



Contents at a Glance

1 Introduction to Oracle Wait Interface	1
2 Oracle Wait Interface Components	15
3 Common Wait Events	49
4 OWI Monitoring and Collection Methods	77
5 Interpreting Common I/O Related Wait Events	103
6 Interpreting Locks-Related Wait Events	141
7 Interpreting Common Latency-Related Wait Events	197
8 Wait Events in a Real Application Clusters Environment	215
9 Performance Management in Oracle Database 10g	237
A Oracle Database 10g Diagnostic Events	275
B Enqueue Waits in Oracle Database 10g	289
C Oracle Dumps and Traces	301
D Direct SGA Access	317
E References	333
Index	339



Contents

FOREWORD	xvii
ACKNOWLEDGMENTS	xxi
INTRODUCTION	xxv
1 Introduction to Oracle Wait Interface	1
The Old Fashion of Oracle Performance Optimization	3
Why Are Cache-Hit Ratios Grossly Inefficient?	3
The New Fashion of Oracle Performance Optimization	5
The OWI Philosophy	6
Database Response Time Tuning Model	8
Paradigm Shift	11
In a Nutshell	13
2 Oracle Wait Interface Components	15
What Is a Wait Event?	16
OWI Components	17
V\$EVENT_NAME View	18
V\$SYSTEM_EVENT View	20
V\$SESSION_EVENT View	24
V\$SESSION_WAIT View	27
Trace Event 10046—The Extended SQL Trace	30
New OWI Views in Oracle Database 10g Release 1	36
V\$SESSION_WAIT_HISTORY View	37
V\$SYSTEM_WAIT_CLASS View	38
V\$SESSION_WAIT_CLASS View	39
V\$EVENT_HISTOGRAM View	40
Types of Wait Events	42
OWI Limitations	45
No CPU Statistics	45
No End-to-End Visibility	45

No Historical Data	46
Inaccuracies	46
In a Nutshell	48
3 Common Wait Events	49
Introduction to Common Wait Events	50
buffer busy waits	50
control file parallel write	51
db file parallel read	53
db file parallel write	53
db file scattered read	54
db file sequential read	55
db file single write	55
direct path read	56
direct path write	57
enqueue	57
free buffer waits	59
latch free	59
library cache pin	60
library cache lock	61
log buffer space	62
log file parallel write	62
log file sequential read	63
log file switch (archiving needed)	63
log file switch (checkpoint incomplete)	63
log file switch completion	64
log file sync	64
SQL*Net message from client	65
SQL*Net message to client	65
Common Wait Events in Oracle Real Application	
Clusters Environment	66
global cache cr request	66
buffer busy global cache	68
buffer busy global cr	68
global cache busy	69
global cache null to x	69
global cache null to s	70
global cache s to x	70
global cache open x	71
global cache open s	71
row cache lock	71
Tracking CPU and Other Statistics	72
In a Nutshell	74

4 OWI Monitoring and Collection Methods	77
Why Is Historical Performance Data Important?	78
Fast and Accurate Root Cause Analysis	79
Why Trace Event 10046 Is Not a Suitable Data Collector	80
Why Statspack Is Not a Suitable Data Collector	81
Database Logoff Trigger as a Data Collector	82
Sampling for Performance Data Using PL/SQL Procedure	86
Data Source	87
Sampling Frequency	88
Repositories	90
Events to Monitor	92
Pros and Cons	100
Sampling for Performance Data with SQL-less SGA Access	101
In a Nutshell	101
5 Interpreting Common I/O Related Wait Events	103
db file sequential read	104
Common Causes, Diagnosis, and Actions	105
db file scattered read	112
Common Causes, Diagnosis, and Actions	113
Why Does db file sequential read event Show Up in a Full Table Scan Operation	116
Why Does a Full Scan Operation Request Fewer Blocks than the MBRC	118
Setting the DB_FILE_MULTIBLOCK_READ_COUNT (MBRC) ..	118
Why Physical I/Os Are Expensive	120
direct path read	120
Common Causes, Diagnosis, and Actions	121
Initialization Parameters of Interest	126
direct path write	128
Common Causes, Diagnosis, and Actions	128
db file parallel write	130
Common Causes, Diagnosis, and Actions	130
log file parallel write	135
Common Causes, Diagnosis, and Actions	135
control file parallel write	138
Common Causes, Diagnosis, and Actions	138
In a Nutshell	140
6 Interpreting Locks-Related Wait Events	141
Latch Free	142
What Is a Latch?	142
Differences between a Latch and a Lock	143

Latch Family	144
Latch Acquisition	145
Latch Classification	146
What Does the Latch Free Wait Event Tell You?	148
Latch Miss Locations	149
Latches in Oracle Database 10g Release 1	149
Common Causes, Diagnosis, and Actions	150
Shared Pool and Library Cache Latches	151
Cache Buffers Chains Latches	159
Cache Buffers Lru Chain Latches	167
Row Cache Objects Latches	169
Enqueue	170
What Is an Enqueue?	170
What Is an Enqueue Resource?	170
What Is an Enqueue Lock?	172
Enqueue Architecture	174
Decoding Enqueue Type and Mode	174
Common Causes, Diagnosis, and Actions	177
Buffer Busy Waits	184
Common Causes, Diagnosis, and Actions	185
In a Nutshell	195
7 Interpreting Common Latency-Related Wait Events	197
log file sync	198
Common Causes, Diagnosis, and Actions	198
log buffer space	204
Common Causes, Diagnosis, and Actions	205
free buffer waits	206
Common Causes, Diagnosis, and Actions	206
write complete waits	210
Common Causes, Diagnosis, and Actions	210
log file switch completion	211
Common Causes, Diagnosis, and Actions	211
log file switch (checkpoint incomplete)	212
Common Causes, Diagnosis, and Actions	212
In a Nutshell	213
8 Wait Events in a Real Application Clusters Environment	215
What's So Special About Waits in Real Application Clusters (RAC)?	216
Global Buffer Cache in Real Application Clusters	217
Global Cache Waits	223
global cache cr request	223
global cache busy	228

RAC Wait Event Enhancements in Oracle Database 10g	230
Enqueue Waits	231
Most Common Enqueues	232
In a Nutshell	236
9 Performance Management in Oracle Database 10g	237
Database Statistics	238
Time Model Statistics	238
Wait Model Statistics	240
Operating System Statistics	240
Additional SQL Statistics	241
Database Metrics	241
New Background Processes	241
Automatic Workload Repository	242
Repository Snapshots	243
Snapshot Baselines	244
Using EM to Manage AWR	244
Manually Managing AWR	247
Active Session History	252
What Is an Active Session?	253
Components of ASH	253
Automatic Database Diagnostic Monitor (ADDM)	259
ADDM Setup	260
Using EM to Access ADDM	261
Manually Running ADDM Report	270
ADDM Views	273
In a Nutshell	273
A Oracle Database 10g Diagnostic Events	275
Oracle Diagnostic Events	276
Types of Diagnostic Events	276
Setting Diagnostic Events	279
Internal Workings of the Events	285
B Enqueue Waits in Oracle Database 10g	289
C Oracle Dumps and Traces	301
Oradebug: The Ultimate Utility for Traces and Dumps	302
Data Block Dump	304
Syntax	304
Buffers Dump	305
Syntax	306
Controlling the Dump Information	306

Buffer Dump	307
Syntax	307
Controlling the Dump Information Using Levels	308
File Headers Dump	308
Syntax	308
Controlling the Dump Information Using Levels	308
Control File Dump	309
Syntax	310
Controlling the Dump Information Using Levels	310
Heap Dump	310
Syntax	311
Controlling the Dump Information Using Levels	311
Library Cache Dump	312
Syntax	312
Controlling the Dump Information Using Levels	312
Processstate Dump	313
Syntax	313
Shared Server State Dump	314
Syntax	314
Controlling the Dump Information Using Levels	314
Systemstate Dump	315
Syntax	315
Controlling the Dump Information Using Levels	315
Redo Log Dump	315
Syntax	316
Controlling the Dump Information Using Levels	316
D Direct SGA Access	317
Overhead	318
Security	318
Speed	318
Concurrency	319
Get Hidden Information	319
Introduction to X\$ Views	319
The Necessary Ingredients	322
Find SGA ID	322
Find SGA Base Address	323
Find the Starting Address of X\$KSUSECST	324
Find the Record Size of the X\$KSUSECST Structure	324
Find Number of Records in the X\$KSUSECST Structure	325
Find the X\$KSUSECST View Columns Offsets	325
Attaching the SGA to a C Program	327

E References	333
Chapter 1	334
Chapter 2	334
Chapter 3	334
Chapter 4	335
Chapter 5	335
Chapter 6	335
Chapter 7	336
Chapter 8	336
Chapter 9	336
Appendix A	337
Appendix B	337
Appendix C	337
Appendix D	337
Index	339