

# Contents at a Glance

About the Author

About the Technical Reviewers

CHAPTER 0.	Introduction.....	1
CHAPTER 1.	Methods and Madness .....	11
CHAPTER 2.	Listening to Oracle's Pain.....	35
CHAPTER 3.	Serialization Control.....	67
CHAPTER 4.	Identifying and Understanding Operating System Contention .....	101
CHAPTER 5.	Oracle Performance Diagnosis.....	135
CHAPTER 6.	Oracle Buffer Cache Internals.....	185
CHAPTER 7.	Oracle Shared Pool Internals.....	245
CHAPTER 8.	Oracle Redo Management Internals.....	291
CHAPTER 9.	Oracle Performance Analysis .....	327

# Contents

## About the Author

## About the Technical Reviewers

<b>CHAPTER 0. Introduction</b> .....	1
Why Buy This Book? .....	2
What Is the Value to Me?.....	3
Who Will Benefit? .....	3
How This Book Is Organized .....	4
What Notations Are Used? .....	5
What Is Not Covered? .....	7
Which Platform and Version? .....	8
About the Tools Used In This Book.....	8
Want More? .....	9
Comments and Questions .....	9
<b>CHAPTER 1. Methods and Madness</b> .....	11
Firefighting 101 .....	12
Stoking the Fire.....	12
Cooling the Fire, Step by Step.....	13
Don't Panic .....	13
Define Your Objectives.....	14
Establish the Scope and Get Reinforcements.....	14
Get a Baseline of Current Performance .....	14
Install Your Tools .....	15
Develop a Simple Communications Strategy .....	15
Pick the Low-Hanging Fruit .....	16
Conduct a Deep Performance Analysis .....	16
Get a Final Baseline of Current Performance.....	17
Document Your Success.....	17
Celebrate! .....	17
OraPub 3-Circle Analysis.....	18
A Couple 3-Circle Analysis Case Studies .....	20
Case Study 1: Quick 3-Circle Analysis.....	20
Case Study 2: More Complex 3-Circle Analysis .....	21
Your Compelling Story .....	22
The Story Components.....	23
The Story Development Process .....	24
Wait-Event Analysis .....	25

# CONTENTS

Oracle Response-Time Analysis .....	27
The Role of the Response-Time Curve .....	27
Case Study of Oracle Response-Time Analysis .....	29
Red Line, Blue Line .....	31
Summary .....	33
<b>CHAPTER 2. Listening to Oracle's Pain .....</b>	<b>35</b>
Performance Diagnosis: The Backstory .....	36
It's All About Instrumentation .....	37
Oracle Instrumentation .....	41
How Oracle Collects Time .....	42
The Wait Event Views .....	45
System-Level Perspective (v\$system_event) .....	45
Session-Level Perspective (v\$session_event) .....	47
Real-Time Session-Level Perspective (v\$session_wait) .....	49
Oracle Time Classification .....	52
Queue Time Classification .....	54
Service Time Classification .....	55
OraPub's Response-Time Analysis Reports .....	56
Instance-Level ORTA Reporting .....	56
Part 1: Workload Metrics .....	57
Part 2: Response Time Summary .....	57
Part 3: IO Wait Time Summary with Event Details .....	61
Part 4: Other Wait Time (Non-IO) Event Details .....	61
Part 5: SQL Activity Details During Probe .....	62
Part 6: Similar SQL Statements .....	62
Part 7: Operating System CPU Utilization .....	62
Session-Level ORTA Reporting .....	62
Profiling a Single Session .....	63
Profiling a Group of Sessions .....	65
Summary .....	65
<b>CHAPTER 3. Serialization Control .....</b>	<b>67</b>
Serialization Is Death .....	68
Serialization and Queuing .....	68
Everyone Gets Involved .....	69
How to Detect and Resolve Contention .....	70
Fundamental Protection Requirements .....	71
Relational Structure Control .....	71
Memory Structure Control .....	72
Oracle Latch Specifics .....	74
How Multiple Latches Are Implemented .....	75
Least Recently Used Lists .....	75
Cache Buffer Chains .....	76
Oracle's General Latching Algorithm .....	78

Shared or Exclusive? .....	80
Immediate or Willing to Wait? .....	80
Spinning on a Latch .....	80
Sleep Time .....	81
Time Accounting .....	83
A Real-Life Latching Acquisition Example .....	84
Should You Increase the <code>_spin_count</code> Parameter? .....	86
How to Detect Significant Latch Contention .....	86
Oracle Mutex Specifics .....	91
What Is a Mutex? .....	91
Benefits of Using Mutexes .....	91
Flexible Creation .....	92
Reduced False Contention .....	92
Control Structure Contention .....	93
Faster Pinning .....	93
Oracle's General Mutex Algorithm .....	94
How to Detect Mutex Contention .....	98
Summary .....	100

## **CHAPTER 4. Identifying and Understanding Operating System Contention**..... 101

The Four Subsystems .....	102
CPU Contention .....	102
How to Model a CPU Subsystem .....	103
Where CPU Time Is Spent .....	104
When Queuing Sets In .....	105
Queue Length Strategies .....	109
OLTP-Centric Systems .....	109
Batch-Centric Systems .....	110
Monitoring CPU Activity .....	110
Utilization .....	110
Run Queue .....	116
Memory Pressure .....	118
Memory Categories .....	118
Real and Virtual Memory .....	119
Shared Memory Segments .....	119
Process-Related Memory .....	119
The General Memory Game .....	120
Swap: A Four-Letter Word .....	120
Memory Page Scanning .....	122
What to Say and What Not to Say About Memory .....	122
IO Contention .....	122
Load-Balancing Still Helps .....	123
Why IO Subsystems Are Expensive .....	123
How We Model an IO Subsystem .....	123
Even the Best IO Subsystems Queue .....	124

## CONTENTS

How to Detect an IO Bottleneck .....	125
Using Oracle's Wait Interface .....	126
Removing Oracle from the Equation .....	128
Using Operating System IO Reports .....	129
Understanding Your Solution Options.....	130
Network Contention.....	132
Network Latency.....	132
Network Collisions.....	133
Dropped Packets.....	134
Summary.....	134

## **CHAPTER 5. Oracle Performance Diagnosis** .....

135

Oracle CPU Consumption and Components .....	135
Oracle's Limited Perspective .....	136
Using Instance Statistics .....	137
Using the System Time Model.....	142
Time Model Superiority .....	142
Time Model Time Classification.....	145
The Ghost IO Bottleneck .....	148
More Than Just an Average .....	149
Wait Event Myths .....	151
Decreasing Wait Time Always Improves Performance.....	151
Decreasing Wait Time Decreases End-to-End Response Time .....	152
End-to-End Response Time Defined.....	153
Some End-to-End Response Time Realities .....	153
The SQL*Net Message from Client Wait Event.....	154
Profiling a Session Is Always the Best Approach.....	156
Profiling Defined.....	157
The Trap .....	157
The Solution .....	158
Modern Architecture Statistics Collection .....	158
Why We Need a Better Collection Facility .....	158
Oracle's Solution: DBMS_MONITOR.....	163
It Helps to Change Our Mindset.....	163
How to Use DBMS_MONITOR.....	164
Criteria Specification: Identify the Session(s) of Interest .....	164
Enable Tracing, Statistics Collection, or Both.....	167
Wait While the Data Is Being Collected .....	168
Query the Appropriate Statistics Collection View .....	168
Disable Tracing, Statistics Collection, or Both.....	169
Consolidate the Trace Files into a Single File .....	169
Tkprof the Trace Files.....	169
Perform Your Analysis.....	169
A DBMS_MONITOR Example .....	169
Active Session History .....	176
Why ASH Is a Big Deal .....	177
A Demonstration of ASH Capabilities .....	177

ASH Data Collection and Architecture .....	182
Summary .....	184
<b>CHAPTER 6. Oracle Buffer Cache Internals .....</b>	<b>185</b>
Big Expectations .....	186
What Is a Buffer? .....	187
Free Buffers .....	187
Dirty Buffers .....	188
Pinned Buffers .....	189
The Role of Buffer Headers .....	189
Cache Buffer Chains .....	191
Introduction to Hashing .....	191
Hash Functions .....	191
Hash Buckets .....	193
Hash Chains .....	194
CBCs in Action .....	194
How to Wreck CBC Performance .....	195
Limiting Concurrency by Decreasing Latches .....	195
Increasing Chain Scan Time by Decreasing Chains .....	198
Increasing Chain Scan Time with Cloned Buffers .....	199
CBC Contention Identification and Resolution .....	202
Least Recently Used Chains .....	203
LRU Chain Changes Over the Years .....	204
Standard LRU Algorithm .....	205
Modified LRU Algorithm .....	206
Oracle's Touch-Count Algorithm .....	207
Midpoint Insertion .....	208
Touch Count Incrementation .....	209
Buffer Promotion .....	210
Hot Region to Cold Region Movement .....	211
About Touch Count Changes .....	212
LRU Chain Contention Identification and Resolution .....	213
The Write List and Database Writer .....	215
The Database Writer in Action .....	216
Database Writer-Related Contention Identification and Resolution .....	218
Free Buffer Waits .....	220
Buffer Busy Waits .....	222
The Four-Step Diagnosis .....	223
Determining If There Is a Parameter Pattern .....	223
Identifying the Buffer Type .....	224
Determining the Header Block .....	226
Implementing the Appropriate Solution Set .....	227
Solutions for a Single Busy Table Block .....	227
Solutions for Multiple Busy Table Blocks .....	228
Solutions for Table Segment Header Blocks .....	229
Solutions for Undo Segment Header Blocks .....	230
Solutions for Index Leaf Blocks .....	231

## CONTENTS

The Situation .....	231
The Solution .....	232
The Good News and the Bad News .....	232
Enqueue Waits .....	233
Diagnosing Enqueue Waits .....	233
TX Enqueue Waits.....	235
Introduction to Interested Transaction Lists.....	236
Undo Segment's Transaction Table .....	237
Deeper into Interested Transaction Lists .....	238
Deeper into Buffer Cloning.....	241
Summary.....	243

## **CHAPTER 7. Oracle Shared Pool Internals .....** 245

Problems in the Shared Pool.....	246
What's in the Shared Pool? .....	247
The Oracle Cursor.....	248
Parent and Child Cursors.....	248
Cursor Building .....	249
Cursor Searching Introduction .....	250
Cursor Pinning and Locking.....	250
Library Cache Management .....	251
Library Cache Architecture .....	251
Library Cache Conceptual Model .....	251
Library Cache Object References .....	252
Object Hashing .....	257
Keeping Cursors in the Cache.....	259
Increase the Likelihood of Caching .....	259
Force Caching.....	260
Private Cursor Caches .....	261
Library Cache Latch/Mutex Contention Identification and Resolution .....	262
Enable Mutexes .....	262
Use Bind Variables to Create Similar SQL .....	263
Use Cursor Sharing .....	264
Take Advantage of the Hash Structure .....	265
Try Mutex-Focused Solutions .....	266
Shared Pool Memory Management .....	268
From Hashing to Subpools .....	268
Memory Allocation and Deallocation.....	270
Shared Pool Latch Contention Identification and Resolution .....	271
Pin Large and Frequently Used Objects .....	271
Flush the Shared Pool.....	275
Increase the Number of Subpools .....	275
Reduce the Shared Pool Size.....	275
4031 Error Resolution .....	276
Flush the Shared Pool.....	276
Increase the Shared Pool Size .....	276
Increase the Shared Pool Reserved Size.....	277

Minimize Cursor Pinning Duration .....	277
Reduce Kept Objects Memory Consumption.....	278
Upgrade to Oracle Database 10g Release 2 .....	278
In-Memory Undo Management.....	278
New Features Bring High Risk.....	279
The Problem: Segment Management .....	279
Introducing In-Memory Undo .....	280
How IMU Works .....	281
Traditional Undo Management.....	281
IMU Management .....	283
A Marked Performance Improvement .....	284
IMU Setup and Monitoring .....	285
Setting Up IMU.....	285
Monitoring IMU.....	286
IMU Contention Identification and Resolution .....	288
Summary.....	289

## **CHAPTER 8. Oracle Redo Management Internals .....** 291

Buffer Cache Changes.....	292
Just Enough Redo Is Generated .....	292
Undo-Related Redo .....	293
Query-Related Redo.....	293
Redo Log Buffer Architecture and Algorithm.....	295
Pre-Oracle9i Release 2 Redo Log Buffer .....	296
Post-Oracle9i Release 2 Redo Log Buffer.....	297
Redo Flow .....	299
Global Temporary Tables .....	300
The Need for True Interim Tables .....	300
Common Characteristics .....	300
Truly Reduced Redo.....	301
Log Writer Background Process Triggers.....	301
Commit Issued .....	302
Commit Write Facility.....	302
Database Writer Posting the Log Writer .....	305
Buffer Fill .....	305
Three-Second Timeout.....	306
Redo-Related Performance Issues .....	307
Log Buffer Space .....	308
Redo Allocation Latch Contention .....	309
Redo Copy Latch Contention.....	310
Log File Sync Contention .....	310
Application-Focused Solutions for Log File Sync Contention.....	311
Operating System-Focused Solutions for Log File Sync Contention .....	314
Oracle-Focused Solutions for Log File Sync Contention .....	314
Log File Parallel Write Contention .....	315
Log Writer Write Challenges .....	315
Gathering Oracle's IO Requirements .....	319

## CONTENTS

Application-Focused Solutions for Log File Parallel Write Contention ..	321
Operating System-Focused Solutions for Log File Parallel Write Contention ..	323
Oracle-Focused Solutions for Log File Parallel Write Contention ..	324
Log File Switch Contention ..	325
Checkpoint Incomplete ..	325
Archive Incomplete ..	326
Summary ..	326

## **CHAPTER 9. Oracle Performance Analysis**..... 327

Deeper into Response-Time Analysis ..	328
Oracle Arrival Rates ..	328
Utilization ..	331
Requirements Defined ..	332
Gathering CPU Requirements ..	332
Gathering IO Requirements ..	334
Capacity Defined ..	336
Gathering CPU Capacity ..	336
Gathering IO Capacity ..	337
Calculating Utilization ..	338
Oracle CPU Utilization ..	338
Operating System CPU Utilization ..	338
IO Utilization ..	339
Oracle Service Time ..	340
Oracle Queue Time ..	341
Oracle Response Time ..	343
The Bridge Between Firefighting and Predictive Analysis ..	343
Total Time and Time Per Workload ..	344
CPU Service and Queue Time ..	344
IO Service and Queue Time ..	345
Oracle Response Time in Reality ..	346
Response-Time Graph Construction ..	348
Selecting the Unit of Work ..	349
Choosing the Level of Abstraction ..	349
The Five-Step Response-Time Graph Creation Process ..	351
Know the System Bottleneck ..	351
Pick an Appropriate Unit of Work ..	352
Determine the Service Time and Queue Time ..	352
If Possible, Compare Utilizations ..	353
Create the Response-Time Graph ..	353
A Response-Time Curve for an IO-Bottlenecked System ..	354
Know the System Bottleneck ..	354
Pick an Appropriate Unit of Work ..	356
Determine the Service Time and Queue Time ..	356
If Possible, Compare Utilizations ..	357
Create the Response-Time Graph ..	357
How to Improve the Performance Situation ..	358

Tuning: Reducing Requirements .....	358
Buying: Increasing Capacity .....	359
Balance: Managing Workload .....	361
Anticipating Solution Impact .....	363
Simplification Is the Key to Understanding .....	363
A Word of Caution .....	364
Full Performance Analysis: Cycle 1 .....	364
Oracle Analysis .....	364
Calculating Averages .....	365
Reducing Queue Time .....	366
Reducing Service Time .....	366
Operating System Analysis .....	366
CPU Subsystem .....	367
IO Subsystem .....	367
Application Subsystem .....	368
Response-Time Graphs .....	369
What-If Analysis .....	370
Full Performance Analysis: Cycle 2 .....	372
Oracle Analysis .....	372
Operating System Analysis .....	372
Application Analysis .....	373
What-If Analysis .....	375
Full Performance Analysis: Cycle 3 .....	376
Oracle Analysis .....	377
Operating System Analysis .....	377
Application Analysis .....	378
Full Performance Analysis: Summary .....	379
Improper Index Impact Analysis Performed .....	381
Proper Use of Work and Time .....	381
Batch Process-Focused Performance Analysis .....	382
Setting Up the Analysis .....	382
Capturing Resource Consumption .....	383
Capturing Client Process Time .....	383
Removing Background Process Time .....	383
Including Parallelism in the Analysis .....	383
Anticipating Elapsed Time .....	384
Scalability Issues .....	384
Operating in the Elbow of the Curve .....	386
Summary and Next Steps .....	386