



# Oracle Data Guard 11g Handbook

ORACLE® 11g  
DATABASE

Undocumented Best Practices and Real-World Techniques

Larry Carpenter

Joe Meeks  
Charles Kim  
Bill Burke

Sonya Carothers  
Joydip Kundu  
Michael Smith  
Nitin Vengurlekar

Foreword by Tom Kyte of [asktom.oracle.com](http://asktom.oracle.com)



**ORACLE®**

*Oracle Press™*

# **Oracle Data Guard 11g Handbook**

## About the Authors

**Larry Carpenter** is a Distinguished Product Manager at Oracle USA and is a member of the Maximum Availability Architecture Product Management team in Server Technologies with a focus on Oracle's High Availability and Disaster Recovery technologies. Larry has 35 years of experience in the computer industry, with the last 20 years focused on the business continuity needs of critical databases and applications. He is recognized by the Oracle user community as a Data Guard expert, an HA Technical Evangelist, and a consultant to diverse Enterprise customers worldwide. Larry's expertise is ensuring the successful deployment of Oracle Disaster Recovery Solutions in diverse computing environments and bringing constantly evolving customer requirements to Oracle's development teams. Larry is conversant in English, Italian, French, and German.

**Joe Meeks** is a Director of Product Management with Oracle's Database High Availability Group in Server Technologies. Joe manages customer programs that focus on data protection and high availability solutions using Oracle Data Guard and the Oracle Maximum Availability Architecture. These programs ensure customer success through knowledge transfer of HA best practices while closely aligning future Oracle development priorities with customer requirements. Joe has 30 years of experience in the computer industry helping customers to address HA requirements of business critical applications in manufacturing, retail, finance, energy, telecommunications, healthcare, and public sectors. He has a BS in Environmental Science and an MBA.

**Charles Kim** is an Oracle ACE and an Oracle Certified DBA. Charles works predominately in the Maximum Availability Architecture (MAA) space (RAC, ASM, Data Guard, and other HA solutions). Charles released his first book, *Oracle Database 11g New Features for DBA and Developers*, in November 2007. Charles also co-authored *Linux Recipes for Oracle DBAs* with APress, published in November 2008. Charles is also the author of the MAA case study at Oracle's web site ([www.oracle.com/technology/deploy/availability/htdocs/FNF\\_CaseStudy.html](http://www.oracle.com/technology/deploy/availability/htdocs/FNF_CaseStudy.html)). He holds certifications in Oracle, Red Hat Linux, and Microsoft; has more than 18 years of IT experience; and has worked with Oracle since 1991. Charles blogs regularly at <http://blog.dbaexpert.com> and provides technical solutions to Oracle DBAs and developers.

**Bill Burke** is a Consulting Technical Director with Oracle's System Performance and Architecture consulting practice. More than half of his 25 years in the IT industry has been committed to volunteer leadership roles. He has served on the board of directors of the International Oracle Users Group, International Oracle Users Council, Oracle Development Tools User Group, has led the first and second IOUG/Oracle Database 10g beta test teams, and has been an active participant on the public boards, forums, and Oracle mailing lists where he was known as the "Kinder and Gentler DBA." Most of his work today in the SP&A Practice is in best practice audits and the implementation and performance tuning of Maximum Availability Architectures including Real Application Clusters (RAC), Data Guard, and their management with Enterprise Manager Grid Control. Bill has been an OCP-certified DBA since version 7 of Oracle.

Mr. Burke is a Certified Flight Instructor—Instrument and has logged hundreds of hours as a commercial pilot and flight instructor over the years. In his free time away from Oracle, he is an accomplished professional photographer who works with local youth sports organizations, non-profit organizations on a pro-bono basis, and specializes in scenic, wilderness, and travel photography with an emphasis on endangered species.

You can reach him at [wburkejr@gmail.com](mailto:wburkejr@gmail.com).

**Sonya Carothers** is a Senior Oracle Database Administrator at PDX, Inc. She has more than 24 years of IT experience in database administration and software development. Sonya has worked as a senior database administrator, IT manager, and technical consultant. She has worked with several relational databases and has been working with Oracle since 1994. In addition, she has worked on a wide variety of projects in multi-platform environments. Her expertise includes high availability architecture, disaster recovery infrastructure, high performance database design, best practice database administration, and systems configuration.

**Joydip Kundu** is currently the Director of Development for Data Guard Logical Standby and LogMiner. He has been with Oracle since 1996 and is one of the original developers of Oracle LogMiner. Joydip is the architect of the log mining engine inside the Oracle RDBMS that underpins Data Guard Logical Standby, Streams Capture, and other redo-based features such as asynchronous Change Data Capture and Audit Vault. Joydip holds a Ph.D. in Computer Science from University of Massachusetts at Amherst.

**Michael Smith** is Principal Member of the technical staff in Oracle's Maximum Availability Architecture (MAA) team in Server Technologies. Mike has been with Oracle for 10 years, previously serving as the Data Guard Global Technical Lead within Oracle Global Support. Mike's current focus is developing, validating, and publishing HA best practices using Data Guard in an integrated fashion across all Oracle Database high availability features. His Data Guard technical specialties focus on network transport, recovery, role transitions, Active Data Guard, and client failover. He has published a dozen MAA Best Practice papers for Oracle 9i, 10g, and 11g. He has been a contributing author to other Oracle Press publications. Mike has also been speaker at the previous three Oracle Open World events held in San Francisco. His "What They Didn't Print in the DOC" best practice presentations covering Data Guard and MAA are a favorite among Oracle users, with attendance at the top of all Oracle Database technology presentations.

**Nitin Vengurlekar**, a consulting member of the technical staff at Oracle, is the author of *Oracle Automatic Storage Management* by Oracle Press. With more than 22 years of IT experience, including OS390 Systems Programming, UNIX Storage Administration, System and Database Administration, Nitin is a seasoned systems architect who has successfully assisted numerous customers to deploy highly available Oracle systems. He has worked for Oracle for more than 14 years, currently in the Real Application Clusters (RAC) engineering group, with specific emphasis on ASM and storage. He has written many papers on ASM usage and deployments on various storage array architectures and serves as a writer of and contributor to Oracle documentation as well as Oracle education material.

## About the Technical Editors

**Michael Powell** is an OCP-certified DBA with more than 15 years of IT experience. He has more than 12 years of experience in implementing and administering Oracle for Fortune 500 companies. Michael has worked as lead DBA for RAC and Data Guard implementations. He is also a contributor to the “Maximum Availability Architecture Implementation Case Study for Fidelity National Financial (FNF)” and has been a participant in Oracle Database Beta programs. Michael specializes in database and Oracle Application implementations. Here’s a link to a case study: [www.oracle.com/technology/deploy/availability/htdocs/FNF\\_CaseStudy.html](http://www.oracle.com/technology/deploy/availability/htdocs/FNF_CaseStudy.html).

**Sreekanth Chintala** is an OCP-certified DBA, has been using Oracle technologies for more than a decade, and has more than 15 years of IT experience. Sreekanth specializes in Oracle high availability, disaster recovery, and grid computing. Sreekanth is an author of many technical white papers and a frequent speaker at Oracle OpenWorld, IOUG, and local user group meetings. Sreekanth is active in the Oracle community and is the current web seminar chair for the community-run Oracle Real Application Clusters Special Interest Group ([www.ORACLERACSIG.org](http://www.ORACLERACSIG.org)).

**ORACLE®**

*Oracle Press™*

# Oracle Data Guard 11g Handbook

Larry Carpenter    Joe Meeks  
Charles Kim        Bill Burke  
Sonya Carothers    Joydip Kundu  
Michael Smith      Nitin Vengurlekar



New York    Chicago    San Francisco  
Lisbon    London    Madrid    Mexico City    Milan  
New Delhi    San Juan    Seoul    Singapore    Sydney    Toronto

Copyright © 2009 by The McGraw-Hill Companies, Inc. All rights reserved. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

ISBN: 978-0-07-162148-9

MHID: 0-07-162148-2

The material in this eBook also appears in the print version of this title: ISBN: 978-0-07-162111-3, MHID: 0-07-162111-3.

All trademarks are trademarks of their respective owners. Rather than put a trademark symbol after every occurrence of a trademarked name, we use names in an editorial fashion only, and to the benefit of the trademark owner, with no intention of infringement of the trademark. Where such designations appear in this book, they have been printed with initial caps.

McGraw-Hill eBooks are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. To contact a representative please e-mail us at [bulksales@mcgraw-hill.com](mailto:bulksales@mcgraw-hill.com).

Information has been obtained by Publisher from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Publisher, or others, Publisher does not guarantee to the accuracy, adequacy, or completeness of any information included in this work and is not responsible for any errors or omissions or the results obtained from the use of such information. Oracle Corporation does not make any representations or warranties as to the accuracy, adequacy, or completeness of any information contained in this Work, and is not responsible for any errors or omissions.

#### TERMS OF USE

This is a copyrighted work and The McGraw-Hill Companies, Inc. ("McGraw-Hill") and its licensors reserve all rights in and to the work. Use of this work is subject to these terms. Except as permitted under the Copyright Act of 1976 and the right to store and retrieve one copy of the work, you may not decompile, disassemble, reverse engineer, reproduce, modify, create derivative works based upon, transmit, distribute, disseminate, sell, publish or sublicense the work or any part of it without McGraw-Hill's prior consent. You may use the work for your own noncommercial and personal use; any other use of the work is strictly prohibited. Your right to use the work may be terminated if you fail to comply with these terms.

THE WORK IS PROVIDED "AS IS." McGRAW-HILL AND ITS LICENSORS MAKE NO GUARANTEES OR WARRANTIES AS TO THE ACCURACY, ADEQUACY OR COMPLETENESS OF OR RESULTS TO BE OBTAINED FROM USING THE WORK, INCLUDING ANY INFORMATION THAT CAN BE ACCESSED THROUGH THE WORK VIA HYPERLINK OR OTHERWISE, AND EXPRESSLY DISCLAIM ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. McGraw-Hill and its licensors do not warrant or guarantee that the functions contained in the work will meet your requirements or that its operation will be uninterrupted or error free. Neither McGraw-Hill nor its licensors shall be liable to you or anyone else for any inaccuracy, error or omission, regardless of cause, in the work or for any damages resulting therefrom. McGraw-Hill has no responsibility for the content of any information accessed through the work. Under no circumstances shall McGraw-Hill and/or its licensors be liable for any indirect, incidental, special, punitive, consequential or similar damages that result from the use of or inability to use the work, even if any of them has been advised of the possibility of such damages. This limitation of liability shall apply to any claim or cause whatsoever whether such claim or cause arises in contract, tort or otherwise.

This book is dedicated to all Oracle Database administrators in the hope that our words will be their guide to success and restful nights. And to those non-Oracle Database administrators, may you wish you, too, were using Oracle Data Guard!  
—Larry Carpenter

A quick shout out to the family—Gretchen, and my kids, Emily, Abby, and Ted. We are all hoping a lot of people buy this book so it can help pay the college bills.  
—Joe Meeks

I dedicate this book to my precious wife, Melissa, and our three boys, Isaiah, Jeremiah, and Noah, for their support during the project and sacrifice of precious family time. Thank you for your unceasing prayers and encouragement.  
—Charles Kim

I'd like to dedicate this book to my loving wife, Sandra, for the commitment of her time with me; without her support and continued motivation, my contribution to this book would not have been possible.  
—Bill Burke

To my son, Julian, thanks for your love, encouragement, and laughter.  
—Sonya Carothers

To my five-year-old daughter, Ria Rajyasri, for making my journey as a father so full of joy and wonder.  
—Joydip Kundu

I would like to dedicate my portion of this book to my wife, Tina, and two of the best daughters a father could ask for, Jessica and Madison. I know having a “computer geek” for a husband and father can at times be tedious (“but Tina, bandwidth is determined by how quickly a medium can change states”) and embarrassing (my T-shirt that has “DAD” spelled out in binary), which makes me love you guys all the more.  
—Mike Smith

I would like to dedicate this book to my kids, Ishan and Nisha; to my wife, Priya; and most importantly to my parents, whose guidance and support have always been invaluable.  
—Nitin Vengurlekar



*This page intentionally left blank*

# Contents

Foreword	xvii
Acknowledgments	xix
Introduction	xxi
<b>1 Data Guard Architecture</b>	<b>1</b>
Data Guard Overview	2
What Is Redo?	2
Redo Transport Services	5
Synchronous Redo Transport	5
Asynchronous Redo Transport	7
Redo Transport Compression	9
Automatic Gap Resolution	9
Apply Services	11
Redo Apply (Physical Standby)	12
SQL Apply (Logical Standby)	15
Can't Decide? Then Use Both!	17
Data Guard Protection Modes	18
Maximum Performance	18
Maximum Availability	18
Maximum Protection	19
Role Management Services	19
Switchover	20
Failover	21
Data Guard Management	24
Active Standby Databases	26
Offload Read-Only Queries and Reporting	26
Offload Backups	27
Testing	27
Data Guard and the Maximum Availability Architecture	29
Conclusion	29

<b>2</b>	<b>Implementing Oracle Data Guard</b> .....	31
	Plan Before You Implement .....	32
	Determining Your Requirements .....	33
	Understanding the Configuration Options .....	35
	Relating the RPO and RTO to the Protection Mode .....	62
	Creating a Physical Standby Database .....	63
	Choosing Your Interface .....	63
	Before You Start .....	64
	Using Oracle Enterprise Manager Grid Control .....	65
	The Power User Method .....	78
	Creating a Logical Standby .....	98
	Data Guard and Oracle Real Application Clusters .....	105
	Conclusion .....	106
<b>3</b>	<b>Redo Processing</b> .....	107
	Important Concepts of Oracle Recovery .....	108
	ACID Properties .....	108
	Oracle Recovery .....	109
	Life of a Transaction .....	111
	Nologging Operations .....	111
	The Components of a Physical Standby .....	114
	Real-time Apply .....	117
	Scaling and Tuning Data Guard Apply Recovery .....	118
	Parallel Media Recovery .....	119
	Tools and Views for Monitoring Physical Standby Recovery .....	120
	Physical Standby Corruption Detection .....	124
	11g New Data Protection Changes .....	124
	Data Protection and Checking on a Physical Standby .....	125
	Conclusion .....	126
<b>4</b>	<b>Logical Standby</b> .....	127
	Characterizing the Dataset Available at the Logical Standby .....	129
	Characterizing the Dataset Replicated from the Primary Database .....	129
	Protecting Replicated Tables on a Logical Standby .....	134
	Customizing Your Logical Standby Database (or Creating a Local Dataset at the Logical Standby) .....	141
	Understanding the Operational Aspects of a Logical Standby .....	145
	Looking Inside SQL Apply .....	145
	Tuning SQL Apply .....	157
	Some Rules of Thumb .....	158
	Determining Whether SQL Apply Is Lagging .....	158
	Determining Whether SQL Apply Is the Bottleneck .....	159
	Determining Which SQL Apply Component Is the Bottleneck .....	159

Troubleshooting SQL Apply	164
Understanding Restarts in SQL Apply	164
Troubleshooting Stopped SQL Apply	167
Conclusion	170
<b>5 Implementing Oracle Data Guard Broker</b>	<b>171</b>
Overview of the Data Guard Broker	172
The Broker Process Model	173
The Broker Process Flow	174
The Broker Configuration Files	176
The Broker CLI	178
Getting Started with the Broker	179
Configuring the Broker Parameters	179
The Broker and Oracle Net Services	183
RAC and the Broker	187
Connecting to the Broker	190
Managing Data Guard with the Broker	193
Creating and Enabling a Broker Configuration	193
Changing the Broker Configuration Properties	200
Changing the State of a Database	211
Changing the Protection Mode	212
Monitoring Data Guard Using the Broker	214
Removing the Broker	216
Conclusion	217
<b>6 Oracle Enterprise Manager Grid Control Integration</b>	<b>219</b>
Accessing the Data Guard Features	220
Configuring Data Guard Broker with OEM Grid Control	221
Verify Configuration and Adding Standby Redo Logs	224
Viewing Metrics	226
Modifying Metrics	227
Viewing the Alert Log File	228
Enabling Flashback Database	230
Reviewing Performance	231
Changing Protection Modes	234
Editing Standby Database Properties	236
Performing a Switchover	238
Performing a Manual Failover	240
Fast-Start Failover	243
Creating a Logical Standby	244
Managing Active Standby	250
Managing Snapshot Standby	250
Removing a Standby Database from Broker Control	250
Keeping an Eye on Availability	252
Conclusion	255

- 7** Monitoring Data Guard Implementations ..... 257
  - Monitoring the Data Guard Environment ..... 258
    - Mining the Alert Log File (PS+LS) ..... 259
    - Gathering Statistical Information from Archive Log History (PS+LS) ..... 264
    - Detecting Archive Log Gaps (PS+LS) ..... 266
    - Identifying Delays in Redo Transport (PS) ..... 268
    - Monitoring Archive Log Destinations (PS+LS) ..... 269
    - Examining Apply Rate and Active Rate (PS) ..... 271
    - Reviewing Transport and Apply Lag (PS+LS) ..... 272
    - Determining the Current Time on the Standby Database (PS) ..... 273
    - Reporting the Status of Managed Recovery Process (PS) ..... 275
  - Data Guard Menu Utility ..... 276
  - Reviewing the Current Data Guard Environment ..... 277
    - Checking the Password File (PS+LS) ..... 278
    - Checking for Nologging Activities (PS+LS) ..... 279
    - Looking at Archivelog Mode and Destinations (PS+LS) ..... 282
    - Checking Standby File Management (PS) ..... 284
    - Revealing Errors in the Data Guard Status View (PS) ..... 284
    - Logical Standby Data Guard Menu ..... 285
  - Conclusion ..... 297
  
- 8** Switchover and Failover ..... 299
  - Introduction to Role Transition ..... 300
    - Switchover ..... 300
    - Failover ..... 302
    - Switchover vs. Failover ..... 309
  - Flashback Technologies and Data Guard ..... 309
  - Performing a Switchover ..... 311
    - Configuration Completeness Check ..... 311
    - Preparatory Checks ..... 311
    - Preprocessing Steps ..... 314
    - Switching over to a Physical Standby ..... 315
    - Switching over to a Logical Standby ..... 320
    - Using the Broker or Grid Control to Switchover ..... 323
    - Switchover Health Check ..... 324
  - Performing a Failover ..... 324
    - Failing over to a Physical Standby ..... 326
    - Failing over to a Logical Standby ..... 328
    - Bringing Back the Old Primary ..... 329
    - Using the Broker or Grid Control to Failover ..... 334
    - Automatic Failover ..... 335
  - A Final Word on Multiple Standbys ..... 348
  - Conclusion ..... 348

<b>9</b>	Active Data Guard	349
	Physical Standby—Open Read-Only	350
	Why Read-Only?	351
	The Downside of Read-Only or Read-Write Mode	352
	Snapshot Standby for QA and Test Environments	353
	Read Write Standby in Oracle Database 10g	353
	Snapshot Standbys in Oracle Database 11g	357
	Real Application Testing	364
	Database Replay	365
	SQL Performance Analyzer	370
	Active Data Guard	371
	Configuring Active Data Guard	374
	Conclusion	376
<b>10</b>	Automating Site and Client Failover	377
	Defining the Problem	378
	Complete Site Failover	378
	Partial Site Failover	379
	The Nitty Gritty	379
	Connection Load Balancing and Connect Time Failover	380
	Outbound Connect Timeout	381
	Transparent Application Failover	382
	Fast Application Notification	384
	The DB_ROLE_CHANGE System Event	386
	Implementing Client Failover	387
	Complete Site Failover Configuration	387
	Conclusion	394
<b>11</b>	Minimizing Planned Downtime Using Data Guard Switchover	395
	Overview of Planned Migration	396
	Leveraging Data Guard Switchover for Planned Migration	397
	Case 1—New Data Center	397
	Case 2—Move to ASM	397
	Performing a Database Rolling Upgrade Using Data Guard	398
	Leveraging Rolling Upgrades Using SQL Apply	399
	Rolling Upgrades Using Transient Logical Standby	402
	Conclusion	408
<b>12</b>	Backup and Recovery Considerations	409
	RMAN Basics	410
	RMAN Integration with Data Guard	411
	Block Change Tracking Support	411
	Control File Management	412
	Resynchronizing the RMAN Catalog	412

- RMAN Configuration in Data Guard ..... 412
  - Example Configuration for a Primary Database ..... 414
  - Example Configuration for a Backup Standby Database ..... 415
  - Example Configuration for Other Physical Standby Databases ..... 415
- Backup Strategies ..... 415
- Backup Scenarios ..... 417
  - Backup Database Not Backed Up ..... 417
  - Full Backups on Primary ..... 417
  - Backup as Copy ..... 419
  - Image Copy Rolled Forward ..... 420
  - Standby Database Creation ..... 423
  - Backups on a Standby Database ..... 423
  - Archive Backups ..... 426
- General Recovery Strategies ..... 426
  - Media Failure ..... 426
  - Block Corruption ..... 426
  - User Errors ..... 429
- Recovery Scenarios ..... 430
  - Loss of a Datafile on a Primary Database ..... 430
  - Loss of a Datafile on a Standby Database ..... 431
  - Loss of Standby Controlfile ..... 432
  - Loss of Primary Controlfile ..... 432
  - Loss of an Online Redo Log File ..... 432
  - Incomplete Recovery of the Primary Database ..... 436
  - Recovering from a Dropped Table ..... 437
  - Recover a Missing Datafile from a Backup Taken on the Standby ..... 437
- General Best Practices ..... 440
- Conclusion ..... 441
- 13** Troubleshooting Data Guard ..... 443
  - Diagnostic Information ..... 444
    - Database Alert Logs ..... 444
    - Observer Log Files ..... 447
    - Data Guard Trace Files ..... 447
    - Data Guard Broker Log Files and Tools ..... 448
    - Dynamic Performance Views ..... 449
  - Data Guard Configuration and Management Errors ..... 450
    - Common Management Issues ..... 450
    - Physical Standby Issues ..... 456
    - Logical Standby Database Failures ..... 459
    - Switchover Issues ..... 461
    - Failover Issues ..... 463
    - Data Guard Broker Issues ..... 464
    - Errors Converting to a Snapshot Standby ..... 468

Helpful Hints and Tips .....	468
Avoid Refreshing the Standby Control File .....	468
Avoid Using the NOLOGGING Clause .....	468
OMF—Copying Control File .....	469
Conclusion .....	470
<b>14</b> Deployment Architectures .....	471
Manufacturing Company: HA Configuration .....	473
Utility Company: Zero Data Loss HA/DR .....	476
Retail Brokerage Firm: HA/DR with Zero Data Loss and Extended Geographic Separation .....	478
Government Agency: Protection from Multi-site Threats .....	480
Pharmaceutical Company: Centralized HA/DR and Data Distribution .....	483
Web Retailer: HA/DR with Reader-farm Scale Out .....	484
Insurance Company: Maximum Availability Architecture .....	486
Conclusion .....	488
<b>A</b> Data Guard vs. Array-based Remote Mirroring Solutions .....	491
The Basics .....	492
Topology .....	493
Performance .....	493
Reliability .....	494
Final Thoughts .....	495
 Index .....	 497



*This page intentionally left blank*



# Foreword

I've often said that there is one thing a DBA is not allowed to get wrong, and that is recovery. To be more general, it is the DBA's job to ensure that data that cannot ever be lost is never lost. If you cannot provide for continuous, no data loss access to all of your corporate data, you have not done the primary job a DBA should do. Providing a solid disaster recovery contingency is part of the job of the DBA, and Oracle Data Guard is the way to provide for it.

Oracle provides many features and functions to facilitate data backup, recovery, and availability. However, there are so many features that at times the implementation and configuration can be daunting. You'll have questions such as "What is the 'best way' to provide continuous availability given my circumstances?" "How do I decide between all of the configurations possible?" "What is the tradeoff of doing it one way versus the other?" "How does it all actually work under the covers?" This book covers in depth all of these questions, plus others. The authors, Larry Carpenter, Joe Meeks, Charles Kim, Bill Burke, Sonya Carothers, Joydip Kundu, Michael Smith, and Nitin Vengurlekar, are experts in the field. They are the people I go to in order to get answers myself.

The book begins by explaining the Data Guard Architecture, starting with the transaction log (REDO) information—what role it plays, how it is transmitted, and how it is ultimately used. The Data Guard architecture is built up, layer by layer, and presented in a manner that's easy to understand. You'll learn not only how the redo is transmitted, but how the receiving disaster recovery site applies (uses) the redo information. You'll learn the differences between a physical standby database and a logical standby database. You'll be introduced to Data Guard's various configuration modes—either for extreme performance on one hand or for guaranteed zero data loss on the other. You'll also learn about some everyday uses for your standby databases; they are not just for failures anymore.

The book progresses to describe the actual physical installation, setup, and configuration of your standby instances. It starts with a section on "before you even think about setting this

up, this is what you need to think about”—an approach I like. Rather than just plowing ahead and making uninformed decisions, you’ll learn about what specifically you need to ask. Important terms such as Recovery Point Objective (RPO) (the point in time to which data must be protected, which is a measure of how much “loss” would be acceptable, say from zero to a lot) and Recovery Time Objective (RTO) (the amount of time you can afford to have the data be unavailable, again from zero to a lot) are introduced and discussed. Unless you can assign some values to those metrics, you’ll find it difficult, if not impossible, to make decisions about how to configure your disaster recovery solution.

After covering how to install and configure your installation, the book addresses performance considerations, including frequently asked questions. (Believe me, I know. On <http://asktom.oracle.com/>, I see them asked frequently.) How do you tune Data Guard? How do you measure Data Guard response times? Where am I spending my time in Data Guard? All of these questions and more are covered with sections on tuning the recovery rate (the rate of application of redo at the disaster recovery site), how to perform Data Guard recovery in parallel, troubleshooting redo apply issues, and understanding the operational aspects (how it all works). To me, that is key. If you understand how something works, you are well equipped to “fix” it.

Next in line is a series of chapters on managing your Data Guard environment, either by using automated tools such as Enterprise Manager or by taking a more “do-it-yourself scripting” approach.

What follows are chapters covering something you hope never to have to do: failover. Well, they actually cover switchover, a graceful, reversible process whereby you can turn production into standby and standby into production, as well as failover. These are areas in which you will need to practice; you don’t want to find out the day you need to failover that either you don’t know how to failover, or, even worse, you cannot failover due to a mistake that was not discovered previously.

The remainder of the book covers other very useful information such as “What else can I use this standby thing for?” “How does this impact my backup and recovery procedures?” “How have other people implemented Data Guard and why did they make the choices they did?” “Why is Data Guard the right way to provide for disaster recovery for my database, and what is wrong with other methods?” And more.

In short, if you need a roadmap describing how to implement disaster recovery, what you need to think about, what are your options, and which ones you should explore, under what circumstances, then this book is for you. It combines the “How does it work?” with “How do I make it work?” in a practical, hands-on way.

—Thomas Kyte  
[asktom.oracle.com](http://asktom.oracle.com)



# Acknowledgments

We want to acknowledge our sponsoring editor, Lisa McClain, for her commitment to this book and her patience with all the authors. Thank you for understanding our busy schedules and personal conflicts while pushing us to deliver in a timely manner. This book would be delayed by another year without her involvement and nurturing.

We also want to acknowledge our acquisitions coordinator, Meghan Riley, editorial supervisor, Janet Walden, the meticulous work of copy editor Lisa Theobald, project manager Vastavikta Sharma, proofreader Paul Tyler, and the entire production and marketing team at Oracle Press. We would also like to extend our personal gratitude to our incredible technical editors, Michael Powell and Sreekanth Chintala, for their great review of all the chapters and contributions.

*—Larry, Joe, Charles, Bill, Sonya, Joy, Mike, and Nitin*

First and foremost, I'd like to thank Bernadette, my wife of 35-plus years, for putting up with my insanity and late nights while we were all working on this book. I would not have made it without her. I would also like to thank Rick Anderson and Mark W. Johnson of Oracle for first introducing me to Database Disaster Recovery, first with Oracle Rdb (originally from Digital and an Oracle product since 1994) and then with Oracle Data Guard starting with Oracle8i. Their dedication to ensuring that our customers were successful was my guide and support in my endeavors to do the same. Finally, my thanks to my manager, Ashish Ray, and our senior VP, Juan Loaiza, for allowing me to contribute to this book.

*—Larry Carpenter*

Many thanks to the development staff who have made Data Guard the best data protection and data availability solution for enterprise databases. Additional thanks to the members of Oracle's Maximum Availability Architecture team who document and validate