managed, 20-1
- PROXY ONLY option
 - of BACKUP, 2-11
- PROXY option
 - of BACKUP, 2-10

R

- RATE option
 - of ALLOCATE CHANNEL, 2-23
 - of CONFIGURE CHANNEL, 2-23
- raw devices
 - backing up to, 17-15
 - restoring to, 18-4
 - UNIX backups, 17-15
 - Windows backups, 17-17
- read-only tablespaces
 - backing up, 6-13
 - backups, 17-9
- Recov, 10-8
- RECOVER clause
 - of ALTER DATABASE, 18-5, 18-14
- RECOVER command, 3-4, 3-6

- unrecoverable objects and standby databases, 19-4
- UNTIL TIME option, 18-24
- USING BACKUP CONTROLFILE clause, 19-5
- RECOVER SQL*Plus statement
 - PARALLEL and NOPARALLEL options, 18-29
- recovery
 - ADD DATAFILE operation, 19-2
 - automatically applying archived logs, 18-11
 - cancel-based, 18-18, 18-23
 - complete, 18-15
 - closed database, 18-16
 - offline tablespaces, 18-18
 - corruption
 - intentionally allowing, 21-5
 - data blocks, 3-7, 7-12
 - guidelines, 3-8
 - database
 - in NOARCHIVELOG mode, 7-1
 - database files
 - how RMAN applies changes, 3-5
 - overview, 3-4
 - datafile without a backup, 7-16
 - datafiles, 19-1
 - ARCHIVELOG mode, 19-2
 - NOARCHIVELOG mode, 19-1
 - determining files needing recovery, 18-3
 - disaster using RMAN, 7-10
 - dropped table, 19-12
 - errors, 21-2
 - interrupting, 18-15
 - media, 18-1, 19-1, 21-1
 - multiple redo threads, 18-12
 - of lost or damaged recovery catalog, 10-19
 - online redo logs, 19-6
 - losing member, 19-7
 - loss of group, 19-8
 - opening database after, 18-24
 - parallel, 18-29
 - parallel processes for, 18-29
 - problems, 21-1
 - fixing, 21-4
 - investigating, 21-3
 - responding to unsuccessful, 18-15
 - setting number of processes to use, 18-29
 - stuck, 21-2
 - time-based, 18-23
 - transportable tablespaces, 19-6
 - trial, 21-6
 - explanation, 21-7
 - overview, 21-6
 - troubleshooting, 21-1
 - user errors, 19-12
 - user-managed, 18-1, 19-1, 21-1
 - using backup control file, 7-7
 - without recovery catalog, 7-8
 - using logs in a nondefault location, 18-14
 - using logs in default location, 18-13
 - using logs in nondefault location, 18-14
 - without a recovery catalog, 1-6

- recovery catalog, 1-6
 - availability, 10-22
 - backing up, 1-7, 10-17
 - compatibility, 1-7
 - contents, 1-6
 - crosschecking, 9-4
 - db identifier problems, 10-6
 - dropping, 10-26
 - managing size of, 10-13
 - moving to new database, 10-20
 - operating with, 1-5
 - operating without, 1-5
 - recovery of, 10-19
 - refreshing, 10-9
 - registering target databases, 1-6, 10-4, 10-5
 - resynchronizing, 10-9
 - snapshot control file, 1-7
 - space requirements, 10-2
 - stored scripts
 - creating, 10-13
 - synchronization, 1-7
 - UNKNOWN database name, 12-23
 - unregistering databases, 10-7
 - updating
 - after schema changes, 10-11
 - upgrading, 10-25
 - views
 - querying, 10-22
- Recovery Manager
 - allocating disk buffers, 11-2
 - allocating tape buffers, 11-2
 - backup sets
 - backing up, 6-5
 - backup types
 - duplexed backup sets, 2-15
 - backups
 - backing up, 2-16
 - batch deletion of obsolete, 2-33
 - control file autobackups, 2-29
 - datafile, 6-5, 6-11
 - image copy, 2-9
 - incremental, 6-2, 6-3
 - long-term, 2-34
 - optimization, 2-35
 - restartable, 2-39
 - tablespace, 6-5, 6-11
 - testing, 2-44, 6-11
 - types, 2-23
 - using tags, 2-19
 - validating, 6-11
 - channels, 2-1
 - generic configurations, 2-5
 - naming conventions, 2-5
 - specific configurations, 2-6
 - commands
 - BACKUP, 2-10, 6-22
 - CATALOG, 10-6
 - CHANGE, 9-4
 - EXECUTE SCRIPT, 10-13
 - interactive use of, 1-3
 - RESYNC CATALOG, 10-20
 - standalone commands, 1-4
 - using command files, 1-3
 - compilation and execution of commands, 1-2
 - configuring
 - default device types, 2-4
 - device types, 2-3
 - corrupt datafile blocks, 2-43
 - handling I/O errors and, 2-41
 - crosschecking recovery catalog, 9-4
 - database connections
 - auxiliary database, 4-2
 - duplicate database, 4-3
 - hiding passwords, 4-4
 - with password files, 4-1
 - DBMS_PIPE package, 4-5
 - duplicate databases
 - how created, 13-2
 - environment
 - definition, 1-1
 - error codes
 - message numbers, 12-3
 - errors, 12-1, 12-2
 - interpreting, 12-5
 - file deletion
 - overview, 9-5
 - fractured block detection in, 2-44
 - hanging backups, 12-13
 - image copy backups, 2-9
 - incremental backups
 - cumulative, 2-26
 - differential, 2-25
 - level 0, 2-25
 - integrity checking, 2-42
 - interactive use of commands, 1-3
 - jobs
 - monitoring progress, 9-13
 - media management
 - backing up files, 1-8
 - Backup Solutions Program (BSP), 1-8
 - crosschecking, 9-4
 - media manager, linking with a, 5-4
 - metadata, 1-4, 10-1, 13-1
 - storing in control file, 1-5
 - monitoring, 9-10, 9-16
 - multiplexing
 - datafiles, 2-12
 - overview, 1-2
 - performance
 - monitoring, 9-10
 - pipe interface, 4-5
 - recovery
 - after total media failure, 7-10
 - recovery catalog, 1-6
 - availability, 10-22
 - backing up, 10-17
 - compatibility, 1-7
 - contents, 1-6
 - crosschecking, 9-4
 - managing the size of, 10-13

- moving to new database, 10-20
- operating with, 1-5
- operating without, 1-5
- recovering, 10-19
- registration of target databases, 1-6, 10-5
- resynchronizing, 10-9
- snapshot control file, 1-7
- synchronization, 1-7
- updating after schema changes, 10-11
- upgrading, 10-25
- reports, 9-1
 - overview, 9-2
- restoring
 - datafiles, 3-1
 - to new host, 7-2
- return codes, 12-7
- RPC calls and, 12-15
- snapshot control file location, 5-20
- standby databases
 - creating, 3-11
- starting, 4-1
- stored scripts, 1-3
- synchronous and asynchronous I/O, 11-3
- tablespace point-in-time recovery, 3-7
- tags for backups, 2-19
- terminating commands, 12-9
- test disk API, 5-5
- types of backups, 2-9
- using RMAN commands, 1-2
- recovery window
 - point of recoverability, 2-31
- recovery windows
 - backup optimization and, 2-38
 - definition, 2-31
- RECOVERY_PARALLELISM initialization
 - parameter, 18-29
- redo logs
 - incompatible format, 21-2
 - listing files for backup, 17-2
 - naming, 18-13
 - parallel redo, 21-2
- redo records
 - problems when applying, 21-2
- REGISTER command, 10-5
- REPORT OBSOLETE command, 2-33
- reports, 9-1
 - obsolete backups, 9-2
 - overview, 9-2
- repository
 - RMAN, 1-4
- RESET DATABASE command
 - INCARNATION option, 10-8
- RESETLOGS operation
 - when necessary, 18-25
- RESETLOGS option
 - of ALTER DATABASE, 18-24, 18-26, 18-28, 18-29
- restartable backups
 - definition, 2-39, 6-10
- restarting RMAN backups, 6-10
- RESTORE command, 3-1
 - FORCE option, 3-3
- restore optimization, 3-3
- restoring
 - archived redo logs, 18-5
 - backup control file
 - using SET DBID, 7-15
 - control files
 - to default location, 18-6
 - to nondefault location, 18-7
 - database
 - to default location, 18-27
 - to new host, 7-2
 - to new location, 18-28
 - database files, 3-1
 - how RMAN chooses, 3-2
 - mechanics, 3-1
 - restore optimization, 3-3
 - datafiles
 - to default location, 18-4
 - to raw devices, 18-4
 - user-managed backups, 18-2
 - keeping records, 17-20
- RESUME clause
 - ALTER SYSTEM statement, 17-14
- resuming recovery after interruption, 18-15
- RESYNC CATALOG command, 10-9
 - FROM CONTROLFILECOPY option, 10-20
- resynchronizing the recovery catalog, 10-9
- retention policies
 - affect on backup optimization, 2-37
 - definition, 2-30
 - disabling, 2-31
 - exempt backups, 2-34
 - recovery window, 2-31
 - redundancy, 2-31, 2-33
- return codes
 - RMAN, 12-7
- RMAN. *See* Recovery Manager

S

- sbtio.log
 - and RMAN, 12-2
- sbttest program, 12-7
- scenarios, Recovery Manager
 - backing up archived redo logs, 6-15
 - duplexing backup sets, 6-2
 - handling backup errors, 6-21
 - maintaining backups and copies, 6-19
 - NOARCHIVELOG backups, 6-18
 - recovering pre-resetlogs backup, 7-1
 - recovery after total media failure, 7-10
 - setting size of backup sets, 6-14
- schemas
 - changes
 - updating recovery catalog, 10-11
- SCN (system change number)
 - use in distributed recovery, 19-14
- server parameter files
 - autobackups, 2-28

- configuring autobackups, 2-28
- server sessions
 - Recovery Manager, 1-2
- session architecture
 - Recovery Manager, 1-2
- SET command
 - MAXCORRUPT option, 6-22
- SET statement
 - AUTORECOVERY option, 18-11
 - LOGSOURCE variable, 18-5, 18-14
- shared server
 - configuring for use with RMAN, 5-22
- short waits
 - definition of, 11-8
- SHOW command, 2-4
- SHUTDOWN statement
 - ABORT option, 18-6, 18-7, 18-20, 18-27, 18-28
- size of backup sets
 - setting, 2-21
- SKIP OFFLINE option
 - of BACKUP, 6-13
- SKIP READONLY option
 - of BACKUP, 6-13
- snapshot control files, 1-7
 - specifying location, 5-20
- split mirrors
 - using as backups, 6-3
- splitting mirrors
 - suspend/resume mode, 17-13
- standalone Recovery Manager commands, 1-4
- standby databases
 - creating using RMAN, 3-11
 - updating with incrementals, 13-24
- starting RMAN
 - without connecting to a database, 4-1
- STARTUP statement
 - MOUNT option, 18-22
- stored scripts
 - creating RMAN, 10-13
 - deleting, 10-16
 - managing, 10-13
 - Recovery Manager, 1-3
- stuck recovery
 - definition, 21-2
- SUSPEND clause
 - ALTER SYSTEM statement, 17-14
- suspending a database, 17-13
- suspend/resume mode, 17-13
- SWITCH command, 2-10
- synchronizing duplicate database
 - using DUPLICATE DATABASE, 13-23
 - using incremental backups, 13-24
- system time
 - changing
 - effect on recovery, 18-20

T

- tables
 - recovery of dropped, 19-12

- tablespace backups
 - using RMAN, 6-5, 6-11
- tablespace point-in-time recovery, 8-3
 - clone database, 20-2
 - introduction, 20-1
 - methods, 20-2
 - performing
 - user-managed, 20-1
 - planning for, 20-3
 - procedures for using transportable tablespace
 - feature, 20-11
 - requirements, 20-3
 - terminology, 20-2
 - transportable tablespace method, 20-2
 - user-managed, 20-2
 - using RMAN, 3-7
 - basic steps, 8-3
 - introduction, 8-1
 - planning, 8-5
 - preparing the auxiliary instance, 8-18
 - restrictions, 8-5
 - why perform, 8-3
- tablespace transport
 - online, 14-1
- tablespaces
 - backups, 17-6
 - offline, 17-4
 - online, 17-6
 - excluding from RMAN backups, 5-18
 - read-only
 - backing up, 6-13, 17-9
 - read/write
 - backing up, 17-5
 - recovering offline in open database, 18-18
- tags, 2-19
- terminating RMAN commands, 12-9
- test databases, creating, 3-9
- test disk API, 5-5
- testing RMAN
 - backups, 2-44, 6-11
 - with media management API, 12-7
- time format
 - RECOVER DATABASE UNTIL TIME
 - statement, 18-24
- time-based recovery, 18-23
 - coordinated in distributed databases, 19-13
- trace files
 - and RMAN, 12-2
 - backing up control file, 17-11
 - control file backups to, 17-11
- transparent backup encryption, 6-7
- transportable tablespace
 - cross-platform, 15-1
 - online, 14-1
- transportable tablespaces
 - and CONVERT DATAFILE/TABLESPACE, 15-1
 - creating with RMAN, 14-1
 - and Data Pump Export, 14-9
 - and past points in time, 14-8
 - auxiliary destination, 14-4, 14-7

- auxiliary instance parameter file, 14-10, 14-11
- Concepts, 14-2
- file locations, 14-12
- initialization parameters, 14-10
- limitations, 14-6
- Shared Pool Size, 14-12, 14-14
- when to use, 14-2
- cross-platform, 15-1
- recovery, 19-6
- TSPITR and, 20-2
- transporting databases across platforms, 15-7
- trial recovery
 - explanation, 21-7
 - overview, 21-6
- TSPITR, 8-3
- TSPITR. *See* tablespace point-in-time recovery
- tuning
 - Recovery Manager
 - V\$ views, 9-10

U

- UNAVAILABLE option
 - of CHANGE, 9-9
- unrecoverable objects
 - and RECOVER operation, 19-4
 - recovery
 - unrecoverable objects and, 19-4
- unregistering a database from the recovery
 - catalog, 10-7
- UNTIL TIME option
 - RECOVER command, 18-24
- upgrading the recovery catalog, 10-25
- user errors
 - recovery from, 19-12
- user-managed backups, 17-1, 17-3
 - backup mode, 17-7
 - control files, 17-10
 - binary, 17-10
 - trace files, 17-11
 - determining datafile status, 17-2
 - hot backups, 17-8
 - listing files before, 17-1
 - offline datafile, 17-4
 - offline tablespaces, 17-4
 - read-only tablespaces, 17-9
 - restoring, 18-4
 - tablespace, 17-6
 - verifying, 17-18
 - whole database, 17-3
- user-managed recovery, 18-20
 - ADD DATAFILE operation, 19-2
 - complete, 18-15
 - incomplete, 18-20
 - interrupting, 18-15
 - opening database after, 18-24
 - scenarios, 19-1
- user-managed restore operations, 18-2
- USING BACKUP CONTROLFILE option
 - RECOVER command, 18-23

V

- V\$ARCHIVED_LOG view, 3-6
 - listing all archived logs, 17-12
- V\$BACKUP view, 17-2
- V\$BACKUP_ASYNC_IO, 9-11
- V\$BACKUP_CORRUPTION view, 2-42
- V\$BACKUP_SYNC_IO, 9-10
- V\$COPY_CORRUPTION view, 2-42
- V\$DATABASE_BLOCK_CORRUPTION view, 6-12, 7-14
- V\$DATAFILE view, 17-2
 - listing files for backups, 17-1
- V\$LOG_HISTORY view
 - listing all archived logs, 18-5
- V\$LOGFILE view, 19-7, 19-8
 - listing files for backups, 17-2
 - listing online redo logs, 17-2
- V\$PROCESS view, 9-10
- V\$RECOVER_FILE view, 18-3
- V\$RECOVERY_LOG view
 - listing logs needed for recovery, 18-5
- V\$SESSION view, 9-10
- V\$SESSION_LONGOPS view, 9-10
- V\$SESSION_WAIT view, 9-10
- V\$TABLESPACE view, 17-2
- validating
 - backups, 6-11
- views
 - recovery catalog, 10-22

W

- wallet, 6-7
- whole database backups
 - ARCHIVELOG mode, 17-3
 - inconsistent, 17-3
 - NOARCHIVELOG mode, 17-3
 - preparing for, 17-3