

Contents

Foreword	xi
About the Author	xiv
About the Technical Reviewers	xv
Acknowledgments	xvi
Introduction	xvii
Setting Up Your Environment	xxv

CHAPTER 1	Developing Successful Oracle Applications	1
	My Approach	2
	The Black Box Approach	4
	How (and How Not) to Develop Database Applications	9
	Understanding Oracle Architecture	9
	Understanding Concurrency Control	15
	Multi-Versioning	20
	Database Independence?	26
	“How Do I Make It Run Faster?”	42
	The DBA–Developer Relationship	46
	Summary	47
CHAPTER 2	Architecture Overview	49
	Defining Database and Instance	50
	The SGA and Background Processes	55
	Connecting to Oracle	57
	Dedicated Server	57
	Shared Server	59
	Mechanics of Connecting over TCP/IP	60
	Summary	62

CHAPTER 3	Files	65
	Parameter Files	66
	What Are Parameters?	67
	Legacy init.ora Parameter Files	69
	Server Parameter Files (SPFILES)	71
	Parameter File Wrap-Up	78
	Trace Files	78
	Requested Trace Files	79
	Trace Files Generated in Response to Internal Errors	83
	Trace File Wrap-Up	85
	Alert File	85
	Data Files	88
	A Brief Review of File System Mechanisms	89
	The Storage Hierarchy in an Oracle Database	90
	Dictionary-Managed and Locally-Managed Tablespaces	94
	Temp Files	96
	Control Files	98
	Redo Log Files	98
	Online Redo Log	99
	Archived Redo Log	101
	Password Files	103
	Change Tracking File	106
	Flashback Log Files	107
	Flashback Database	107
	Flash Recovery Area	108
	DMP Files (EXP/IMP Files)	108
	Data Pump Files	110
	Flat Files	113
	Summary	114
CHAPTER 4	Memory Structures	115
	The Process Global Area and User Global Area	115
	Manual PGA Memory Management	116
	Automatic PGA Memory Management	123
	Choosing Between Manual and Auto Memory Management	133
	PGA and UGA Wrap-Up	135
	The System Global Area	135
	Fixed SGA	139
	Redo Buffer	140

Block Buffer Cache	141
Shared Pool	148
Large Pool	150
Java Pool	151
Streams Pool	152
Automatic SGA Memory Management	152
Summary	154

■ CHAPTER 5 Oracle Processes 155

Server Processes	156
Dedicated Server Connections	156
Shared Server Connections	158
Connections vs. Sessions	159
Dedicated Server vs. Shared Server	165
Dedicated/Shared Server Wrap-Up	169
Background Processes	170
Focused Background Processes	171
Utility Background Processes	178
Slave Processes	181
I/O Slaves	181
Parallel Query Slaves	182
Summary	182

■ CHAPTER 6 Locking and Latching 183

What Are Locks?	183
Locking Issues	186
Lost Updates	186
Pessimistic Locking	187
Optimistic Locking	189
Optimistic or Pessimistic Locking?	200
Blocking	200
Deadlocks	203
Lock Escalation	208
Lock Types	209
DML Locks	209
DDL Locks	217
Latches	220
Manual Locking and User-Defined Locks	229
Summary	230

CHAPTER 7	Concurrency and Multi-Versioning	231
	What Are Concurrency Controls?	231
	Transaction Isolation Levels	232
	READ UNCOMMITTED	234
	READ COMMITTED	235
	REPEATABLE READ	237
	SERIALIZABLE	239
	READ ONLY	241
	Implications of Multi-Version Read Consistency	242
	A Common Data Warehousing Technique That Fails	242
	An Explanation for Higher Than Expected I/O on Hot Tables	244
	Write Consistency	246
	Consistent Reads and Current Reads	247
	Seeing a Restart	249
	Why Is a Restart Important to Us?	252
	Summary	253
 CHAPTER 8	 Transactions	 255
	Transaction Control Statements	256
	Atomicity	257
	Statement-Level Atomicity	257
	Procedure-Level Atomicity	259
	Transaction-Level Atomicity	262
	Integrity Constraints and Transactions	262
	IMMEDIATE Constraints	262
	DEFERRABLE Constraints and Cascading Updates	263
	Bad Transaction Habits	265
	Committing in a Loop	266
	Using Autocommit	272
	Distributed Transactions	273
	Autonomous Transactions	275
	How Autonomous Transactions Work	275
	When to Use Autonomous Transactions	277
	Summary	281

CHAPTER 9	Redo and Undo	283
	What Is Redo?	283
	What Is Undo?	284
	How Redo and Undo Work Together	287
	Example INSERT-UPDATE-DELETE Scenario	287
	Commit and Rollback Processing	291
	What Does a COMMIT Do?	292
	What Does a ROLLBACK Do?	298
	Investigating Redo	300
	Measuring Redo	300
	Redo Generation and BEFORE/AFTER Triggers	302
	Can I Turn Off Redo Log Generation?	308
	Why Can't I Allocate a New Log?	313
	Block Cleanout	314
	Log Contention	317
	Temporary Tables and Redo/Undo	319
	Investigating Undo	323
	What Generates the Most and Least Undo?	323
	ORA-01555: snapshot too old Error	325
	Summary	336
CHAPTER 10	Database Tables	337
	Types of Tables	337
	Terminology	339
	Segment	339
	Segment Space Management	341
	High-Water Mark	342
	FREELISTS	344
	PCTFREE and PCTUSED	347
	LOGGING and NOLOGGING	350
	INITRANS and MAXTRANS	351
	Heap Organized Tables	351
	Index Organized Tables	354
	Index Organized Tables Wrap-Up	369
	Index Clustered Tables	370
	Index Clustered Tables Wrap-Up	378
	Hash Clustered Tables	378
	Hash Clustered Tables Wrap-Up	387
	Sorted Hash Clustered Tables	388

Nested Tables	390
Nested Tables Syntax	391
Nested Table Storage	399
Nested Tables Wrap-Up	402
Temporary Tables	403
Temporary Tables Wrap-Up	410
Object Tables	410
Object Tables Wrap-Up	418
Summary	418

CHAPTER 11 Indexes	421
An Overview of Oracle Indexes	422
B*Tree Indexes	423
Index Key Compression	426
Reverse Key Indexes	429
Descending Indexes	435
When Should You Use a B*Tree Index?	437
B*Trees Wrap-Up	447
Bitmap Indexes	448
When Should You Use a Bitmap Index?	449
Bitmap Join Indexes	453
Bitmap Indexes Wrap-Up	455
Function-Based Indexes	455
Important Implementation Details	455
A Simple Function-Based Index Example	456
Indexing Only Some of the Rows	464
Implementing Selective Uniqueness	466
Caveat on CASE	467
Caveat Regarding ORA-01743	468
Function-Based Indexes Wrap-Up	469
Application Domain Indexes	469
Frequently Asked Questions and Myths About Indexes	471
Do Indexes Work on Views?	471
Do Nulls and Indexes Work Together?	471
Should Foreign Keys Be Indexed?	474
Why Isn't My Index Getting Used?	475
Myth: Space Is Never Reused in an Index	482
Myth: Most Discriminating Elements Should Be First	485
Summary	488

CHAPTER 12 Datatypes	489
An Overview of Oracle Datatypes	489
Character and Binary String Types	492
NLS Overview	492
Character Strings	495
Binary Strings: RAW Types	502
Number Types	504
NUMBER Type Syntax and Usage	507
BINARY_FLOAT/BINARY_DOUBLE Type Syntax and Usage	510
Non-Native Number Types	511
Performance Considerations	511
LONG Types	513
Restrictions on LONG and LONG RAW Types	513
Coping with Legacy LONG Types	515
DATE, TIMESTAMP, and INTERVAL Types	520
Formats	521
DATE Type	522
TIMESTAMP Type	529
INTERVAL Type	537
LOB Types	540
Internal LOBs	541
BFILES	553
ROWID/UROWID Types	555
Summary	556
 CHAPTER 13 Partitioning	 557
Partitioning Overview	558
Increased Availability	558
Reduced Administrative Burden	560
Enhanced Statement Performance	565
Table Partitioning Schemes	567
Range Partitioning	567
Hash Partitioning	570
List Partitioning	575
Composite Partitioning	577
Row Movement	579
Table Partitioning Schemes Wrap-Up	581

Partitioning Indexes	582
Local Indexes vs. Global Indexes	583
Local Indexes	584
Global Indexes	590
Partitioning and Performance, Revisited	606
Auditing and Segment Space Compression	612
Summary	614

CHAPTER 14 Parallel Execution

When to Use Parallel Execution	616
A Parallel Processing Analogy	617
Parallel Query	618
Parallel DML	624
Parallel DDL	627
Parallel DDL and Data Loading Using External Tables	628
Parallel DDL and Extent Trimming	630
Parallel Recovery	639
Procedural Parallelism	639
Parallel Pipelined Functions	640
Do-It-Yourself Parallelism	643
Summary	648

CHAPTER 15 Data Loading and Unloading

SQL*Loader	649
Loading Data with SQLLDR FAQs	653
SQLLDR Caveats	679
SQLLDR Summary	680
External Tables	680
Setting Up External Tables	681
Dealing with Errors	687
Using an External Table to Load Different Files	690
Multiuser Issues	691
External Tables Summary	692
Flat File Unload	692
Data Pump Unload	701
Summary	703

INDEX