



# Contents

ACKNOWLEDGMENTS .....	xv
INTRODUCTION .....	xvii
<b>I The Right Approach to Building Applications .....</b>	<b>I</b>
It's a Team Effort .....	2
DBA and Developer Roles .....	3
Read the Documentation .....	6
A Guide to the Guides .....	7
Road Maps to Reading .....	10
Avoid the Black Box Syndrome .....	11
Database Independence versus Database Dependence .....	11
Dangers of Black Box Syndrome .....	12
It's a Database, Not a Data Dump .....	23
Use Primary and Foreign Keys .....	24
Test the Overhead of Referential Integrity .....	24
Middle Tier Checking Is Not a Panacea .....	26
Build a Test Environment .....	29
Test Against Representative Data .....	30
Don't Test with a Single User .....	33
Don't Test in a Dust-Free Lab .....	33
Design to Perform; Don't Tune to Perform .....	34
Don't Use Generic Data Models .....	34
Design Your Data Model for Efficiency .....	37
Define Your Performance Goals from the Start .....	40
Work to Clear, Specific Metrics .....	41
Collect and Log Metrics over Time .....	41
Don't Do It Because "Everyone Knows You Should" .....	42
Benchmark, Benchmark, Benchmark .....	43
Small-Time Benchmarking .....	44
Big-Time Benchmarking .....	47
Instrument the System .....	49
Trace from asktom.oracle.com .....	50
Instrument for Remote Debugging .....	51

Use DBMS_APPLICATION_INFO Everywhere	52
Use DEBUG.F in PL/SQL	53
Turn on SQL_TRACE in Your Application	54
Use Industry-Standard APIs	55
Build Your Own Routines	56
Audit Is Not a Four-Letter Word	56
Question Authority	57
Beware of Universal “Bests”	58
Suspect Ratios and Other Myths	58
Don’t Look for Shortcuts	60
Keep It Simple	61
Consider Alternate Approaches	62
Let the Database Do What It Does Best	63
Use Supplied Functionality	66
We Heard Feature X Is Slow	66
We Heard Feature X Is Complicated	69
We Don’t Want to	69
We Didn’t Know	71
We Want Database Independence	71
Summary	73
<b>2 Your Performance Toolkit</b>	<b>75</b>
SQL*Plus	76
Set Up SQL*Plus	77
Customize the SQL*Plus Environment	78
Read the Documentation!	80
EXPLAIN PLAN	80
Setup for EXPLAIN PLAN	81
Use EXPLAIN PLAN	82
How to Read a Query Plan	84
Avoid the EXPLAIN PLAN Trap	87
Use DBMS_XPLAN and V\$SQL_PLAN	90
AUTOTRACE	92
Set Up AUTOTRACE	92
Use AUTOTRACE	93
Format the AUTOTRACE Output	93
Understand the AUTOTRACE Output	95
What Are You Looking for in AUTOTRACE Output?	96
TKPROF	120
Enable TKPROF	121
Run TKPROF	122
Read a TKPROF Report	124
TKPROF for the Masses	128
Runstats	135
Set Up Runstats	136
Use Runstats	140

Statspack	144
Set Up Statspack	145
Use Statspack	146
What People Do Wrong with Statspack	147
Statspack at a Glance	148
DBMS_PROFILER	153
Why You Want to Use the Profiler	153
Profiler Resources	156
JDeveloper (and Debugging)	156
Summary	160
<b>3 Architectural Decisions</b>	<b>161</b>
Understand Shared Server vs. Dedicated Server Connections	162
How Do Dedicated Server Connections Work?	163
How Do Shared Server Connections Work?	165
Common Misconceptions about Shared Server Connections	169
Dedicated Server vs. Shared Server Wrap-Up	170
Take Advantage of Clustering	171
How Does RAC Work?	172
What Are the Benefits of RAC?	176
Clustering Wrap-Up	177
Know When to Use Partitioning	178
Partitioning Concepts	178
The Partitioning Myth	180
Why Use Partitioning?	185
Partitioning Wrap-Up	187
Know When to Use Parallel Operations	187
The Parallel Myth	188
Parallel Administration	191
Parallel Query	193
Parallel DML	195
DIY Parallelism	196
Parallel Processing Wrap-Up	199
Summary	199
<b>4 Effective Administration</b>	<b>201</b>
Use SPFILES to Start Your Database	202
The Problems with PFILES	202
How SPFILES Work	203
Convert a Database to Use SPFILE	203
Save System Parameter Changes	204
Are PFILES Obsolete?	204
Help, My SPFILE Is Broken and I Cannot Start	205
SPFILE Wrap-Up	207
Let Oracle Manage Your Datafiles	208
When Is OMF Useful?	208

How OMF Works .....	209
OMF Wrap-Up .....	211
Bulletproof Your Recovery .....	212
Backup Guidelines .....	212
Backup and Recovery Wrap-Up .....	216
Use Locally Managed Tablespaces .....	216
Why Are DMTs Obsolete? .....	216
Use System-Managed LMTs When You Do Not Know How Big Your Objects Will Become .....	217
Use Uniform Extent Sizes When You Know the Ultimate Size of an Object ..	220
Some LMT Caveats .....	221
LMT and DMT Wrap-Up .....	226
Let Oracle Manage Your Segment Space .....	226
Understand Freelists and Freelist Groups .....	227
How PCTFREE and PCTUSED Control Freelists .....	231
The Case for ASSM .....	232
ASSM Wrap-Up .....	234
Let Oracle Manage Your Rollback Segments .....	235
Set UNDO_RETENTION .....	236
An UNDO Tablespace Caveat .....	238
UNDO Tablespace Wrap-Up .....	239
Summary .....	239
<b>5 Statement Processing .....</b>	<b>241</b>
Understand the Types of SQL Statements .....	242
How Are Statements Executed? .....	242
Parsing .....	243
Optimization and Row-Source Generation .....	249
Execution .....	251
Statement Execution Wrap-Up .....	252
Queries from Start to Finish .....	252
A Quick-Return Query .....	253
A Slow-Return Query .....	255
Consistent Reads .....	257
Modification DML from Start to Finish .....	261
DDL Processing .....	262
Use Bind Variables .....	264
What Are the Advantages of Using Bind Variables? .....	265
Use Bind Variables with Java and VB .....	273
There Are Exceptions to Every Rule .....	279
Bind Variable Peeking .....	283
Parse as Little as Possible .....	287
The Cost of Parsing .....	287
Use PL/SQL to Reduce Parses .....	290
Move SQL Out of Triggers to Reduce Parsing .....	298
Prepare Once; Execute Many .....	301
Summary .....	302

<b>6 Getting the Most Out of the Cost-Based Optimizer</b>	<b>303</b>
Why the RBO Is Dead	304
Make the CBO Do Its Best	307
Adjust the OPTIMIZER_INDEX_CACHING and OPTIMIZER_INDEX_COST_ADJ Parameters	307
Use SYSTEM Statistics	311
Optimize the CBO	320
Set COMPATIBLE for Upgrades	320
Set DB_FILE_MULTIBLOCK_READ_COUNT to Reduce Full-Scan Costs	321
Set HASH_JOIN_ENABLED to Control Hash Joins	326
Set OPTIMIZER_DYNAMIC_SAMPLING to Gather Statistics Dynamically	326
Set OPTIMIZER_FEATURES_ENABLE to Control Feature Choices	333
Set OPTIMIZER_MAX_PERMUTATIONS to Control Permutations	335
Set OPTIMIZER_MODE to Pick a Mode	338
Rewrite Queries with QUERY_REWRITE_ENABLED and QUERY_REWRITE_INTEGRITY	348
Control PGA Memory with BITMAP_MERGE_AREA_SIZE, SORT_AREA_SIZE, and HASH_AREA_SIZE	355
Use STAR_TRANSFORMATION_ENABLED for Star Queries	359
Set Others Parameters that Affect the Optimizer	360
Use the 10053 Event to Trace CBO Choices	361
Summary	365
<b>7 Effective Schema Design</b>	<b>367</b>
Fundamental Schema Design Principles	368
Let the Database Enforce Data Integrity	368
Use the Correct Datatype	374
Optimize to Your Most Frequently Asked Questions	379
Overview of Table Types	380
B*Tree Index Clustered Tables	382
Create Clusters	383
Use Clusters	386
Clusters Wrap-Up	397
Index-Organized Tables (IOTs)	398
Use IOTs as a Space-Saving Alternative to Association Tables	398
Use IOTs to Colocate Randomly Inserted Data	400
IOTs Wrap-Up	404
External Tables	404
Set Up External Tables	405
Modify External Tables	408
Use External Tables for Direct-Path Loading	409
Use External Tables for Parallel Direct-Path Loading	411
Use External Tables for Merging	411
Handle Errors with External Tables	413
Indexing Techniques	415
Use FBIs—Think Outside the Box	416
Use Domain Indexes	422

Compression	425
Use Index Key Compression	426
Use Table Compression for Read-Only/Read-Mostly Tables	432
Compression Wrap-Up	442
Summary	442
<b>8 Effective SQL</b>	<b>445</b>
What You Need to Write Efficient SQL	446
Understand Access Paths	447
Full Scans	447
ROWID Access	454
Index Scans	456
Cluster Scans	465
Understand Joins	465
Nested Loops	465
Hash Joins	468
Sort-Merge Joins	471
Cartesian Joins	473
Anti-Joins	477
Full Outer Joins	481
Schema Matters (Physical)	484
Really Know SQL	487
The ROWNUM Pseudo Column	490
Scalar Subqueries	504
Analytics	514
Don't Tune a Query	535
Understand the Question	535
A Proof-of-Concept Example	536
Overview of Other SQL Techniques	540
Summary	541
<b>9 Effective PL/SQL Programming</b>	<b>543</b>
Why PL/SQL?	544
PL/SQL Is the Most Efficient Language for Data Manipulation	544
PL/SQL Is Portable and Reusable	546
Write as Little as You Can	548
Not Doing It Procedurally	551
Fit Your Code on the Screen	552
Use Packages	553
Advantages of Packages	553
Break the Dependency Chain	553
Packages Wrap-Up	559
Use Static SQL	559
Advantages of Static SQL	559
Look for Opportunities to Replace Dynamic SQL	560
Static SQL Wrap-Up	562
Bulk Processing	563
Use Bulk Processing When It Has Dramatic Effects	563

Use Bulk Processing for ETL Operations	565
Bulk Processing Wrap-Up	571
Returning Data	571
Advantages of Ref Cursors	572
Use Ref Cursors to Return Result Sets	572
Use %TYPE and %ROWTYPE	577
Base Record Types on a Table	577
Base Record Types on a Cursor	581
Base Datatypes on a Column	583
Using Invoker Rights	585
Invoker Rights and Multiple Schemas	585
Criteria for Invoker Rights Routines	586
Make Your Lookups Work Efficiently	588
Single-Row Fetching for Lookups	590
Bulk Processing for Lookups	593
Single-Statement Operations for Lookups	595
Lookup Wrap-Up	596
Be Careful with Autonomous Transactions	597
Criteria for Autonomous Transactions	597
Autonomous Transactions Can Impact Data Integrity	598
Choose Whether to Use Implicit or Explicit Cursors	600
Use Implicit Cursors for Single-Row Selections	601
Use Implicit Cursors for Result Sets with a Limited Number of Rows	607
Implicit/Explicit Cursors Wrap-Up	609
Summary	609
<b>10 So, You Had an Accident</b>	<b>611</b>
Find Out What's Different	612
Start Collecting History Today	615
Detective Work	616
Change One Thing at a Time	617
Have a Sound Reason for Changing that One Thing	618
Have a Goal	618
Validate Your Hypothesis	619
Be Able to Unchange Things	622
Build a Test Case	622
Test Case Requirements	622
Keep Your Test Cases as Small as Possible	623
Summary	624
<b>A Setting Up and Some Scripts</b>	<b>627</b>
Setting Up BIG_TABLE	628
Frequently Used Scripts	630
PRINT_TABLE	630
SHOW_SPACE	634
COLS_AS_ROWS	638
GEN_DATA	645
Index	649