

Expert Oracle Practices

Oracle Database Administration
from the Oak Table



**Melanie Caffrey, Pete Finnigan, Randolph Geist, Alex Gorbachev,
Tim Gorman, Connie Green, Charles Hooper, Jonathan Lewis,
Niall Litchfield, Karen Morton, Robyn Sands, Jože Senegačnik,
Uri Shaft, Riyaj Shamsudeen, Jeremiah Wilton, Graham Wood**

Apress®

Expert Oracle Practices: Oracle Database Administration from the Oak Table

Copyright © 2010 by Melanie Caffrey, Pete Finnigan, Randolph Geist, Alex Gorbachev, Tim Gorman, Connie Green, Charles Hooper, Jonathan Lewis, Niall Litchfield, Karen Morton, Robyn Sands, Jože Senegačnik, Uri Shaft, Riyaj Shamsudeen, Jeremiah Wilton, Graham Wood

All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval system, without the prior written permission of the copyright owner and the publisher.

ISBN-13 (pbk): 978-1-4302-2668-0

ISBN-13 (electronic): 978-1-4302-2669-2

Printed and bound in the United States of America 9 8 7 6 5 4 3 2 1

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

President and Publisher: Paul Manning

Lead Editor: Jonathan Gennick

Technical Reviewers: Melanie Caffrey, Arup Nanda, Peter Sharman

Editorial Board: Clay Andres, Steve Anglin, Mark Beckner, Ewan Buckingham, Gary Cornell, Jonathan Gennick, Michelle Lowman, Matthew Moodie, Jeffrey Pepper, Frank Pohlmann, Ben Renow-Clarke, Dominic Shakeshaft, Matt Wade, Tom Welsh

Coordinating Editor: Fran Parnell

Copy Editors: Sharon Wilkey, James A. Compton

Compositor: Molly Sharp

Indexer: Brenda Miller

Artist: April Milne

Cover Designer: Anna Ishchenko

Distributed to the book trade worldwide by Springer-Verlag New York, Inc., 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax 201-348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com.

For information on translations, please e-mail info@apress.com, or visit 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/info/bulksales.

The information in this book is distributed on an “as is” basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author(s) nor Apress shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this work.

The source code for this book is available to readers at www.apress.com. You will need to answer questions pertaining to this book in order to successfully download the code.

Contents at a Glance

Foreword	xix
About the Authors.....	xxiii
About the Technical Reviewers	xxix
■ Chapter 1: Battle Against Any Guess	1
■ Chapter 2: A Partly Cloudy Future	17
■ Chapter 3: Developing a Performance Methodology	35
■ Chapter 4: The DBA as Designer	73
■ Chapter 5: Running Oracle on Windows	111
■ Chapter 6: Managing SQL Performance.....	131
■ Chapter 7: PL/SQL and the CBO	153
■ Chapter 8: Understanding Performance Optimization Methods	173
■ Chapter 9: Choosing a Performance Optimization Method.....	297
■ Chapter 10: Managing the Very Large Database	347
■ Chapter 11: Statistics	369
■ Chapter 12: Troubleshooting Latch Contention	399
■ Chapter 13: Measuring for Robust Performance	441
■ Chapter 14: User Security.....	467
■ Chapter 15: Securing Data.....	507
■ Index.....	533

Contents

Foreword	xix
About the Authors.....	xxiii
About the Technical Reviewers	xxix
■ Chapter 1: Battle Against Any Guess	1
Guess Hunting	1
Why Do We Guess?.....	3
Understanding a Problem	5
Logical Conclusions vs. Historical Observations	6
Knowledge Is Power	8
RTFM	9
Facing the Unknown.....	11
Paradigm Shifts	11
Experience Is Danger.....	12
Fixing the Root Cause?	13
Best Practices and Myths	14
BattleAgainstAnyGuess.com.....	15
■ Chapter 2: A Partly Cloudy Future	17
What Is Cloud Computing?	17
Software as a Service (SAAS).....	18
Platform as a Service (PAAS).....	18
Infrastructure as a Service (IAAS).....	18
Who Are the Cloud Providers?	19
Sun	19

Salesforce.com.....	19
Google.....	19
Microsoft	20
Amazon.com.....	20
Running Oracle on Amazon’s Cloud.....	20
But Is It Supported?	21
Making the Cloud Concrete	21
Prerequisites	23
How Do You Work This Thing?.....	23
Starting Out: Getting a Suitable Operating System Running	26
Persistent Storage	28
Simple Storage Service (S3).....	28
Elastic Block Storage (EBS)	29
EBS Performance for Oracle.....	29
Attaching and Configuring EBS Storage	31
Persistence Approaches	32
Method A: Generic AMI and EBS File System	32
Method B: Custom AMI	33
Method C: Boot from EBS	33
Oracle Backup on EC2: The OSB Cloud Module.....	33
Summary	34
Chapter 3: Developing a Performance Methodology	35
What Is Performance?	35
The Early Days.....	35
Time-Based Performance Analysis.....	36
Performance Strategy.....	36
Design and Development.....	37
Common Design Pitfalls	38
Lightweight Performance Measures.....	39

Quality Assurance.....	39
Testing for Performance.....	39
Capturing Resource Utilization and Outlines	40
New Software and Upgrades.....	41
Know Your Hardware.....	41
Verify Statistics Collection.....	42
Back Up the Optimizer Statistics	43
Implement Change Incrementally.....	43
Post Installation or Upgrade	43
Reactive Tuning.....	44
Step 1: Define the Problem.....	45
Step 2: Examine the Performance Data.....	49
Step 3: Formulate a Theory	52
Step 4: Implement and Verify the Solution	53
Diagnostic Tools	54
Using and Interpreting ADDM	55
Using and Interpreting the ASH Report.....	56
Using and Interpreting the AWR and Statspack Instance Reports.....	60
Meaning of Key Statistics.....	67
Time-Based Statistics	67
% Activity and Average Active Sessions.....	68
ASH-Estimated DB time.....	69
V\$OSSTAT	70
Wait Classes	70
The Optimizer	70
Managing Statistics.....	71
Locking Statistics	71
Execution Plan Stability and Profiles	71
Summary	71

Chapter 4: The DBA as Designer	73
When to Get Involved in Application Design	74
Be Approachable	74
Ask for Periodic Sign-off on Design and Application Milestones	75
Attend Code Reviews.....	75
Hold Postmortems	76
Partnership Between DBAs and Developers.....	76
Hold Brown Bag Sessions.....	77
Sit Near Each Other	77
Be Open to New Ideas	77
Be on the Same Side	78
Design-First Methodologies vs. Agile Techniques.....	78
Design-First Approach.....	79
Agile Software Development	80
Pros and Cons of Each Methodology	80
Schema Design.....	83
Choose Your Datatypes Carefully	83
When Bigger Is Not Better	91
Heaps of Trouble	92
Faster, Not Harder	100
Other Design Considerations	102
Middle Tier vs. Database	102
Flexibility, Security, Speed	103
The Importance of Having Integrity	103
Don't Be High Maintenance	105
The DBA as Database Evangelist.....	105
Reading the Documentation and Keeping Current	107
Knowing, Testing, and Teaching Your Software's Features.....	108
Learning from Your Mistakes and Experiences	110

Triages and Postmortems.....	110
Constant and Iterative Knowledge Sharing	110
■ Chapter 5: Running Oracle on Windows	111
Architecture.....	111
CPU Resources	112
Memory	117
Disk.....	119
Management.....	119
The Registry.....	119
Services.....	122
Scripting	126
Summary	130
■ Chapter 6: Managing SQL Performance.....	131
Adopting a Performance Mindset.....	131
Defining and Measuring Performance.....	133
EXPLAIN PLAN.....	133
DBMS_XPLAN	135
Extended SQL Trace Data	138
Interpreting Performance Data.....	141
Case 1: The Lack of a <i>Good</i> Index.....	141
Case 2: The Presence of Unidentified Data Skew.....	142
Case 3: SQL That Should Be Rewritten.....	144
Case 4: SQL That Unnecessarily Invokes PL/SQL	148
Summary	151
Further Reading.....	152
■ Chapter 7: PL/SQL and the CBO	153
Reviewing the Basics	153
Parsing Phase.....	153
Execution Plan Preparation	154

Using the Extensible Optimizer.....	156
User-Defined Statistics.....	156
User-Defined Selectivity.....	157
User-Defined Cost.....	157
Creating an Example.....	157
Creating Some Example Objects.....	157
Running an Example Query.....	160
Giving the CBO Better Information.....	162
Understanding How It Works.....	162
Indicating Default Selectivity and Default Cost.....	164
Specifying Defaults (Syntax).....	165
Determining a Default Cost.....	165
Breaking the Association.....	165
Influencing the Execution Plans.....	166
Example 1: Influence of the Increased Cost.....	166
Example 2: Influence on the Order of Operations with Default Statistics.....	168
Example 3: Influence on the Order of Operations.....	170
Summary.....	172
Chapter 8: Understanding Performance Optimization Methods.....	173
Blindly Changing Parameters.....	174
Monitoring and Reacting to the BCHR.....	174
Monitoring Delta Values of System/Session Stats.....	182
Monitoring File Activity.....	184
Monitoring the Delta Values of System/Session Waits.....	190
Monitoring CPU Utilization.....	196
CPU Load Generators.....	197
Determining the CPU Run Queue.....	198
Determining CPU Utilization.....	201
Sampling Performance with Low Overhead.....	203

Capturing Some Statistics	203
Decision Tree for Quickly Interpreting the Statistics	209
Creating Statspack or AWR Reports	212
Monitoring the Delta Values for SQL Statements	215
Examining Execution Plans and Plan Statistics	219
Examining Optimizer Parameters Affecting Plans	227
Generating 10053 Cost-Based Optimizer Traces.....	230
Activating and Deactivating the Optimizer Trace	230
Query Blocks.....	231
Peeked Bind Variables.....	231
Optimizer Parameters Used.....	233
Transformations	233
System Statistics.....	235
Base Statistical Information	236
Dynamic Sampling.....	238
Single Table Access Path	241
General Plans.....	242
Plan Table.....	243
Query Block Registry	244
Hints	245
The Query	245
Generating 10046 Extended Traces.....	246
Brief Summary of a Raw 10046 Extended Trace File's Contents	247
Enabling a 10046 Extended Trace	249
Disabling 10046 Tracing.....	253
Sample Trace File Analysis with Oracle 11.1.0.7	253
Examining Server Stack Traces.....	260
Generating a Trace File on Error.....	260
Initiating a Trace with SQL*Plus ORADEBUG	261

Operating-System-Generated Stack Traces	275
Reviewing the Enterprise Manager ADDM Findings	275
Examining Network Packets.....	279
Examining Client-Side Traces.....	283
SQL*Net Tracing	283
Process Monitor Tracing.....	285
Spy++ Tracing.....	286
Investigating Enqueue Waits	286
Summary	291
Chapter 9: Choosing a Performance Optimization Method.....	297
Decision Tree for Performance Monitoring.....	297
Performance Problems Not Yet Reported.....	298
Problems Reported by End Users	298
Problems Reported by IT Staff.....	299
Sample Investigations	300
Quick Checkup.....	300
Problem After Upgrading the Oracle Release Version	306
Problem After Upgrading the ERP Version.....	313
Performance Optimization Issues.....	321
Inefficient SQL	321
Verify the Inefficiency.....	321
Collect Additional Data	322
Verify That the Trace File Covers Only One Test.....	322
Verify That the Trace File Is Complete.....	323
Verify That the Issue Is a Database Issue.....	325
Determine Whether It Is a Parse or Execution Problem.....	327
Parse Performance Issues.....	328
Majority of Parse Time Spent on the CPU.....	329
Majority of Parse Time Spent on Wait Events.....	331

High Execution Time or Fetch Time Issues.....	331
General Optimizer Settings and Object Statistics.....	331
Histogram Issues.....	332
Common Parameters Influencing Optimizer.....	333
Statement and Physical Design Issues.....	334
Data Access Issues.....	335
Optimizer Not Using (Correct) Index	337
Pagination (Top N) Queries	338
Processing Large Result Sets.....	339
Join Issues.....	340
Parallel Processing Issues.....	341
Shared Pool Abuse	342
Resolving Shared Pool Abuse	343
General Guidelines for Investigating Shared Pool Abuse.....	344
Chapter 10: Managing the Very Large Database	347
Designing (or Retrofitting) a VLDB.....	348
Infinity Is So Imprecise... ..	349
Partitioning	351
Everything Is a Segment.....	353
Data Manipulation with Partitioning	353
Partition Pruning.....	357
Partition Configuration.....	358
Information Life Cycle Management.....	360
Backup Optimization and Guaranteed Recovery	362
Further Notes on Storage	364
Limits of Which to Be Aware	365
Database Block Size	365
Number of Files in a Database	365
Storage That Can Migrate.....	366

Parameter READ_ONLY_OPEN_DELAYED	367
Summary	367
Chapter 11: Statistics	369
It Can't Be Done!	369
Subquery Anomaly	370
Partition Elimination	374
Lack of Knowledge	378
Problems with Statistics	385
Timing	387
Multinationals	387
Partitioning	387
Batch Jobs	388
Creating Statistics	388
Other Stats	394
Baseline	396
Summary	397
Chapter 12: Troubleshooting Latch Contention	399
Latches and Why We Need Them	399
Solitaire, Parent, and Child Latches	400
Operational Specifics	401
Immediate Mode	401
Willing-to-Wait Mode	403
Latch-Wait Posting Mode	404
Identifying and Analyzing Latch Contention	404
Step 1: Identify Latches Causing Contention	404
Step 2: Review Distribution of Gets	406
Step 3: Check the Code Path	406
Cache Buffers Chains Latch Contention	407
Common Causes of CBC Latch Contention	409

Analyzing CBC Latch Contention.....	410
Resolving CBC Latch Contention	415
Shared Pool Latch Contention	419
Structures in the Shared Pool.....	419
Common Causes of Shared Pool Latch Contention	421
Analyzing Shared pool Latch Contention	422
Resolving Shared Pool Latch Contention.....	428
Library Cache Latch Contention.....	429
Common Causes of Library Cache Latch Contention.....	431
Analyzing Library Cache Latch Contention	431
Library Cache Latches and Mutexes.....	432
Resolving Library Cache Latch Contention	432
Enqueue Hash Chains Latch Contention.....	434
Common Causes of Enqueue Hash Chains Latch Contention	435
Analyzing Enqueue Hash Chains Latch Contention	436
Resolving Enqueue Hash Chains Latch Contention	438
Advanced Help for Latch Contention Problems	439
The v\$latch_parent View.....	439
The spin_count Parameter	439
The _latch_classes and _latch_class_N Parameters.....	439
The _latch_wait_posting and _enable_reliable_latch_waits Parameters	440
Summary	440
■ Chapter 13: Measuring for Robust Performance	441
Finding the Red Rocks.....	442
Understanding the Properties of Performance	443
Response Time Is Key	443
Throughput Counts	443
Meeting Expectations Matters.....	444
All Together Now...	444

“Tuning” a Data Warehouse	445
Initial Tuning	445
Repeating the Analysis	447
Exploring What Variance Can Tell Us About a Process	448
Distribution Analysis	449
Distribution of Elapsed Time Data	452
Variance	452
The Index of Dispersion	453
What About Standard Deviation?	453
Elapsed Time Data Sources	454
Achieving Robust Performance	455
Designing an Experiment	456
Using Instrumentation	457
Measuring the Results	459
Tolerance Ranges and Process Capability	463
What is “Too Much” Variation?	464
Measuring Variance Within Oracle Sample Sets	464
Sampling from Samples	464
Summary	465
Chapter 14: User Security	467
Securing User Accounts	468
User Enumeration	469
Splitting the Task in Two	470
Dealing with Oracle Database 10g and Prior	472
Feature Analysis	475
Accounts That Can Definitely Be Removed	477
Accounts That Definitely Have to Remain in the Database	478
Accounts to Analyze Individually	479
Reduction of Accounts	482

Account Password Strength	487
Cracking Passwords with a PL/SQL Password Cracker.....	488
Cracking Passwords with a “Real” Password Cracker	490
Fixing Weak Passwords.....	497
Roles and Privilege Assessment.....	498
Have Accounts Been Used?.....	498
Have Accounts Been Shared?	499
Password Management.....	500
Audit Settings	503
Summary	504
Quiz Answer.....	505
■ Chapter 15: Securing Data.....	507
Identifying Key Data.....	508
Locating the Database Table	509
Direct Table Privileges.....	510
Understand the Hierarchy.....	513
Other Methods to Read Data	519
Access to Access.....	525
Duplicating Data	527
Generalizing Across the Database.....	530
Summary	531
Quiz Answer.....	532
■ Index.....	533

Foreword

Like many red-blooded Americans, my friend—let’s call him John—dabbles in carpentry whenever he can motivate himself to move from the computer to produce something more tangible. Recently he discovered that although motivation can be a great catalyst, it can never replace skills. In a grand testimony to that truth, he has never been able to produce anything more than a hole in a board or two pieces of wood out of one. His most recent story parallels the teachings of this book.

Embarrassed by the incessant boasting of neighbors and friends about their great accomplishments in the fine art of carpentry, John woke up one day determined to turn the tide—become skilled at carpentry. He even picked out a project—a garden shed to store all the mysterious but seemingly useful stuff that occupied the modest real-estate of his garage, forcing his cars to become refugees on the driveway. No way, no sir, he told the cars—pretty soon they would be able to go back to their promised homeland, just as soon as he banished the clutter to the garden shed, which he would build with his newly acquired skill. The cars seemed to honk in agreement, or perhaps he imagined it.

Charged with new passion, he rushed off to a store for homebuilders full of other trumped-up newbies like himself. He chose a book on do-it-yourself garden sheds. He bought the materials and the all-important tools. He did everything by the book. But two years have passed by, and all he has been able to produce are numerous cuts, bruises, scratches, and countless pieces of wood wasted as a result of not being cut properly—either cut too short or angled too acutely. In the course of these years, my friend added more tools to the collection—tools that supposedly make a carpenter out of anyone with the right inclination and bent of mind—and more wasted wood, drill bits, and saw dust. The shed, not even a remote resemblance of it, never saw the light of the day.

What went wrong? The missing ingredient was the right foundation. My friend was not a skilled carpenter. All the good intentions and excellent tools didn’t do a thing to get him the right skills. What he should have invested in early on was time in learning the trade. He should have spent time in small but sure ways, learning how to use the tools he acquired. He should have learned how to measure pieces of wood before cutting them. No, he should have learned how to measure twice before cutting! He should have learned how to master the tools before actually using them on a project. But, above all, he should have realized that sophisticated tools and do-it-yourself books are not substitutes for a conceptual foundation on anything that requires expertise—the building of garden sheds included.

I tried to persuade my friend to give up on that project and hire a professional. But John doesn’t give up easily. Instead, he began anew. This time he invested in learning the craft, in making small cuts and completing small projects to build up his skills. Once he was confident, he started on the shed. Finally, late one fall and after an embarrassing amount of time, the shed was done. John put it on concrete slabs and stored all his lawn and garden tools in it. For the first time in years, John’s cars went inside his garage. The neighbors and I gathered. Everyone was smiling. John was beaming as he showed his handiwork to us all.

Then, the snow came.

You see, it was fall 2008, and the 2008–2009 winter turned into one of worst we had seen, with snowfalls going up to 13 inches or so every other day. The snow and ensuing cold spell left a huge pile of ice on John’s beloved shed. One day, after days of enduring under the heavy accumulation, the shed buckled under the weight. We all watched helplessly as his blood and sweat just went up in a pile of wood, exposing everything inside it.

We neighbors tried to enlighten John on the futility of his attempts, but he won't budge. He spent hours and hours fixing the broken shed, but after every effort it would break. His attempts drew mixed emotions from the neighbors. Most people applauded his attempts at a positive attitude by making these efforts to fix the shed. But there were a few well-wishers who advised him to give up, to admit failure, and to have a prebuilt shed installed by professionals.

Where did John go wrong this second time? At least he started off in the right way, but he didn't have the expertise in the business of sustenance. Building a shed is only half the story; making it tough enough to withstand the elements is the other half—one that proverbially separates the men from the boys.

How does John's story parallel the teachings in the book? In many ways, John's experience paraphrases the activity of many professionals engaged in the technology trade—Oracle technology included. Even some so-called experts are not immune to the temptations of a shortcut. The dependence on tools is a foregone conclusion in many cases, with greater emphasis on seemingly sophisticated ones. What's worse, it is a common belief in the corporate world that a go-getter attitude coupled with the right tools enables one to do any job, and do it well. That belief often leads to a reinforcement of the pattern of behavior that rewards an effort—not necessarily the right effort.

The expert must have the right tools, but the sophistication of the tools is not as important as their right use. Using the handle end of a screwdriver to drive a nail when a hammer is lying right in front of you is not just idiotic but goes against the “no screwdrivers were harmed in the making of this project” theme. Using the right tool for the right job is what makes an expert an expert. Consider Chapter 3, in which Connie, Uri, and Graham show you how to use ADDM to resolve some thorny issues in performance. Many people I meet often balk at using a tool such as ADDM. Why? Their reasons are often a mixture of ignorance, skepticism, and masochistic viewpoints that a tool can't be better than their bare hands. But for a DBA to avoid using ADDM is like a nurse in a busy hospital wanting to throw away the automatic blood-pressure monitor. Sure, a nurse can get by with a manual cuff. So can an expert DBA get by without ADDM. But ADDM makes life easier and more efficient, and an expert will take advantage of ADDM in order to get those benefits.

Getting back to the story, Alex shows in Chapter 1 how the behavior of doing just *something*, even doing something completely ineffective, is often rewarded. Work that does not get results represents lost productivity and delayed execution. Alex painstakingly demonstrates how it hurts productivity to make assumptions rather than to take accurate measurements, and that a major impediment to the resolution of any problem is the presence of assumptions. Many chapters in this book try to drive home the point that the practice of measuring instead of assuming should be the first order of business in any resolution process. In Chapter 6, Karen shows how to instrument and measure code for hidden issues, and to rely on proper measurement techniques. In Chapters 8 and 9, Charles and Randolph describe how to develop a system that measures performance and practically puts the culprit of a performance problem right in front of you. And Robyn shows in Chapter 13 how to use statistics to find the proverbial “red rock” representing a performance problem buried under a mountain of collapsed rubble.

About eight years ago, I was delivering a technical session on buffer busy waits, how to diagnose them, and how resolve them. At the end of the session, during the question-and-answer period, an attendee asked me a very interesting question: SQL Server didn't suffer conditions such as buffer busy waits, so did that mean that SQL Server was better than Oracle? I resorted to an allegory as a response. I said that I took my car for oil changes, but not my bicycle; so did it mean that my bicycle is technologically superior to my car? The response, as expected, drew a chuckle from the audience and drove home the point about the futility of comparison between those two databases. In the audience was Tim Gorman, one of the authors of this book, who also responded to the question from the attendee. Tim asked the attendee how he knew that SQL Server did not indeed have the condition known as buffer busy waits. Perhaps SQL Server did suffer from those, but he just didn't know how to measure them and resolve them.

While the audience was mulling over the response, I realized the enormous universal truth behind Tim's response. Most of the audience, even seasoned DBAs, weren't even aware of the buffer busy waits

in Oracle, or were quite hazy in their understanding of them. The point was not about the specific condition or what it is called. The point was about the universal acceptance of the absence of something that has not been characterized yet. The questioning of accepted knowledge is the key to getting closer to the truth. But it is not just the act of skepticism, but the examination of the knowledge that matters. Centuries ago, many believed the world was flat. Had Newton not discovered gravity, the concept of a round earth might have defied logic and been relegated to the category of myth. Instead, the ability to understand gravity and correlate that understanding with other observations about the world led to widespread acceptance that the world was, in fact, round.

In Chapter 12, Riyaj explains the little known and often misunderstood concept of latch contention. Lack of understanding of latch contention leads to design of systems that do not scale well, and to issues that remain unresolved for quite long periods of time. Riyaj's chapter should also educate readers enough to explain some of the other mysteries that have plagued them. The concept of deeper understanding to dispel myths is reinforced in Jonathan's Chapter 11 on statistics. These chapters highlight the need to build a system for scalability, which was the precise cause of the failure in the case of the garden shed made by my friend John.

Many of the authors, such as Jonathan and Jože, portray themselves not as consultants or DBAs extraordinaire, but merely as researchers. This self-attribution is a testimony to the fact that they are still learning and will continue to do so. They have epitomized the belief that knowledge is never a destination; it's a journey. It is not supposed to end. Any pretense of accomplishing the end is exactly that—a pretense. The other authors hold different titles, but all of them espouse the same principle: they never seek the end of knowledge; they create it, encourage others to create ever more, and then they absorb it all like a sponge. As a famous contemporary author once wrote, the worst part of knowledge is that the more you know, the more you learn how much you don't know. The best part is that this realization results in the creation of knowledge exponentially. The enlightened group of authors behind this book have attempted to kick off that very process. By filling in some voids in your knowledge, they hopefully tickle your imagination and inspire you to explore more, to add to the entire gamut of knowledge as a whole.

Going back to the original discussion, what makes an expert? Is it winning some titles or passing some certification tests? Many of these authors have been honored for their excellence in some form or other. Some are formally recognized by Oracle (as Administrator Certified Expert and ACE Directors, and as an Author of the Year, for example). Others are recognized by formal certification authorities (as Oracle Certified Professionals and Oracle Certified Masters). All are recognized by the strictest body of all—the OakTable Network itself. But all these recognitions pale in comparison to the biggest accomplishment of all—their recognition by their peer group: the user community. Ultimately, what makes the difference is the acceptance of their excellence by you, the reader and the user of Oracle technology. All of these authors participate heavily in the user community, give knowledge back to the community—internationally or in their local spheres—and shape the progress of accumulation of collective knowledge. When they come together to pen their thoughts in the format of a book, the result can be nothing short of sheer pleasure.

An expert also must be well rounded, albeit is expected to be an expert in a subset of the technologies. The dissemination of knowledge is not complete without a deeper understanding of some of the supporting technologies. In Chapter 5, Niall shows how to effectively manage Oracle on Windows through a thorough explanation of the subtle nuances of Windows and how to get around them. In Chapter 7, Jože shows how to exploit a little known but powerful feature of the cost-based optimizer called the *extensible optimizer*. In Chapters 14 and 15, world-renowned security expert Pete shows how to deviate from performance aspects and bulletproof your database from security threats—a goal that attained considerable importance lately. In Chapter 10, another renowned expert named Tim takes a detour to the world of superlatives—big databases, longer load times, and so on, to educate readers in the fine art of managing beasts of databases.

Technology changes faster than you can spell it, and most of us are stuck in a loop of understanding deeply and then becoming obsolete. However, there is no substitute to solid foundation in design, as Melanie describes in Chapter 4. Changes are inevitable, and keeping up with the pace of change often poses challenges for experts in any domain. Pace of change is one of the factors leading to the pollution of knowledge, as performance-tuning guru Jonathan explains in Chapter 11 in relation to the gathering of statistics. And Jeremiah's Chapter 2 is the result of change; it explains the newest phenomenon to hit the experts—Oracle in the cloud.

As a part of my extra-curricular activities, I often travel outside my native United States to other countries to provide seminars, trainings, sessions, and to participate in discussion forums. In fact, I am writing this foreword in the middle of a two-week, four-country tour across Europe. While on the plane from Estonia to Germany, I was reflecting on the lesser-appreciated aspects of knowledge processing, the ones too subtle to notice but too important to ignore. Cultural aspects play a huge role in the collection and dissemination of knowledge. Mankind is far from being homogeneous; it's submerged in plurality—languages, cultures, customs, and expressions shaped by traditions. Ignoring this plethora of differences is not only detrimental to synergy, but also to the very essence of the knowledge-begets-knowledge belief. Not understanding cultural and customary differences can only lead to the inefficiency of the process, if not to a more dangerous reversal and creation of ignorance puddles. A vital characteristic of this cast of authors is that they represent many parts of the globe. The ideas they propound, the mannerisms they convey, are shaped and influenced—some more than the other—by the customs and traditions they uphold. A profound example is the United States, a melting pot of many cultures in which diversity gives strength to the society and is celebrated. It's diversity of the authors that gives this book an all-round perspective, something that is literally world-class.

Finally, it's the supporting cast that makes the last run of the technical aspect of the book a huge success. The authors didn't just get a free rein in whatever they wanted to write; their work was thoroughly scrutinized by a team of technical reviewers—also members of the OakTable Network and respected professionals in their own right. Even the editor—Jonathan Gennick—is an Oakie as well, making it an Oakie production end to end.

While I have been honored to write the foreword, I stoop under the weight of the huge expectation of summarizing the essence of these chapters from the stalwarts of the Oracle user community into a few pages. In conclusion, I will reiterate the message that has been delivered by many authors in many shapes or forms—excellence is not an end or even a journey. It is not about gathering as much knowledge as possible; it's about building a very strong foundation on whatever you know, expanding the horizon but not at the cost of a weakened foundation. It is about refining your foundation, and reinforcing it by constantly subjecting it to assessment and analysis, and finally providing evidence (quantifiable or otherwise), not opinions based on perceptions.

Happy reading!

Arup Nanda
Chief database architect, Starwood Hotels and Resorts
OakTable Network member
Oracle ACE Director

About the Authors

Melanie Caffrey, Pete Finnigan, Randolph Geist, Alex Gorbachev, Tim Gorman, Connie Green, Charles Hooper, Jonathan Lewis, Niall Litchfield, Karen Morton, Robyn Sands, Jože Senegačnik, Uri Shaft, Riyaj Shamsudeen, Jeremiah Wilton, Graham Wood

Melanie Caffrey

Melanie Caffrey is a senior development manager for Oracle Corporation, providing front-end and back-end Oracle solutions for the business needs of various clients. She is co-author of several technical publications, including *Oracle Web Application Programming for PL/SQL Developers*, *Oracle DBA Interactive Workbook*, and *Oracle Database Administration: The Complete Video Course*, all published by Prentice Hall. She has instructed students in Columbia University's Computer Technology and Applications program in New York City, teaching advanced Oracle database administration and PL/SQL development. She is a frequent Oracle conference speaker.

Pete Finnigan

Pete is a world-renowned expert in the area of Oracle security, providing consultancy, design expertise, security audits, and training, all in the area of Oracle security. Pete is a member of the OakTable Network. He has spoken regularly all over the world at various conferences such as those of the United Kingdom Oracle Users Group (UKOUG), PSOUG, Black Hat, and Risk. Pete is a published author on Oracle security, and researches and writes about the subject regularly. Pete also runs his website, www.petefinnigan.com, dedicated to Oracle security and providing help and free tools to the community.

Randolf Geist

Randolf Geist has been working with Oracle software for 15 years now. Since 2000 he has operated as a freelance database consultant focusing primarily on performance-related issues, and in particular helping people to understand and unleash the power of the Oracle cost-based optimizer (CBO). He is writing on his blog about CBO-related issues and is also regularly contributing to the official Oracle Technology Network (OTN) forums. Randolf is a member of the OakTable Network, the Oracle ACE program, and is an Oracle Certified Professional DBA for Oracle Versions 8i, 9i, and 10g. He also maintains SQLTools++, an open source Oracle GUI for Windows.

Alex Gorbachev

Alex Gorbachev is a respected figure in the Oracle world, and a sought-after leader and speaker at Oracle conferences around the globe. He is an OakTable Network member, and has been recognized as an Oracle ACE Director for his contributions to the community and unmatched set of skills. He is the founder of the *Battle Against Any Guess* movement, promoting scientific troubleshooting techniques. He is currently the chief technology officer at The Pythian Group.

Alex has worked for The Pythian Group in several roles. He began by leading a team of database experts in Ottawa. He then moved to Australia to build the company's presence in the East Asia Pacific region. Now he is back in Ottawa as The Pythian Group's chief technology officer. In all his work, Alex continues to work toward bridging the gap between business and technology. The search for the perfect fit between technology, engineering talents, and business process is what keeps him up at night.

Tim Gorman

Tim Gorman began his IT career in 1984 as a C programmer on Unix and VMS systems, working on medical and financial systems as an application developer, systems programmer, and systems administrator. He joined Oracle Corporation in 1990 as a consultant, became an independent consultant in 1998, and has worked for SageLogix since 2000. Gorman is the co-author of *Essential Oracle8i Data Warehousing* and *Oracle8 Data Warehousing*. He specializes in performance-tuning applications, databases, and systems, as well as data warehouse design and implementation, backup and recovery, architecture and infrastructure, and database administration. Gorman still considers himself a pretty good coder, although the market for C programs has dried up somewhat lately.

Connie Green

Connie Green has 20 years of experience with Oracle products, the last 12 years working for Oracle U.S. Server Development. Connie's expertise is in database performance tuning. She designed and developed Statspack from its initial release. Over the years Connie has been privileged to work with many talented kernel developers, performance engineers, and support engineers, having been involved with library cache development, buffer cache development, and the query optimizer.

Charles Hooper

Charles Hooper is the IT manager and an Oracle database administrator at K&M Machine Fabricating, where he has been employed since 2000. His role in the company extends well beyond Oracle database administration responsibilities, providing opportunities for database performance tuning, network administration, programming, hardware/software troubleshooting, mentoring fellow IT staff, and end-user training for the Infor ERP Visual system as well as other custom-developed applications. Charles is well known among the user community of Infor Visual Enterprise because of his years of participation in various user forums answering technical questions, providing the only Oracle presentation at Infor's 2008 user's conference, and providing presentations to regional user groups. Prior to joining K&M, he was a computer/technology instructor and Novell NetWare administrator. He holds a bachelor of science in mathematics and computer science. Much has changed in the computer industry since his first home computer in 1981, which offered 5KB of memory, a surprisingly small amount today as manufacturers offer home computers with as much as 24GB of memory.

Jonathan Lewis

Jonathan Lewis has been working in the IT industry for nearly 25 years, and has been using the Oracle RDBMS for more than 20. For the last 16 years he has been working as a freelance consultant, often spending only one or two days at a time with any client to address critical performance problems. He also advises on design and implementation problems, and on how to make best use of the most appropriate Oracle features for a given project.

Jonathan is renowned throughout the world (having visited 42 countries at last count) for his tutorials and seminars about the Oracle database engine and how to make best use of it. His exceptional ability has earned him an O-1 visa from the United States, allowing him to do consultancy and lecture work there.

Jonathan has written two books about Oracle (*Practical Oracle8i*, Addison-Wesley, 2000; *Cost-Based Oracle Fundamentals*, Apress, 2005) as well as contributing to two others (*Oracle Insights*, Apress, 2004; *Oracle Database 10g New Features*, McGraw-Hill, 2004). He also writes regularly for the UKOUG magazine, and occasionally for other publications around the world. In the limited amount of time he has left over, Jonathan also publishes high-tech Oracle articles on his blog at <http://jonathanlewis.wordpress.com>.

Niall Litchfield

Niall Litchfield is a DBA of 15 years standing, with extensive experience of running databases on various x86(64) platforms, and with a particular bias toward Microsoft Windows. Niall began professional life as a "bean-counter" for KPMG; consequently he is a DBA who respects auditors, not least because they have to stand in unheated warehouses in the middle of winter watching other people count things, something he never wants to do again. His dad was a *real* engineer crafting genuinely useful objects with fine tolerances and that, coupled with an early experience in which he showed that the two competing macro-economic theories of the time both explained actual behavior equally well, and that neither explained it adequately, rather explains his evidence-based approach to databases in general and performance tuning in particular.

Karen Morton

Karen Morton is a consultant and educator specializing in application optimization in both shoulder-to-shoulder consulting engagements and classroom settings. She is a senior principal database engineer for Agilex Technologies (<http://agilex.com>), a Chantilly, Virginia company founded by Bob LaRose and Jay Nussbaum. Agilex offers customers innovative thinking on leveraging advanced technologies within the healthcare, federal, intelligence, commercial, and public sectors to fully realize the value of information.

For over 20 years, Karen has worked in information technology, starting out as a mainframe programmer, developer, DBA, data architect, and now as a researcher, educator, and consultant. Having used Oracle since the early 90s, she began teaching others how to use Oracle over a decade ago.

She is a frequent speaker at conferences and user groups, an Oracle ACE, and a member of the OakTable network. She blogs at <http://karenmorton.blogspot.com>.

Robyn Sands

Robyn Anderson Sands is a software engineer for Cisco Systems. In a previous incarnation, she worked in industrial engineering, manufacturing development, and quality engineering at Lockheed Martin, supporting the P-7, F-22, and C-130J programs. Robyn has been working with databases and Oracle software since around 1996. She began her work with Oracle by developing tools for scheduling, composite fabrication capacity modeling, and engineering workflow, and progressing to the implementation and administration of data warehouse, PeopleSoft, and SAP systems. Current projects include “architecting” and managing the development of embedded database systems for Cisco customers, and searching for new ways to design and develop database systems with consistent performance and minimal maintenance requirements. She has been a speaker at UKOUG, Miracle conferences, Oracle OpenWorld, and the Hotsos Symposium. She occasionally posts random blog entries at adhdoddba.blogspot.com.

Jože Senegačnik

Jože Senegačnik has more than 20 years of experience in working with Oracle products. He began in 1988 with Oracle Database version 4 while working for the City of Ljubljana, Slovenia, where he had charge over the city’s municipal and geographic information systems. From 1993 to 2003, he worked in developing GIS systems for the Surveying and Mapping Authority of the Republic of Slovenia, and in the development of applications for other governmental institutions, all based on the Oracle database. More recently, he has specialized in performance optimization, having developed his own toolset for monitoring performance and analyzing trace files.

Jože is an internationally recognized speaker, and a member of the highly respected OakTable Network (oaktable.net). He is a regular speaker at user-group conferences, especially those put on by the Slovenian Oracle Users Group (SIOUG), the Independent Oracle Users Group (IOUG), and the United Kingdom Oracle Users Group (UKOUG). He also speaks routinely at the Hotsos Symposium and Oracle OpenWorld. In addition to sharing his knowledge through conference talks, Jože conducts technical seminars organized either by Oracle University or himself. He was awarded Oracle ACE membership for his long record of positive contributions to the Oracle community.

Uri Shaft

Uri Shaft received a Ph.D. from the University of Wisconsin-Madison, specializing in database systems. He was the lead developer of the innovative QUIQ hybrid database and information retrieval system. Currently, Uri works for Oracle USA in the server manageability group. He is the developer in charge of the Automatic Database Diagnostic Monitor (ADDM), and of components that analyze the Active Session History (ASH) sampled data. Uri's areas of expertise include multimedia support in database systems, multidimensional indexing theory (and practice), and database performance diagnosis and tuning.

Riyaj Shamsudeen

Riyaj Shamsudeen is the principal DBA and president of OraInternals (www.orainternals.com), a performance/recovery/EBS11i consulting company. He specializes in RAC, performance tuning, and database internals. He also frequently blogs about these technology areas in his blog, <http://orainternals.wordpress.com>. He is a regular presenter in many international conferences such as Hotsos, COLLABORATE, RMOUG, SIOUG, and UKOUG. He is a proud member of OakTable network. He has more than 16 years of experience using Oracle technology products and more than 15 years as an Oracle DBA/Oracle Applications DBA.

Jeremiah Wilton

Jeremiah Wilton has worked with Oracle technology since 1994. His main claim to fame is having been Amazon.com's first database administrator, back in the pre-IPO days. For seven years, he helped Amazon.com survive exponential scaling and a wide variety of nearly catastrophic technology failures. Jeremiah owned and ran ORA-600 Consulting for a number of years, until it was acquired by Blue Gecko, a global provider of remote administration for Oracle, MySQL, and E-Business Suite. Jeremiah also teaches the Oracle certificate program for the University of Washington. Jeremiah is an Oracle Certified Master, a member of the OakTable, and a frequent presenter at industry conferences and user groups. His publications and whitepapers can be found at www.bluegecko.net.

Graham Wood

Graham Wood is an architect in the database development group at Oracle. Most of his 20 years at Oracle have been spent in performance-related areas, including designing and tuning large high-performance systems, building monitoring tools such as Statspack, and in architecting performance and scalability features into the database engine itself. More recently Graham was the architect of the Oracle Manageability team tasked with simplifying the process of tuning the operation of the database, which resulted in the development of AWR, ASH, and ADDM to provide automatic tuning.

About the Technical Reviewers

Melanie Caffrey

Melanie Caffrey is a senior development manager for Oracle Corporation, providing front-end and back-end Oracle solutions for the business needs of various clients. She is co-author of several technical publications, including *Oracle Web Application Programming for PL/SQL Developers*, the *Oracle DBA Interactive Workbook*, and *Oracle Database Administration: The Complete Video Course*, all published by Prentice Hall. She has instructed students in Columbia University's Computer Technology and Applications program in New York City, teaching advanced Oracle database administration and PL/SQL development. She is a frequent Oracle conference speaker.

Arup Nanda

Arup Nanda has been an Oracle DBA for more than 16 years—touching all aspects of database management and architecture—from modeling to performance tuning and disaster recovery. He has written more than 300 articles, co-authored 4 books, spoken at 150 technical conferences, and delivered a number of full-length training sessions. In 2003, he was awarded DBA of the Year by Oracle. He is an Oracle Certified Professional DBA, an OTN ACE Director, and a member of the OakTable Network. He lives in Connecticut with his wife Anu and son Anish.

Peter Sharman

Peter Sharman is a curriculum developer with Server Technologies Curriculum Development at Oracle Corporation. He has over 20 years of Oracle experience, and has been a speaker at several Oracle OpenWorld, Miracle Database Forum, and RMOUG Training Days conferences. Pete has multiple OCP and OCM certifications, and is currently responsible for developing training on the Oracle Enterprise Manager product family.