

# Expert Oracle Indexing and Access Paths

Maximum Performance for Your Database

Second Edition



Darl Kuhn  
Sam R. Alapati  
Bill Padfield



Apress®

## ***Expert Oracle Indexing and Access Paths***

Darl Kuhn  
Morrison, Colorado, USA

Sam R. Alapati  
Flower Mound, Texas, USA

Bill Padfield  
Aurora, Colorado, USA

ISBN-13 (pbk): 978-1-4842-1983-6  
DOI 10.1007/978-1-4842-1984-3

ISBN-13 (electronic): 978-1-4842-1984-3

Library of Congress Control Number: 2016953748

Copyright © 2016 by Darl Kuhn, Sam R. Alapati and Bill Padfield

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director: Welmoed Spahr

Lead Editor: Jonathan Gennick

Development Editor: Laura Berendson

Technical Reviewer: Michelle Malcher

Editorial Board: Steve Anglin, Pramila Balan, Laura Berendson, Aaron Black, Louise Corrigan,

Jonathan Gennick, Robert Hutchinson, Celestin Suresh John, Nikhil Karkal, James Markham,

Susan McDermott, Matthew Moodie, Natalie Pao, Gwenan Spearing

Coordinating Editor: Jill Balzano

Copy Editor: Kim Burton-Weisman

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit [www.springer.com](http://www.springer.com). Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a Delaware corporation.

For information on translations, please e-mail [rights@apress.com](mailto:rights@apress.com), or visit [www.apress.com](http://www.apress.com).

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales-eBook Licensing web page at [www.apress.com/bulk-sales](http://www.apress.com/bulk-sales).

Any source code or other supplementary material referenced by the author in this text is available to readers at [www.apress.com](http://www.apress.com). For detailed information about how to locate your book's source code, go to [www.apress.com/source-code/](http://www.apress.com/source-code/).

Printed on acid-free paper

*To Deni Luella Rossiter Staheli, she inspired greatness*

*—Darl Kuhn*

*With love to Oyuna, Evan, and my family*

*—Bill Padfield*

# Contents at a Glance

<b>About the Authors.....</b>	<b>xv</b>
<b>About the Technical Reviewer .....</b>	<b>xvii</b>
<b>Acknowledgments .....</b>	<b>xix</b>
<b>■ Chapter 1: Introduction to Oracle Indexes.....</b>	<b>1</b>
<b>■ Chapter 2: B-tree Indexes.....</b>	<b>23</b>
<b>■ Chapter 3: Bitmap Indexes .....</b>	<b>57</b>
<b>■ Chapter 4: Index-Organized Tables.....</b>	<b>77</b>
<b>■ Chapter 5: Specialized Indexes .....</b>	<b>93</b>
<b>■ Chapter 6: Partitioned Indexes .....</b>	<b>125</b>
<b>■ Chapter 7: Tuning Index Usage.....</b>	<b>157</b>
<b>■ Chapter 8: Maintaining Indexes.....</b>	<b>185</b>
<b>■ Chapter 9: SQL Tuning Advisor .....</b>	<b>211</b>
<b>■ Chapter 10: In-Memory Column Store.....</b>	<b>235</b>
<b>Index.....</b>	<b>253</b>

# Contents

<b>About the Authors.....</b>	<b>xv</b>
<b>About the Technical Reviewer .....</b>	<b>xvii</b>
<b>Acknowledgments .....</b>	<b>xix</b>
<b>■ Chapter 1: Introduction to Oracle Indexes.....</b>	<b>1</b>
Index Basics .....	2
Performance Without an Index .....	2
Implementing an Index to Improve Performance .....	4
Accessing Only the Index.....	6
Oracle Ignoring an Index.....	8
Index Basics Wrap-up.....	9
Determining Which Type of Index to Use .....	10
B-tree Indexes .....	11
Specialized Index Types.....	13
Determining Which Columns to Index .....	18
Indexes on Primary Key and Unique Key Columns .....	18
Indexes on Foreign Key Columns.....	18
Other Suitable Columns .....	19
Indexing Guidelines .....	20
Summary.....	21

<b>■ Chapter 2: B-tree Indexes</b> .....	<b>23</b>
Understanding How Oracle Uses B-tree Indexes.....	24
Prepping for B-tree Indexes .....	29
Estimating the Size of an Index Before Creation .....	30
Creating Separate Tablespaces for Indexes.....	31
Inheriting Storage Parameters from the Tablespace .....	32
Naming Standards.....	32
<b>Implementing B-tree Indexes</b> .....	<b>33</b>
Creating a B-tree Index.....	33
Reporting on Indexes.....	34
Displaying Index Code .....	35
Dropping a B-tree Index .....	36
<b>Managing B-tree Indexes with Constraints</b> .....	<b>37</b>
Creating B-tree Index on Primary Key Columns .....	37
Creating a B-tree Index on Unique Key Columns.....	42
Indexing Foreign Key Columns .....	46
<b>Multiple Indexes on the Same Column Combinations</b> .....	<b>49</b>
<b>Extended Data Types</b> .....	<b>50</b>
<b>Indexing Extended Columns</b> .....	<b>52</b>
Virtual Column Solution .....	53
Function-Based Index Solution.....	55
<b>Summary</b> .....	<b>56</b>
<b>■ Chapter 3: Bitmap Indexes</b> .....	<b>57</b>
Understanding Bitmap Indexes .....	58
Creating a Bitmap Index.....	61
Creating a Partitioned Bitmap Index.....	62
Creating a Bitmap Index on an Index-Organized Table .....	63
Performance Implications of Querying with Bitmap Indexes .....	64
Performance Implications of Loading Data with Bitmap Indexes.....	68

Understanding Bitmap Join Indexes..... 71

Creating a Bitmap Join Index ..... 72

Duplicating Indexes for Performance ..... 73

Reporting on Bitmap Indexes ..... 74

Summary..... 75

■ **Chapter 4: Index-Organized Tables..... 77**

    Understanding the Structure ..... 77

    Understanding the Advantages ..... 78

    Creating an Index-Organized Table ..... 79

    Adding an Overflow Segment ..... 81

    Compressing an Index-Organized Table ..... 84

    Building Secondary Indexes..... 85

    Rebuilding an Index-Organized Table ..... 88

    Converting to or from an Index-Organized Table ..... 89

    Reporting on Index-Organized Tables..... 90

    Summary..... 91

■ **Chapter 5: Specialized Indexes ..... 93**

    Invisible Indexes..... 93

        When to Create an Invisible Index ..... 93

        Creating an Invisible Index ..... 94

        Finding Invisible Indexes in Your Database ..... 95

        Making an Invisible Index Available to the Optimizer ..... 95

        Maintaining an Invisible Index..... 96

    Descending Indexes ..... 96

    Function-Based Indexes..... 98

        Creating a Function-Based Index ..... 99

        Limitations of Function-Based Indexes ..... 102

        Collecting Statistics for Function-Based Indexes ..... 103

    Indexes on Virtual Columns ..... 104

<b>Key Compressed Indexes .....</b>	<b>106</b>
When Key Compression is Useful .....	106
Creating a Compressed Index.....	107
Key Compression and Storage.....	109
Advanced Index Compression.....	110
<b>Composite Indexes .....</b>	<b>112</b>
Understanding Index Skip Scans and Composite Indexes.....	113
Ordering the Columns in a Composite Index .....	114
Choosing Keys for Composite Indexes.....	115
<b>Creating Virtual Indexes .....</b>	<b>117</b>
<b>Reverse Key Indexes .....</b>	<b>119</b>
Disadvantages of a Reverse Key Index.....	120
When to Use a Reverse Key Index .....	121
Creating a Reverse Key Index.....	122
Application Domain Indexes .....	122
<b>Summary.....</b>	<b>123</b>
<b>■ Chapter 6: Partitioned Indexes .....</b>	<b>125</b>
<b>Understanding Partitioned Indexes .....</b>	<b>126</b>
<b>Creating a Local Partitioned Index .....</b>	<b>126</b>
The Simplest Form .....	127
Partition-Level Requirements.....	128
Prefixed and Non-Prefixed Options.....	128
<b>Managing Primary Keys and Unique Indexes .....</b>	<b>129</b>
<b>Creating a Global Partitioned Index .....</b>	<b>131</b>
<b>Partial Indexes.....</b>	<b>135</b>
Creating Local Partial Indexes .....	136
Creating Global Partial Indexes.....	137
Performance Implications of Partial Indexes.....	138
<b>Choosing the Type of Index for Your Application.....</b>	<b>140</b>
<b>Maintaining Indexes on Partitioned Tables.....</b>	<b>141</b>



Adding a Partition .....	142
Truncating a Partition .....	143
Moving a Partition.....	143
Splitting a Partition .....	144
Exchanging a Partition.....	145
Dropping a Partition.....	146
Merging a Partition .....	147
<b>Rebuilding Global Partitioned and Non-Partitioned Indexes.....</b>	<b>147</b>
<b>Setting Index Partitions as Unusable and then Rebuilding.....</b>	<b>150</b>
<b>Index Implications for Interval Partitioning .....</b>	<b>152</b>
<b>Making Older Data Read-Only .....</b>	<b>153</b>
<b>Reporting on Partitioned Indexes .....</b>	<b>153</b>
<b>Summary.....</b>	<b>155</b>
<b>■ Chapter 7: Tuning Index Usage.....</b>	<b>157</b>
Optimizer Access Paths.....	157
Index Scans.....	158
Index Unique Scan.....	158
Index Range Scan.....	159
Index Skip Scan .....	161
Index Full Scan .....	162
Index Fast Full Scan .....	163
Determining Whether a Query Uses an Index.....	163
Avoiding an Index .....	165
Avoiding All Use of an Index .....	165
Avoiding Only the Fast Full Scan .....	166
Forcing a Table Scan .....	166
Choosing Between an Index and a Table Scan.....	166
Why the Optimizer May Ignore Indexes.....	167
Number of Distinct Rows.....	168
Index Clustering Factor.....	168

<b>How Index Access Paths Can Change Without New Statistics .....</b>	<b>169</b>
Using the NOT EQUAL Condition .....	169
Querying with Wild Characters .....	171
Referencing Null Values in Predicates .....	172
Writing Functions in a Query .....	173
Skipping the Leading Portion of an Index .....	174
<b>Forcing the Optimizer to Use an Index .....</b>	<b>175</b>
Applying the INDEX Hint .....	175
Applying Related Hints .....	176
Troubleshooting a Failed INDEX Hint .....	177
Adjusting the optimizer_index_cost_adj Parameter .....	179
Collecting Accurate Statistics for an Index .....	180
<b>Parallelizing Index Access .....</b>	<b>181</b>
<b>Summary .....</b>	<b>183</b>
<b>■ Chapter 8: Maintaining Indexes.....</b>	<b>185</b>
<b>Gathering Statistics for Indexes .....</b>	<b>185</b>
The DBMS_STATS Package .....	185
The METHOD_OPT Parameter .....	187
<b>Working with Unusable Indexes .....</b>	<b>188</b>
Making an Index Unusable .....	189
Specifying the SKIP_UNUSABLE_INDEXES Parameter .....	191
Asynchronous Global Index Maintenance .....	192
<b>The Role of the INDEX_STATS View in Index Rebuilds .....</b>	<b>193</b>
Benefits of the INDEX_STATS View .....	193
Problems with the INDEX_STATS View .....	195
<b>Index Rebuilding: The Debate .....</b>	<b>198</b>
Arguments for Rebuilding .....	198
Arguments Against Rebuilding .....	199
<b>Coalescing Indexes to Reduce Fragmentation .....</b>	<b>199</b>
<b>Shrinking Indexes to Reduce Fragmentation .....</b>	<b>200</b>
<b>Moving Tables and Indexes .....</b>	<b>201</b>

Improving Index Creation Efficiency .....	202
Parallelizing Index Creation .....	202
Avoiding Redo Generation During Index Creation .....	204
Using Larger Block Sizes .....	204
Compressing Indexes .....	205
Using Multiple Options Together .....	205
Generating the DDL for Creating an Index.....	205
Using the DBMS_METADATA Package .....	205
Using the SESSION_TRANSFORM Procedure.....	207
Using the SET_FILTER Procedure.....	207
Using Data Pump .....	208
Dropping an Index .....	208
Dropping an Index Online .....	209
The Hazards of Dropping an Index.....	209
Summary .....	210
<b>■ Chapter 9: SQL Tuning Advisor .....</b>	<b>211</b>
Automatic SQL Tuning Job.....	212
Verifying Automatic Jobs Running.....	213
Viewing Automatic SQL Tuning Job Advice .....	213
Generating a SQL Script to Implement Automatic Tuning Advice .....	216
Disabling and Enabling Automatic SQL Tuning .....	216
Managing SQL Tuning Sets.....	218
Viewing Resource-Intensive SQL in the AWR.....	218
Viewing Resource-Intensive SQL in Memory.....	220
Populating SQL Tuning Set from High-Resource SQL in AWR .....	222
Populating a SQL Tuning Set from High-Resource SQL in Memory .....	223
Populating SQL Tuning Set with All SQL in Memory .....	224
Displaying the Contents of a SQL Tuning Set.....	225
Selectively Deleting Statements from a SQL Tuning Set .....	227
Adding Statements to an Existing SQL Tuning Set.....	228
Dropping a SQL Tuning Set.....	228

<b>Running the SQL Tuning Advisor .....</b>	<b>228</b>
Creating a Tuning Task.....	230
Execute DBMS_SQLTUNE and View the Advice .....	232
Viewing and Dropping Tuning Tasks .....	232
Running SQL Tuning Advisor from SQL Developer .....	233
Running SQL Tuning Advisor from Enterprise Manager .....	233
<b>Summary.....</b>	<b>234</b>
<b>■ Chapter 10: In-Memory Column Store .....</b>	<b>235</b>
Traditional Row Store Format.....	235
Columnar Format.....	236
Use Cases for IM Column Store.....	238
Enabling a Database for the IM Column Store.....	238
Enabling a Table for IM Column Store .....	240
Excluding Columns .....	242
Enabling New Tables to Automatically Use In-Memory .....	243
Determining if a Query Is Used in the IM Column Store .....	243
Viewing the Execution Plan .....	244
Displaying Session-Level Statistics.....	245
Querying the Data Dictionary.....	246
Observing Performance Differences.....	248
Setting a Priority.....	248
Sizing the IM Column Store Memory Area.....	249
Adjusting Compression .....	249
Utilizing IM Column Store with RAC .....	250
Summary.....	251
<b>Index.....</b>	<b>253</b>

# About the Authors



**Darl Kuhn** is a DBA/developer working for Oracle. He also teaches Oracle classes at Regis University in Denver, Colorado. He is an active member of the Rocky Mountain Oracle Users Group. Darl enjoys sharing knowledge, which has led to several book projects over the years.



**Sam R. Alapati** is a principal administrator at Sabre Holding in Southlake, Texas. He has written several Oracle and WebLogic Server books. Sam is passionate about data and databases, and he enjoyed writing his part of this book.

■ ABOUT THE AUTHORS



**Bill Padfield** is an Oracle Certified Professional. He works for a large telecommunications company in Denver, Colorado, as a senior database administrator. Bill helps administer and manage a large data warehouse environment consisting of more than 90 databases. He has been an Oracle database administrator for more than 19 years and he has worked in the IT industry since 1985. Bill also teaches graduate database courses at Regis University.

# About the Technical Reviewer



**Michelle Malcher** (@malcherm) is a well-known volunteer leader in security and database communities. She is currently on the board of directors for the FUEL Palo Alto Network User Group and has served on the board of directors for the Independent Oracle User Group (IOUG). As an Oracle ACE director, she has had the opportunity to present around the world at user group conferences. She continues to participate on the conference committee for the IOUG. Her technical expertise on topics from database to cybersecurity, as well as her senior-level contributions as a speaker and author, have aided many corporations across architecture and risk assessment, purchasing and installation, and ongoing systems oversight. These corporations include Wells Fargo & Company, where Michelle is currently team lead for Enterprise Database Security.

# Acknowledgments

Thanks to Jonathan Gennick, Jill Balzano, and the Apress staff; it takes a quality team to produce a quality book. Thanks also to Michelle Malcher who made numerous suggestions and modifications to improve the quality of this book. Also thanks to the many people I've learned from over the years: Dave Jennings, Scott Schulze, Andy Brown, Mark Lutze, Nitin Mittal, Venkatesh Ranganathan, Valerie Eipper, Mike Tanaka, Simon Ip, Laurie Bourgeois, Scott Grover, John Finocchiaro, Kevin O'Grady, Jack Chung, Akshay Trivedi, Suresh Raju, Mosharof Hossain, Brent Davis, Evangelia "Valia" Antonatou, Peter Compel, Petr Fojtik, Tim Spangler, Jan Pewner, Michael Rohrbacher, Mingyan Fan, Ashwin Rao, Chandupavan Panditi, Naveen Shankaranarayana, Dustin Garvey, Sampanna Salunke, Rahul Khandelwal, Siraj Sherriff, Gaurav Mehta, Janet Bacon, Naomi Kuhn, Peter Schow, Todd Sherman, Patrick David, Carson Vowles, Aaron Isom, Tim Gorman, Tom Kyte, Jaden Jackson, Brad Stucki, Scott Lovell, Terri Cannon, and Jim Stark.

—Darl Kuhn

I wish to acknowledge the great help provided by Jill Balzano, the coordinating editor, in seeing this book through all the stages without a problem. I appreciate Jill's patience and good cheer, to say nothing about her efficiency in keeping track of all the changes in a book co-written by three authors.

As with all of my other books, I wish to thank the great help and support (and love) from my family: Valerie, Nina, and Nicholas, as well as Shannon, Keith, Shawn, and Dale. I wish to acknowledge the great debt to my other family as well: my mother, Swarna Kumari; my brothers, Hari and Siva Sankara Prasad; my sisters-in-law, Aruna and Vanaja; and my nieces and nephews, Ashwin, Teja, Appu, and Soumya.

—Sam R. Alapati

As always, thanks to Apress for the opportunity to contribute to their great library of publications. Thanks to Jonathan Gennick and Jill Balzano for their input, support, and patience. Thanks to my great co-authors, Darl and Sam.

Doing this helps me reflect on all the help I've had in my career. This includes, but is not limited to the following individuals. To some of my managers over the years—thanks to Bob Ranney, Beth Bowen, Larry Wyzgala, John Zlamal, Linda Scheldrup, Amy Neff, Maureen Frazzini, Billie Ortega, and Chad Bower. Thanks to my current team, which includes Dave Carter, Sandy Hass, Atul Shirke, Kevin Tomimatsu, John Townley, Paul Watkins, Sunkarsh Guduthur, Rajendra Mishra, Ulagammal Kaliyappan, and Gopinath Kanagal.

Over the years, I've learned an awful lot from countless people. Some of the people that have helped me along the way include Mark Nold, Mick McMahon, Sandra Montijo, Roby Sherman, Adric Norris, Brent Wagner, Tony Arlt, Jerry Sanderson, Mike Hammontre, Pat Cain, Gary Whiting, Ron Fullmer, Becky Enter, John Weber, Dianna Belangee, Randy Rezac, John Chatto, Pankaj Guleria, Avanish Gupta, Scott Bunker, Paul Mayes, Rick Barry, Sue Wagner, Glenn Balanoff, Linda Lee Burau, Deborah Lieou-McCall, Kristi Sargent, George Huner, Pad Kail, Curtis Gay, Ross Bartholomay, Carol Rosenow, Sheryl Gross, John Piel, Kim Lake, Rob Grote, Zane Warton, Pat Wuller, Wil Minor, Jim Barclay, Jason Hermstad, Shari Plantz-Masters, Denise Duncan, Bob Mason, Rob Coates, Debbie Chartier, Blair Christensen, Meera Ganesan, Kedar Panda, Kirsten Franz, Valarie Martinez, Mallick Robyeat, Liping Zhang, Satish Pasupathy, John Lungerhausen, Richard Rawles, Cindy Powell, Mary Ann Wooten, Deb Kingsley, Paula Obering, and Deb Sutton.

—Bill Padfield