



Professional Expertise Distilled

# Oracle Database 12c Backup and Recovery Survival Guide

A comprehensive guide for every DBA to learn recovery and backup solutions

Francisco Munoz Alvarez  
Aman Sharma

**[PACKT]** enterprise   
PUBLISHING professional expertise distilled

# Table of Contents

[Oracle Database 12c Backup and Recovery Survival Guide](#)

[Credits](#)

[About the Author](#)

[Acknowledgement](#)

[About the Author](#)

[Acknowledgement](#)

[About the Reviewers](#)

[www.PacktPub.com](#)

[Support files, eBooks, discount offers and more](#)

[Why Subscribe?](#)

[Free Access for Packt account holders](#)

[Instant Updates on New Packt Books](#)

[Preface](#)

[What this book covers](#)

[What you need for this book](#)

[Who this book is for](#)

[Conventions](#)

[Reader feedback](#)

[Customer support](#)

[Downloading the example code](#)

[Errata](#)

[Piracy](#)

[Questions](#)

[1. Understanding the Basics of Backup and Recovery](#)

[Purpose of backup and recovery](#)

[Testing backups](#)

[Protecting data](#)

[Media failure](#)

[Hardware failure](#)

[Human error](#)

[Application error](#)

[Types of backup](#)

[A physical backup](#)

[A logical backup](#)

[Backup strategies](#)

[Restore versus recovery](#)

[What is redo?](#)

[Redo generation and recoverability](#)

[The NOARCHIVELOG mode](#)

[The ARCHIVELOG mode](#)

[Understanding the ARCHIVELOG mode](#)

[Preparing for the ARCHIVELOG mode](#)

- [Checking the status of the ARCHIVELOG mode](#)
- [Specifying parameters](#)
- [Viewing the status of archival destinations](#)
- [Placing a database into the ARCHIVELOG mode](#)
- [Differences between redo and undo](#)
- [Facing excessive redo generation during an online backup?](#)

[Summary](#)

## [2. NOLOGGING Operations](#)

[LOGGING versus NOLOGGING](#)

[Disabling redo generation \(NOLOGGING\)](#)

[NOLOGGING operations](#)

[Indexed organized tables – an exception](#)

[Reducing redo generation](#)

[Tips when LOGGING is in effect \(not using NOLOGGING\)](#)

[Backups](#)

[Bulk inserts](#)

[Bulk deletes](#)

[Bulk updates](#)

[Partitioning](#)

[Tips for developers](#)

[Tips when NOLOGGING is in effect](#)

[Partitioning](#)

[Direct path inserts](#)

[Bulk inserts](#)

[Bulk deletes](#)

[Bulk updates](#)

[Backups and NOLOGGING](#)

[Redo-related wait events](#)

[The 'log file parallel write' event](#)

[The 'log file sync' event](#)

[The 'redo log space request' event](#)

[The 'log buffer space' event](#)

[Block corruption due to NOLOGGING](#)

[Repairing NOLOGGING changes on physical and logical standby databases](#)

[Finding sessions that generate lots of redo](#)

[Some other important facts](#)

[Redo and undo for DML](#)

[Redo and temporary tables](#)

[Redo generation and materialized views](#)

[Flashback and NOLOGGING](#)

[Performance and recovery considerations](#)

[Direct path load using SQL\\*Loader](#)

[Some useful scripts](#)

[Redo generated since instance startup](#)

[Redo generated since session startup](#)

[Redo generated by current user sessions](#)

[Current status for redo logs](#)  
[Redo log group and log switch information](#)  
[NOLOGGING objects in the database](#)

[Summary](#)

### [3. What is New in 12c](#)

[Pluggable database](#)

[RMAN new features and enhancements](#)

[Container and pluggable database backup and restore](#)

[Enterprise Manager Database Express](#)

[Backup privileges](#)

[SQL and DESCRIBE](#)

[Multi-section backups for incremental backups](#)

[Network-based recovery](#)

[Active Duplicate](#)

[Support for the third-party snapshot](#)

[Cross-platform data transport](#)

[Table recovery](#)

[Data Pump's new features and enhancements](#)

[Disabling LOGGING on Data Pump Import](#)

[Full transportable Export/Import](#)

[Exporting views as tables](#)

[Extended character data types](#)

[Encryption password](#)

[Compressing tables on Import](#)

[Exporting data from the data vault](#)

[Creating SecureFile LOBs on Import](#)

[Auditing Data Pump commands](#)

[Summary](#)

### [4. User-managed Backup and Recovery](#)

[Cold backup](#)

[Offline backup](#)

[Hot backups](#)

[Hot backup of a whole database](#)

[Hot backup of tablespaces](#)

[Hot backup of a container database](#)

[Whole container database](#)

[Root only or individual pluggable database](#)

[Check datafile status](#)

[Control file backup](#)

[Binary backup](#)

[Text file backup](#)

[Flashback database](#)

[Recovering from a user-managed backup](#)

[Other recovery scenarios](#)

[Losing all copies of the current control file](#)

[Losing one copy of a multiplexed control file](#)

[Loss of archived redo logs or online redo logs](#)

[Loss of SPFILE](#)

[Summary](#)

## [5. Understanding RMAN and Simple Backups](#)

[Why RMAN?](#)

[Getting started with RMAN](#)

[RMAN architecture](#)

[Target database](#)

[Target Control File](#)

[RMAN channels](#)

[Auxiliary database](#)

[Recovery Catalog](#)

[RMAN client](#)

[Oracle Enterprise Manager \(EM\) Cloud Control 12c](#)

[Media Management Library \(MML\)](#)

[Oracle secure backup](#)

[Memory requirements for RMAN](#)

[Configuring ARCHIVELOG mode and FRA](#)

[Introducing RMAN backup formats](#)

[Backup sets](#)

[Image copy](#)

[Full backup](#)

[Incremental backups](#)

[Differential incremental backup](#)

[Cumulative incremental backup](#)

[Power of one](#)

[Getting introduced to RMAN commands](#)

[Taking your first RMAN backup](#)

[File sections for backups of large data files](#)

[Using RMAN for performing incremental backups](#)

[Fast incremental backups using the block change tracking \(BCT\) file](#)

[Multi-section incremental backups](#)

[Incrementally updated backups](#)

[Performing backups of the control file, the SPFILE, and archived redo logs](#)

[Using RMAN compression for backups](#)

[RMAN for multitenant container databases](#)

[Performing a full backup of a multitenant CDB](#)

[Partial backup of a multitenant CDB](#)

[Performing backup of a pluggable database and root](#)

[Performing backups of backups](#)

[Restarting RMAN backups](#)

[Useful RMAN views related to backups](#)

[Summary](#)

## [6. Configuring and Recovering with RMAN](#)

[RMAN configuration – an introduction](#)

[Using the V\\$RMAN CONFIGURATION view](#)

- [Using the SHOW ALL command](#)
- [Configuring the backup retention policy](#)
  - [Redundancy retention policy](#)
  - [Recovery window retention policy](#)
- [Using backup optimization](#)
- [Configuring the device types for the backup](#)
- [Configuring auto backup for the control file and SPFILE](#)
- [Configuring RMAN channels](#)
- [Creating duplexed backups](#)
- [Configuring encrypted backups](#)
  - [Transparent encryption of backups](#)
  - [Creating and using Oracle Software Keystore](#)
  - [Password encryption](#)
  - [Dual-mode encryption](#)
- [Configuring compression for backups](#)
- [Configuring the snapshot control file](#)
- [Configuring the archived log deletion policy](#)
- [Configuring the FRA](#)
- [Configuring authentication for RMAN](#)
  - [Operating system authentication](#)
  - [Authenticating the password file](#)
- [Crash and media recovery](#)
- [Key terms related to recovery](#)
  - [Overview of stages in the database startup](#)
  - [Steps involved in the crash/instance recovery](#)
  - [Instance recovery in container and pluggable databases](#)
  - [Performing media recovery](#)
  - [Recovery in the NOARCHIVELOG mode](#)
  - [Loss of a temporary file](#)
  - [Loss of non-system data files](#)
  - [Loss of system data files](#)
  - [Recovering whole CDBs, PDBs, and root container databases](#)
    - [Complete recovery of a container database](#)
    - [Complete recovery of a pluggable database](#)
    - [Complete recovery of the root container](#)
  - [Performing control file recovery](#)
  - [Performing Block Media Recovery \(BMR\)](#)
  - [Performing point-in-time recovery](#)
    - [Database Point-in-time Recovery \(DBPITR\)](#)
    - [Tablespace Point-in-time Recovery \(TSPITR\)](#)
  - [Table and table partition-level recovery from RMAN backups](#)
  - [Data recovery advisor](#)
- [Summary](#)
- [7. RMAN Reporting and Catalog Management](#)
  - [Using the control file for RMAN metadata](#)
  - [Using the CONTROLFILE\\_RECORD\\_KEEP\\_TIME parameter](#)

## What is a recovery catalog

### Creating the recovery catalog

Sizing and creating a database for the recovery catalog

Creating the default tablespace for the catalog owner

Creating the catalog owner schema

Granting privileges to the catalog owner

Creating the recovery catalog

Using the recovery catalog

Resynchronizing the recovery catalog with the control file

Merging multiple recovery catalogs into one

Using virtual private catalogs

Creating and managing stored scripts

Making a recovery catalog highly available

Upgrading the recovery catalog

Unregistering databases from the recovery catalog

Dropping a recovery catalog

Views related to the recovery catalog

## Reporting in RMAN

Using the LIST command

Using the REPORT command

## Summary

## 8. RMAN Troubleshooting and Tuning

### Getting started with RMAN troubleshooting

Using CHECKSYNTAX

Reading the RMAN error stack

Debugging RMAN using the DEBUG clause

Using the alert log and operating system trace files

### RMAN tuning – an introduction

I/O and RMAN – two sides of one coin

Number and size of the input and output buffers

Synchronous and asynchronous I/O modes

Setting the Large Pool memory

Monitoring RMAN I/O performance using dictionary views

V\$BACKUP\_ASYNC\_IO (for asynchronous I/O)

V\$BACKUP\_SYNC\_IO (for synchronous I/O)

Tuning SBT (tape) performance

Monitoring RMAN sessions and operations

Stopping RMAN from being uncontrollable

Using incremental, multi-section, multiplexing, and parallelism

Troubleshooting RMAN performance using tracing

## Summary

## 9. Understanding Data Pump

What is Data Pump?

The Data Pump architecture

New concepts with Data Pump

Methods to move the data

[Datafile copy](#)

[Direct path](#)

[External tables](#)

[Conventional path](#)

[Network link](#)

[Data Pump files](#)

[Roles for Data Pump export and import](#)

[Directory objects](#)

[Creating directory objects](#)

[Data Pump scenarios](#)

[Schema export and import](#)

[Exporting and importing tables](#)

[Exporting and importing a whole database/pluggable database](#)

[Using Export to estimate space](#)

[Parallel full database export and interactive-command mode](#)

[Importing tables with only metadata](#)

[Exporting views as tables](#)

[Importing data via a network link](#)

[Summary](#)

## [10. Advanced Data Pump](#)

[Data masking](#)

[Metadata repository and version control](#)

[Using SOURCE\\_EDITION and TARGET\\_EDITIONS](#)

[Cloning a user](#)

[Creating smaller copies of production](#)

[Creating your database in a different file structure](#)

[Moving all objects from one tablespace to another](#)

[Moving an object to a different schema](#)

[Migrating data for upgrade](#)

[Downgrading an Oracle Database](#)

[Transporting a tablespace](#)

[Data Pump flashback](#)

[Monitoring Data Pump job status](#)

[Some performance tuning tips](#)

[Summary](#)

## [11. OEM12c and SQL Developer](#)

[Configuring backup, recovery, and catalog settings](#)

[Backup settings](#)

[Recovery settings](#)

[Catalog settings](#)

[Scheduling an RMAN backup](#)

[Using the Oracle-Suggested Backup strategy option](#)

[Using the Customized Backup option](#)

[Restore points](#)

[Export/Import with OEM12c](#)

[Executing an export operation](#)

[Monitoring the job](#)

[SQL developer 3.2](#)

[RMAN operations](#)

[Data Pump operations](#)

[Summary](#)

## [A. Scenarios and Examples – A Hands-on Lab](#)

[Configuring the database](#)

[Making sure an spfile is used to start the database](#)

[Placing the database in the archivelog mode and activating Flashback](#)

[Creating a new redo log group and associated files](#)

[Configuring RMAN](#)

[Creating the target DB RMAN backup account](#)

[Configure RMAN using the configure command](#)

[Backup database](#)

[Checking and deleting obsolete backups](#)

[Creating RMAN catalog user](#)

[Creating recovery catalog](#)

[Register your DB in the recovery catalog](#)

[Creating a virtual private catalog](#)

[Enabling Block Change tracking](#)

[Playing with RMAN, FRA, and catalog views](#)

[Monitoring a backup](#)

[Incremental backups](#)

[Multisection backups](#)

[FRA – checking number of redo switches](#)

[Check for alerts](#)

[Check FRA usage](#)

[See the archived log generated by the DB target](#)

[See the control file backups](#)

[See the corrupted list that exists in datafile backups](#)

[See block corruption in the DB, populated when backup or backup validate](#)

[See all RMAN configurations \(equivalent to show all\)](#)

[Monitor backup outputs \(RMAN\)](#)

[Offline backups with RMAN](#)

[Offline backup without using configured defaults](#)

[Using backup limits \(duration\)](#)

[Modifying the retention policy for a backup set \(archival backups\)](#)

[Archive deletion policy](#)

[Using RMAN to scan DB for physical and logical errors](#)

[Configuring tablespaces for exclusion from whole database backups](#)

[Skipping offline, inaccessible, or read-only datafiles](#)

[Forcing backups of read-only datafiles](#)

[Backup of newly added datafiles](#)

[Backup files not backed up during a specific period](#)

[General backup examples](#)

[Backup copies](#)

## Advanced RMAN

Information about fully-completed backups

Summary of the active session history

How long does it take?

V\$BACKUP ASYNC IO

Tablespace Point-in-time Recovery (TSPITR)

Reporting from a catalog

Duplex backup

Check if the database is recoverable

Recover advisor

## Magic with Data Pump

Preparing Data Pump

Data masking

Metadata repository

Cloning a user

Creating smaller copies of production

Creating your database in a different structure

Time-based flashback

## Backup and recovery scenarios

Active duplication of a database to a different server with the same structure

(non-OMF and non-ASM)

Duplicating a PDB

ASM backup and restore

Recovering from the loss of the SYSTEM tablespace

Recovering a lost datafile using an image from an FRA

## Index

# Oracle Database 12c Backup and Recovery Survival Guide

---

# Oracle Database 12c Backup and Recovery Survival Guide

Copyright © 2013 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: September 2013

Production Reference: 1170913

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham B3 2PB, UK.

ISBN 978-1-78217-120-1

[www.packtpub.com](http://www.packtpub.com)

Cover Image by Dilleshwar (<[dil.snowfire@gmail.com](mailto:dil.snowfire@gmail.com)>)

# Credits

## **Authors**

Francisco Munoz Alvarez

Aman Sharma

## **Reviewers**

Gokhan Atil

Wissem EL Khlifi

Bjorn Naessens

Arup Nanda

Alessandro Parisi

Laurent Schneider

## **Acquisition Editor**

Vinay Argekar

## **Lead Technical Editor**

Azharuddin Sheikh

## **Technical Editors**

Manan Badani

Aparna Chand

Vivek Pillai

## **Project Coordinator**

Wendell Palmer

## **Proofreader**

Ameesha Green

## **Indexers**

Rekha Nair

Tejal Soni

**Graphics**

Abhinash Sahu

Ronak Dhruv

**Production Coordinator**

Shantanu Zagade

**Cover Work**

Shantanu Zagade

# About the Author

**Francisco Munoz Alvarez** has over two decades of experience in consulting, analysis, support, implementation, and migration of Oracle products. He is also an expert in most phases of a database life cycle, for example, development, stabilization, security, backup and recovery, tuning, installations, and data warehouse (ETL) with excellent implementation and support methodologies. He is a popular speaker at many Oracle conferences around the world.

He is also the President of CLOUG (Chilean Oracle Users Group), LAOUC (Latin American Oracle Users Group Community, which is the umbrella organization for all of Latin America), and NZOUG (New Zealand Oracle Users Group). He worked as an Associate Technologist for Oracle in Brazil and Chile. He was an Oracle instructor for the New Horizons Centre (Chile) and for Oracle Education (Brazil and Chile). He also worked in the first team to introduce Oracle to South America (Oracle 6 and the beta version of Oracle 7). He was also the first Master Oracle 7 Database Administrator in South America, as well as the first Latin American Oracle professional to be awarded a double ACE (ACE in 2008 and ACE Director in 2009) by Oracle HQ. In 2010, he had the privilege to receive a prestigious Oracle Magazine Editor's Choice Award as the Oracle Evangelist of the Year—a huge recognition for his outstanding achievements in the Oracle world that includes the creation and organization of the already famous OTN Tours that are the biggest Oracle evangelist events in the world.

Currently, Francisco works for Revera Limited, which is a leading provider of utility computing infrastructure and enterprise data management in New Zealand, as the Oracle Professional Services Manager since June 2011. He also maintains an Oracle blog (<http://www.oraclenz.org>) and you can always contact him through this or Twitter (@fcomunoz) regarding any questions about this book.

# Acknowledgement

Writing a book is not the accomplishment of one or two people, it involves a huge group of people behind the scenes to make it possible. Due to this, I would like to take this opportunity to acknowledge some important people in my life that without their support and help, writing this book would have been an impossible journey. I would like to start with my wonderful wife, mother of my four wonderful children; honey, without your enormous patience and support, nothing could be possible—even though at times you don't understand this fascinating journey and this fascinating world, my love, I dedicate this book to you! Secondly, my mother, my father and my well-remembered grandmother for everything you all taught me in my early years and the things I have seen because of you. And finally my good friends, Thomas Kyte, Arup Nanda, Aman Sharma (your support, help, and friendship is invaluable), Gokhan Atil, Wissem EL Khifi, Bjorn Naessens, Alessandro Parisi, Laurent Schneider, all editors, and everyone involved in the book, without your help and support throughout writing this book and even during my career, I would never have been able to finish writing this book.

# About the Author

**Aman Sharma** is an Oracle Database consultant and instructor. He holds a Master's degree in Computer Applications and has been working with Oracle Database for over a decade. His main focus is to understand how Oracle Database works internally. Besides the core database, he has a very strong knowledge of Linux, Solaris, Oracle RAC, Data Guard, RMAN, Oracle Exadata and Oracle Enterprise Manager.

He is an Oracle Certified Professional and also an Oracle Certified Expert in Oracle RAC, SQL and Linux. He is also an Oracle Certified Implementation Specialist and a certified Solaris System Administrator. He loves to share his learning with the Oracle community, so when he is not delivering an Oracle related session, he spends a lot of his time participating in OTN (Oracle Technology Network) discussion forums. He also maintains an Oracle blog (<http://blog.aristadba.com>), is also on Twitter (@amansharma81), you can reach to him using either ways. He is a strong supporter of user groups and is a member of India Oracle Users Group (<http://www.aioug.org>) and has been a speaker at various Oracle conferences organized by AIOUG. He is also an associate member of IOUG (<http://www.ioug.org>). In 2010, he was awarded the prestigious Oracle ACE award from Oracle Corporation.

# Acknowledgement

I dedicate this book to my entire family for their so much love, affection, and constant encouragement.

A child is just like a plain canvas. It's his parents who fill colors into it with their love, blessings, guidance, and teachings. The same is true for me as well. This book is a result of the blessings, love, affection, guidance, and support of my mom and dad. Thank you so very much mom and dad! Dad, how so I wish that you were here today, but I know you are watching over me from the heavens, your blessings are always with me and you must be very happy today! We all miss you so very much—every day, in every moment!

Writing a book is a team effort and requires a lot of hard work but every journey starts from that very first step. For me, getting aboard to work on this book became possible only because of my dear friend and co-author Francisco Munoz Alvarez, who invited me to become a co-author with him. It was a really great team effort and working with him was an absolute delight, as well as a very learning and rewarding experience. I thank him for not only making me a part of this project but also for his support, help and mentoring the entire time while we were working on this project.

Though a book is written by its authors, there are many other people who work behind the scenes and do tons of things that finally result in letting a book see the sunshine. The most important among these are the technical reviewers who do the painful job of reading the raw writing of an author, line by line and correct it. It's due to their corrections, suggestions, feedback, and questions that convert those sentences written in a Word document into the chapters of an actual book. I am so very thankful to our most excellent team of the technical reviewers which we got: Arup Nanda, Wissem EL Khlifi, Gokhan Atil, Bjorn Naessens, Laurent Schneider, and Alessandro Parisi. Without their input, feedback and corrections, it wouldn't have been possible to find out what was missing, what needed to be removed, which topics required more explanation, and above all, to spot the technical errors and correct them. So when you would be reading about a topic in any chapter and would like it, remember that a major contribution to it has come in some way or the other from the technical excellence and input given by the reviewers.

I also want to say a big thanks to the entire team at Packt who did a splendid job in every aspect of the book, especially to my editor Vinay Argekar, and Wendell Palmer, the coordinator for the entire project, for being so supportive in this entire journey. Thanks to our Lead Technical Editor Azharuddin Sheikh along with his entire team: Manan Badani, Aparna Chand, and Vivek Pillai. They all have spent a lot of time ironing out the final wrinkles from the chapters, giving them that marvellous polished look which you will see throughout. I also want to pass on my gratitude to all those people with whom I might have not interacted directly but who worked very hard behind the scenes to bring this book into your hands.

I would like to take this opportunity to say thanks to all those people in the Oracle community who were not involved in this project but who are always an inspiration and constantly motivated me—not only to correct and improve myself by learning as much as possible, but also to share that learning with others and help them. The list is far too long and I am going to miss more than I will be able to put here but still, to name a few (and in no particular order): Jonathan Lewis, Hans Forbrich, Dan Morgan, Hemant Chitale, Sam Alapati, Eddie Awad, Uwe Hesse, Syed Jaffar Hussain, Kellyn Pot'Vin, Tim Gorman, Steve Adams, Kyle Hailey, Tim Hall, Mark Bobak, Cary Millsap, Howard Rodgers, Steve Karam, Tanel Poder, Alex Gorbachev, Kai Yu, Riaz Shamsudeen, Murali Vallath, Gaja Krishna Vaidyanatha, Rama Velpuri, and so many more as the list goes on and on and on. Many thanks to all of you!

Thanks to my friend Kamran Agayev for giving me the very first writing opportunity by inviting me to be his co-author and Don Burelson, and entire team of Rampant Press for their support in the entire time of writing it. Thanks to Zahida Praveen who gave me the first opportunity to step in and prove myself in the real world. I also wish to express my gratitude and thanks to all my superordinates for their constant support and motivation: Arvind Dikshit, Gautham Malareddy, Usha George, Joann Too, Ajai Sahni, Shantnu Tandan, Latha Gangadharan, Kalyan Chakravarthy, and Ritesh Jain. I would like to offer my gratitude to all the delegates that I get to meet in person and to those whom I interact with virtually over the OTN forums, for asking their questions and doubts. It's your questions which make me learn (and many times, even re-learn) new things every single day.

Finally, thanks a ton to *you* for making this book a part of your library. I really hope and wish that you are going to like the book and it's going to be useful for you in some way.

# About the Reviewers

**Gokhan Atil** is a Senior Applications DBA working at one of the leading IT consultancy companies in Turkey. He has a strong background in UNIX systems and hands-on experience in Oracle Database versions 11g/10g/9i/8i and Oracle E-Business Suite R12/11i. He is an active member of the Oracle community and a founding member of the Turkish Oracle User Group (TROUG). He has written and presented papers at various conferences. He's one of the co-authors of the book *Expert Oracle Enterprise Manager 12c*, published in 2013.

Gokhan Atil holds various Oracle certifications such as Oracle Database 10g/11g Administrator Certified Professional, Oracle E-Business Suite R12 Applications DBA Certified Professional, and Oracle Exadata Database Machine X3 Administrator Certified Expert. He was honored with the Oracle ACE award in 2011. He has a blog where he shares his experiences with Oracle since 2008. You can get in touch with Gokhan at <http://www.gokhanatil.com>.

**Wissem El Khlifi** is the first Oracle ACE in Spain and an Oracle Certified Professional DBA with over 12 years of IT experience.

He earned a Computer Science Engineer's degree from FST, Tunisia, and a Master's in Computer Science from the UPC, Barcelona.

His areas of interest are Linux System Administration, Oracle Database including high availability (RAC, Dataguard), and Oracle applications.

He has worked with Oracle Database since version 7.

His career has included the roles of Oracle programmer, Java analyst/programmer, Oracle DBA, architect, project manager, and team leader. He currently works as the Senior Oracle DBA for Schneider Electric/APC.

He writes numerous articles on his Oracle website: <http://www.oracle-class.com>.

**Bjorn Naessens** is a Senior Oracle DBA at Uptime Technologies N.V. based in Belgium. Prior to his work with Uptime, he was a system engineer/DBA at a leading Belgian media company. Most of Bjorn's projects are for new infrastructure setups, hence using RMAN for backup/recovery has been constant in these projects.

Next to Oracle databases, Bjorn is specialized in Oracle VM and Oracle Linux and has a passion for the Oracle Database Appliance.

**Arup Nanda** has been an Oracle DBA for the last 19 years and counting. He works as the Principal Database Architect at a New York multinational company, has co-authored four books, has written around 300 articles, and presented at 150 different conferences around the world. Honoring professional accomplishments and