



# Contents

|                       |      |
|-----------------------|------|
| ACKNOWLEDGMENTS ..... | xvii |
| INTRODUCTION .....    | xix  |

## PART I

### Overview of the Method and the Tools

|  |           |
|--|-----------|
| <b>I Overview of Oracle Tuning .....</b>     | <b>3</b>  |
| The Overall Tuning Approach .....            | 4         |
| Server Tuning .....                          | 9         |
| Network Tuning .....                         | 10        |
| Disk Tuning .....                            | 11        |
| Instance Tuning .....                        | 12        |
| Object Tuning .....                          | 13        |
| SQL Tuning .....                             | 14        |
| Tuning with Oracle Parallel Query .....      | 15        |
| Tuning Real Application Clusters (RAC) ..... | 15        |
| Oracle and STATSPACK .....                   | 16        |
| Conclusion .....                             | 19        |
| <b>2 Overview of STATSPACK .....</b>         | <b>21</b> |
| The STATSPACK Architecture .....             | 23        |
| How STATSPACK Collects Data .....            | 23        |
| The STATSPACK Table Structures .....         | 26        |
| STATSPACK Control Tables .....               | 27        |
| STATSPACK Parameter Tables .....             | 29        |
| Uses for STATSPACK Information .....         | 31        |
| Database Tuning with STATSPACK .....         | 31        |
| Resource Planning .....                      | 32        |
| Predictive Modeling .....                    | 33        |
| Conclusion .....                             | 34        |

|          |   |           |
|----------|---|-----------|
| <b>3</b> | <b>Installing and Configuring STATSPACK</b>           | <b>35</b> |
|          | Overview of the STATSPACK Scripts                     | 36        |
|          | STATSPACK Scripts for Oracle8 and Oracle8i            | 37        |
|          | STATSPACK Scripts for Post-8.1.6 STATSPACK            | 38        |
|          | Step 1: Create the perfstat Tablespace                | 40        |
|          | Step 2: Run the Create Scripts                        | 41        |
|          | Run the Pre-8.1.7 Install Scripts                     | 41        |
|          | Install Prerequisites                                 | 41        |
|          | Step 3: Test the STATSPACK Install                    | 43        |
|          | Step 4: Schedule Automatic STATSPACK Data Collections | 43        |
|          | STATSPACK Configuration and Maintenance               | 45        |
|          | Viewing STATSPACK Snapshots                           | 45        |
|          | Adjusting the STATSPACK Collection Thresholds         | 47        |
|          | Snapshot Levels                                       | 47        |
|          | Snapshot Thresholds                                   | 48        |
|          | Back-Porting STATSPACK for Oracle 8.0 Through 8.1.5   | 50        |
|          | Removing Old STATSPACK Snapshots                      | 50        |
|          | Using the STATSPACK purge Utility                     | 50        |
|          | Manually Removing STATSPACK Snapshots                 | 51        |
|          | Removing Ranges of Snapshots                          | 51        |
|          | An Intelligent UNIX Script to Purge Snapshots         | 54        |
|          | Handy STATSPACK Shell Scripts                         | 59        |
|          | A Quick Elapsed-Time STATSPACK Script                 | 59        |
|          | Monitoring STATSPACK Table Extents                    | 65        |
|          | Conclusion  | 67        |
| <b>4</b> | <b>Data Inside the STATSPACK Tables</b>               | <b>69</b> |
|          | What Is Missing from STATSPACK?                       | 70        |
|          | STATSPACK Subordinate Table Structures                | 70        |
|          | STATSPACK Summary Tables                              | 71        |
|          | Changes in STATSPACK Tables for Oracle9i              | 71        |
|          | The stats\$latch_misses_summary Table                 | 74        |
|          | The stats\$sgastat_summary Table (Oracle8i Only)      | 75        |
|          | The stats\$sql_summary Table                          | 77        |
|          | The stats\$parameter Table                            | 79        |
|          | STATSPACK System Tables                               | 80        |
|          | The stats\$rollstat Table                             | 80        |
|          | The stats\$latch Table                                | 82        |
|          | The stats\$latch_children Table                       | 83        |
|          | The stats\$librarycache Table                         | 83        |
|          | The stats\$waitstat Table                             | 84        |
|          | The stats\$enqueuestat Table                          | 85        |

|   |     |
|---|-----|
| The stats\$sysstat Table .....                                  | 87  |
| The stats\$sesstat Table .....                                  | 89  |
| The stats\$sgastat Table .....                                  | 90  |
| STATSPACK Transaction Tables .....                              | 91  |
| The stats\$buffer_pool Table (Pre-Oracle9i Only) .....          | 91  |
| The stats\$buffer_pool_statistics Table .....                   | 92  |
| The stats\$filestatxs Table .....                               | 94  |
| STATSPACK Event Tables .....                                    | 95  |
| The stats\$system_event Table .....                             | 96  |
| The stats\$session_event Table .....                            | 98  |
| The stats\$bg_event_summary Table .....                         | 99  |
| The stats\$idle_event Table .....                               | 100 |
| Oracle Parallel Server Tables (Real Application Clusters) ..... | 100 |
| The stats\$rowcache_summary Table .....                         | 101 |
| The stats\$sgaxs Table .....                                    | 101 |
| Conclusion .....  | 102 |

## PART II

### Tuning the Oracle Database with STATSPACK

|  |            |
|--|------------|
| <b>5 Extending STATSPACK to Collect Server Statistics .....</b>                            | <b>105</b> |
| Overview of the vmstat Utility .....   | 106        |
| Dialect Differences in vmstat .....  | 107        |
| What to Look for in vmstat .....   | 108        |
| Identifying CPU Bottlenecks with vmstat .....  | 109        |
| Identifying High CPU Usage with vmstat .....   | 112        |
| Identifying RAM Memory Bottlenecks .....   | 112        |
| Understanding UNIX RAM Memory Paging .....   | 112        |
| Capturing Server Performance Data Inside STATSPACK .....                                   | 119        |
| A Script to Capture vmstat Information .....   | 119        |
| Internals of the vmstat Capture Script .....   | 123        |
| Reporting vmstat Information on Other Oracle Servers .....                                 | 124        |
| Conclusion .....   | 125        |
| <b>6 Tuning the Server Environment .....</b>   | <b>127</b> |
| The Relationship Between the Database Administrator<br>and the Systems Administrator ..... | 129        |
| Online Server Monitor Tools .....  | 129        |
| Using glance .....   | 129        |
| Using top to Monitor the Server .....  | 130        |
| Using sar to Monitor Server Statistics .....   | 132        |
| Monitoring Server CPU Consumption .....  | 135        |

|   |            |
|---|------------|
| Upgrading an Entire Server .....                            | 136        |
| Adding Additional CPU Processors .....                      | 137        |
| Load Balancing of Server Tasks .....                        | 137        |
| Using nice and priocntl to Change Execution Priority .....  | 141        |
| Monitoring Server Memory Consumption .....                  | 143        |
| Server Memory Settings .....                                | 144        |
| Very Large Memory and Oracle .....                          | 144        |
| Making Oracle Memory Nonswappable .....                     | 145        |
| Reporting on Server Statistics .....                        | 146        |
| Server Exception Reports .....                              | 146        |
| Daily vmstat Trend Reports .....                            | 150        |
| Long-Term Server Analysis and Trending .....                | 154        |
| Conclusion .....  | 154        |
| <b>7 Tuning the Network Environment .....</b>               | <b>155</b> |
| Optimizing Oracle NET Configuration .....                   | 156        |
| The tcp.nodelay Parameter in protocol.ora .....             | 157        |
| The automatic_ipc Parameter of sqlnet.ora .....             | 158        |
| SDU and TDU Parameters in tnsnames.ora .....                | 158        |
| The queuesize Parameter in listener.ora .....               | 160        |
| The break_poll_skip Parameter of sqlnet.ora .....           | 160        |
| The disable_oob Parameter of sqlnet.ora .....               | 161        |
| The epc_disabled Environment Variable .....                 | 161        |
| Other Oracle Features that Affect Network Behavior .....    | 162        |
| Using Array Fetches to Improve Network Throughput .....     | 162        |
| Using the Multi-Threaded Server .....                       | 162        |
| Connection Pooling and Network Performance .....            | 164        |
| ODBC and Network Performance .....                          | 166        |
| Tuning with Oracle Replication .....                        | 167        |
| Monitoring Network Performance from Oracle STATSPACK .....  | 169        |
| Tuning the Distributed Network .....                        | 172        |
| Using netstat to Monitor Network Activity .....             | 172        |
| Conclusion .....  | 174        |
| <b>8 Tuning the Disk I/O Subsystem with STATSPACK .....</b> | <b>175</b> |
| Oracle Tuning Factors that Influence Disk I/O .....         | 177        |
| Oracle Internals and Disk I/O .....                         | 178        |
| Oracle File Organization Techniques .....                   | 178        |
| Transient Disk Hot Spots .....                              | 180        |
| Mapping Oracle Disk Architectures .....                     | 186        |
| The Multiple RAM Buffer Issue .....                         | 187        |
| File Striping with Oracle .....                             | 188        |
| Using RAID with Oracle .....                                | 189        |

|   |     |
|---|-----|
| Using Oracle with Raw Devices .....                 | 190 |
| Load Balancing Disks with Oracle Databases .....    | 191 |
| Configuring Oracle Tablespaces and Datafiles .....  | 191 |
| Building the Oracle File-to-Disk Architecture ..... | 194 |
| Reporting on the Oracle Disk Architecture .....     | 195 |
| STATSPACK Reports for Oracle Datafiles .....        | 196 |
| Detailed Disk and File I/O with STATSPACK .....     | 203 |
| A STATSPACK Report on Specific I/O Activity .....   | 204 |
| A STATSPACK Script to Identify Hot Datafiles .....  | 206 |
| The Approach to Locating Hot Disks .....            | 207 |
| Extending STATSPACK for Disk I/O Data .....         | 208 |
| The Basic iostat Utility .....                      | 208 |
| Defining the STATSPACK Table .....                  | 210 |
| Capturing the iostat Information .....              | 211 |
| Generating iostat Reports .....                     | 213 |
| Viewing I/O Signatures with STATSPACK .....         | 217 |
| Conclusion .....                                    | 222 |

### PART III

## **Tuning the Oracle Database with STATSPACK**

|   |            |
|---|------------|
| <b>9 Tuning the Oracle Database Instance .....</b>              | <b>225</b> |
| An Overview of the Oracle Database Instance .....               | 227        |
| The Oracle SGA .....  | 228        |
| Oracle Background Processes .....                               | 232        |
| Blocksize and Oracle Disk I/O .....                             | 235        |
| The db_file_multiblock_read_count and Oracle .....              | 236        |
| Tuning the Oracle7 through Oracle8i Data Buffers .....          | 236        |
| Introduction of Data Block Caching .....                        | 237        |
| Full Data Caching .....   | 239        |
| The Data Buffer Hit Ratio .....                                 | 240        |
| Data Buffer Pool Internals .....                                | 242        |
| Monitoring Data Buffer Pool Usage with STATSPACK .....          | 245        |
| Overview of the Oracle8i Data Pools .....                       | 248        |
| Locating Tables and Indexes for the KEEP Pool .....             | 249        |
| Tuning the RECYCLE Pool .....                                   | 256        |
| Identifying Candidates for the RECYCLE Pool .....               | 257        |
| Trend Reports of the Data Buffer Hit Ratio with STATSPACK ..... | 261        |
| Tuning the Oracle9i Data Buffer Pools .....                     | 267        |
| The Seven Data Buffer Hit Ratios .....                          | 267        |
| Tuning the Oracle8.0 Database Writer Processes .....            | 274        |
| Oracle 8.0 Parameters .....                                     | 275        |

|   |            |
|---|------------|
| Monitoring Database Writer Contention in Oracle8i and Oracle9i . . . . .        | 282        |
| Tuning the Shared Pool . . . . .  | 285        |
| Tuning the Library Cache . . . . .  | 292        |
| Monitoring the Library Cache Miss Ratio . . . . .                               | 293        |
| Monitoring Objects Within the Library Cache<br>with STATSPACK . . . . .         | 294        |
| Tuning the Dictionary Cache . . . . .   | 299        |
| Tuning Oracle Sorting . . . . .   | 305        |
| Tuning the Undo Records (Rollback Segments) . . . . .                           | 312        |
| Monitoring Dedicated Connections to Oracle . . . . .                            | 316        |
| UNIX Interaction with the Multi-Threaded Server . . . . .                       | 319        |
| Prerequisites for Using the MTS . . . . .                                       | 320        |
| Oracle9i Dynamic RAM and UNIX . . . . .   | 329        |
| Oracle9i and UNIX Granules . . . . .  | 331        |
| Changing Dynamic SGA and PGA Components . . . . .                               | 333        |
| Oracle9i PGA Memory Allocation for Dedicated Connections . . . . .              | 334        |
| Automatic RAM Memory Management in Oracle9i . . . . .                           | 334        |
| Moving Toward a Self-Tuning Oracle9i Database . . . . .                         | 342        |
| Conclusion . . . . .  | 351        |
| <b>10 Tuning Oracle Tables and Indexes . . . . .</b>                            | <b>353</b> |
| Automatic Space Management in Oracle9i . . . . .                                | 354        |
| The Evolution of Bitmaps in Oracle . . . . .                                    | 355        |
| Oracle9i Freelists Algorithms . . . . .   | 356        |
| Characteristics of Bitmap Segment Management . . . . .                          | 357        |
| Oracle9i Freelist Internals . . . . .   | 358        |
| Oracle9i Tools for Automatic Space Management . . . . .                         | 362        |
| Traditional Oracle Storage Parameters and Performance . . . . .                 | 365        |
| The pctfree Storage Parameter . . . . .   | 365        |
| The pctused Storage Parameter . . . . .   | 366        |
| The Freelists Storage Parameter . . . . .                                       | 367        |
| The Freelist Groups Storage Parameter for OPS . . . . .                         | 367        |
| Summary of Storage Parameter Rules . . . . .                                    | 367        |
| Traditional Freelist Management and Oracle Objects . . . . .                    | 369        |
| Linking and Unlinking from the Freelists . . . . .                              | 370        |
| Reducing Freelist Relinks . . . . .   | 372        |
| Table Internals and Freelists . . . . .   | 373        |
| Long Data Columns and Freelist Behavior . . . . .                               | 375        |
| Setting <i>pctfree</i> and <i>pctused</i> Based on Average Row Length . . . . . | 377        |
| Buffer Busy Waits and Freelist Contention . . . . .                             | 378        |
| Using STATSPACK to Find Wait Contention . . . . .                               | 379        |
| Finding Buffer Busy Waits with STATSPACK . . . . .                              | 382        |
| Reorganizing Oracle Tables . . . . .  | 390        |
| Using CTAS to Reorganize a Table . . . . .                                      | 392        |

|   |            |
|---|------------|
| Identifying Oracle Tables with Chained Rows .....             | 398        |
| Identifying Tables with Long Rows .....                       | 403        |
| Identifying Sparse Tables .....                               | 404        |
| Resequencing Oracle Table Rows for High Performance .....     | 406        |
| Index Rebuilding Techniques .....                             | 409        |
| When to Rebuild Indexes .....                                 | 415        |
| Automating Index Rebuilds .....                               | 415        |
| Identifying Unused Indexes in Oracle9i .....                  | 417        |
| A Sample Index Monitoring Session .....                       | 417        |
| Tuning Index Contention with Hidden Parameters .....          | 418        |
| Monitoring Oracle Tables and Indexes with STATSPACK .....     | 419        |
| Allocating the STATSPACK Extension Tables .....               | 419        |
| Collecting the STATSPACK Snapshot for Tables and Indexes .... | 421        |
| Reports on Tables and Indexes .....                           | 424        |
| Distributing the Table Reports via E-mail .....               | 442        |
| Conclusion .....  | 443        |
| <b>II Tuning Oracle SQL .....</b>                             | <b>445</b> |
| Goals of SQL Tuning .....                                     | 446        |
| The Problem of Declarative SQL Syntax .....                   | 447        |
| The Oracle SQL Optimizers .....                               | 449        |
| optimizer_mode = RULE .....                                   | 449        |
| optimizer_mode = FIRST_ROWS .....                             | 449        |
| optimizer_mode = ALL_ROWS .....                               | 449        |
| optimizer_mode = CHOOSE .....                                 | 450        |
| Tuning with Rule-Based Optimization .....                     | 453        |
| Changing the Rule-Based Driving Table .....                   | 453        |
| When the Rule-Based Optimizer Fails to Use                    |            |
| the Correct Index .....                                       | 454        |
| Tuning with Cost-Based Optimization (CBO) .....               | 456        |
| Invoking the Cost-Based Optimizer .....                       | 458        |
| Gathering Statistics for the CBO .....                        | 458        |
| Determining the Default optimizer_mode .....                  | 460        |
| Miscellaneous Tuning Techniques .....                         | 461        |
| Tuning with CBO SQL Hints .....                               | 462        |
| Tuning SQL Subqueries .....                                   | 462        |
| The Problem of Literal SQL Statements .....                   | 463        |
| Using cursor_sharing in Oracle8i .....                        | 466        |
| Tuning SQL with Temporary Tables .....                        | 467        |
| Tuning SQL by Adding Indexes .....                            | 468        |
| General Rules for Writing Efficient SQL .....                 | 469        |

|   |            |
|---|------------|
| The SQL Tuning Process .....  | 470        |
| Step 1: Identify High-Impact SQL in the Library Cache .....                 | 471        |
| Step 2: Extract and Explain the SQL Statement .....                         | 481        |
| Step 3: Tune the SQL Statement .....  | 487        |
| An Actual Case-Study in SQL Tuning .....                                    | 488        |
| Get the Full Table Scan Report .....  | 488        |
| Advanced SQL Execution Plan Analysis .....                                  | 490        |
| Table Histograms .....  | 494        |
| Making Permanent Changes to Tuned SQL .....                                 | 495        |
| Using Stored Outlines for SQL .....   | 495        |
| Using the v\$sql_plan and the v\$sql_workarea Views .....                   | 496        |
| Conclusion .....  | 499        |
| <b>I2 Tuning with Oracle Parallel Features .....</b>                        | <b>501</b> |
| Using Oracle Parallel Query .....   | 503        |
| Parallel Query Parameters .....   | 504        |
| Setting the Optimal Degree of Parallelism .....                             | 505        |
| Using Parallel Query Hints .....  | 509        |
| Monitoring Oracle Parallel Query .....                                      | 511        |
| Monitoring Oracle Parallel Query with STATSPACK .....                       | 511        |
| Monitoring Oracle Parallel Query with v\$ Views .....                       | 512        |
| Parallel Queries and Distributed Objects .....                              | 513        |
| Finding Candidate Tables for Oracle Parallel Query .....                    | 514        |
| Using Parallel DML .....  | 516        |
| Turning On Parallel DML .....   | 516        |
| Parallelizing Oracle Table Reorganizations .....                            | 517        |
| Parallel Index Rebuilding .....   | 518        |
| Conclusion .....  | 519        |
| <b>I3 Tuning the Oracle Parallel Server Environment .....</b>               | <b>521</b> |
| Introduction to Oracle Cluster Server Architecture .....                    | 522        |
| Partitioning Data for RAC .....   | 525        |
| The Integrated Distributed Lock Manager (Oracle7<br>Through Oracle8i) ..... | 526        |
| Configuring the IDLM within the Oracle8i OPS Environment ...                | 527        |
| Oracle Table Settings for OPS .....   | 529        |
| Tuning the OPS Environment .....  | 532        |
| STATSPACK Tables for Monitoring OPS .....                                   | 534        |
| The stats\$rowcache_summary Table .....                                     | 534        |
| The stats\$sgaxs Table .....  | 535        |
| The stats\$sysstat Table .....  | 535        |
| Extending STATSPACK for OPS Information .....                               | 537        |
| Querying the v\$ Views for Oracle Parallel Server .....                     | 540        |



|   |     |
|---|-----|
| Comparing Real Application Clusters to Oracle Parallel Server ..... | 543 |
| Monitoring Transaction Application Failover in Oracle9i .....       | 545 |
| Conclusion .....  | 546 |

## PART PART IV

### Database Reporting with STATSPACK

|  |            |
|--|------------|
| <b>14 Monitoring Oracle with STATSPACK .....</b>                     | <b>549</b> |
| The Standard STATSPACK Report .....                                  | 550        |
| Introductory Section .....   | 550        |
| Wait Event Section .....   | 553        |
| SQL Section .....  | 554        |
| Tablespace Activity Section .....                                    | 559        |
| Overview of the Alert Scripts .....                                  | 568        |
| Customized Exception Alert Reports for the DBA .....                 | 569        |
| Daily STATSPACK Alert Report .....                                   | 570        |
| Daily Server Alert Report .....                                      | 579        |
| A Real-Time Check for Oracle Problems .....                          | 580        |
| Weekly Object Growth Report .....                                    | 582        |
| Trace Alert Report .....   | 585        |
| Web Server Alert Report .....  | 588        |
| Buffer Busy Waits Alerts .....                                       | 589        |
| A STATSPACK Reactive Report .....                                    | 591        |
| Scheduling and Customizing Oracle Alert Reports .....                | 596        |
| A Sample UNIX Crontab to Schedule Oracle Reports<br>and Alerts ..... | 596        |
| Conclusion .....   | 597        |
| <b>15 Trend Analysis with STATSPACK .....</b>                        | <b>599</b> |
| Plotting STATSPACK Data Using Microsoft Excel .....                  | 600        |
| Plotting STATSPACK Data with a Spreadsheet Chart Wizard ....         | 600        |
| Step 1: Customize the STATSPACK Report .....                         | 601        |
| Step 2: Run the Report in SQL*Plus .....                             | 603        |
| Step 3: Highlight and Copy the Results .....                         | 603        |
| Step 4: Open Excel and Paste the Data .....                          | 604        |
| Step 5: Partition the Data into Columns .....                        | 604        |
| Step 6: Column Delimit the Data .....                                | 604        |
| Step 7: Start the Chart Wizard .....                                 | 606        |
| Step 8: Choose a Line Chart .....                                    | 606        |
| Step 9: Complete the Chart Wizard and View the Chart .....           | 606        |
| Step 10: Add a Trend Line .....                                      | 608        |
| STATSPACK Reports for Forecasting .....                              | 609        |
| Basic STATSPACK Metrics for Trend Analysis .....                     | 609        |

|   |     |
|---|-----|
| STATSPACK Extensions for Database Server Trend Analysis . . . . | 610 |
| Checking Server Trends by Hour . . . . .                        | 613 |
| Plotting Server Statistics by Day of the Week . . . . .         | 616 |
| Web-Based Plotting of STATSPACK Data . . . . .                  | 617 |
| The RRDtool Product . . . . .                                   | 617 |
| Other Web-Based Graphing Tools . . . . .                        | 618 |
| The STATSPACK Viewer Product . . . . .                          | 620 |
| Conclusion . . . . .  | 625 |
| Index . . . . .   | 627 |