

Table of Contents

| | |
|---|--------------|
| List of Figures | xv |
| List of Tables | xxi |
| Foreword | xxiii |
| Preface | xxv |
| Chapter 1: Tuning for Performance | 1 |
| 1.1 WHAT IS PERFORMANCE? | 2 |
| 1.1.1 Basic Performance Metrics | 2 |
| 1.1.2 Performance Metrics | 3 |
| 1.1.3 Relevance of Benchmarks | 5 |
| 1.2 WHAT IS PERFORMANCE TUNING? | 6 |
| 1.2.1 Project Development | 7 |
| 1.2.2 Server and Hardware Configuration | 9 |
| 1.2.3 Maximizing Hardware Utilization | 11 |
| 1.2.4 Application Development | 12 |
| 1.2.5 Troubleshooting | 13 |
| 1.3 STRATEGIES FOR SUCCESS | 14 |
| 1.3.1 How to Do It Right | 15 |
| 1.3.2 How to Do It Wrong | 17 |
| 1.4 MORE PRECIOUS THAN PERFORMANCE | 19 |
| 1.4.1 Assuring Availability | 19 |
| 1.4.2 Recovery Time | 21 |
| 1.4.3 Hot Standby | 22 |

| | | |
|---|---|-----------|
| 1.4.4 | Modularity | 24 |
| 1.4.5 | Transactional Semantics | 25 |
| 1.4.6 | Recommendations | 26 |
| Chapter 2: Hardware and System Software | | 27 |
| 2.1 | SERVER SYSTEM | 27 |
| 2.1.1 | Disk Drives | 29 |
| 2.1.2 | Disk Controllers | 32 |
| 2.1.3 | Disk I/O | 33 |
| 2.1.4 | Disk Mirroring | 37 |
| 2.1.5 | Disk Striping | 38 |
| 2.1.6 | Memory | 39 |
| 2.1.7 | Central Processing Unit | 41 |
| 2.1.8 | Network | 42 |
| 2.2 | SYSTEM SOFTWARE | 44 |
| 2.2.1 | Operating System Symmetry | 45 |
| 2.2.2 | Operating System Efficiency | 45 |
| 2.3 | MULTIPROCESSORS AND BEYOND | 46 |
| 2.3.1 | Shared Memory Systems | 46 |
| 2.3.2 | Clusters | 47 |
| 2.3.3 | Massively Parallel Processing | 48 |
| Chapter 3: Sybase Product and Feature Overview | | 49 |
| 3.1 | RDBMS ENGINES | 50 |
| 3.1.1 | Sybase SQL Server | 51 |
| 3.1.2 | Sybase MPP | 56 |
| 3.1.3 | Sybase IQ | 56 |
| 3.2 | CLIENT AND SERVER LIBRARIES | 57 |
| 3.2.1 | Client Libraries and Language Preprocessors | 58 |
| 3.2.2 | Open Server | 60 |
| 3.3 | AUXILIARY SERVERS | 61 |
| 3.3.1 | Backup Server | 62 |
| 3.3.2 | Monitor Server | 64 |
| 3.3.3 | Replication Server | 64 |

| | | |
|---|---|------------|
| 3.4 | PUTTING IT ALL TOGETHER | 66 |
| 3.5 | USING OTHER PEOPLE'S SOFTWARE | 66 |
| Chapter 4: SQL Server — Form and Structure | | 69 |
| 4.1 | TASK SCHEDULING | 69 |
| 4.2 | MEMORY ALLOCATION | 74 |
| 4.3 | MULTIPROCESSOR ISSUES | 76 |
| 4.4 | NETWORK I/O | 76 |
| 4.5 | DEVICES, SEGMENTS, AND PAGES | 77 |
| 4.5.1 | Mapping Segments and Databases to Disks | 78 |
| 4.5.2 | Mapping Database Pages to Device Pages | 85 |
| 4.5.3 | Logical Page Organization | 85 |
| 4.5.4 | How SQL Server Allocates Pages | 90 |
| 4.6 | INDEXES | 92 |
| 4.6.1 | Clustered Index | 93 |
| 4.6.2 | Clustered Index and Duplicate Keys | 93 |
| 4.6.3 | Nonclustered Index | 96 |
| 4.6.4 | Fill Factor | 97 |
| 4.6.5 | Maximum Rows per Page | 98 |
| 4.6.6 | Update Statistics | 98 |
| 4.7 | TEXT AND IMAGE DATA | 100 |
| Chapter 5: SQL Server — Methods and Features | | 103 |
| 5.1 | CACHE MANAGEMENT | 103 |
| 5.1.1 | Dirty Buffers | 104 |
| 5.1.2 | Named Caches and Buffer Pools | 105 |
| 5.1.3 | Buffer Replacement Strategy | 106 |
| 5.2 | LOCK MANAGEMENT | 108 |
| 5.2.1 | Page-Level Locks | 108 |
| 5.2.2 | Table-Level Locks | 110 |
| 5.2.3 | Lock Upgrades and Demand Locks | 111 |
| 5.2.4 | Lock Escalation | 112 |
| 5.2.5 | Isolation Levels and Locks | 112 |

| | | |
|-------|--|-----|
| 7.5.2 | Cost of Communication | 208 |
| 7.5.3 | Special Considerations for PC Clients | 209 |
| 7.5.4 | Performance of the Communication Subsystem | 211 |

Chapter 8: Monitoring SQL Server 221

| | | |
|------------|-----------------------------------|------------|
| 8.1 | ANALYSIS AND MONITORING | 221 |
| 8.1.1 | Analysis | 222 |
| 8.1.2 | Solving a Performance Problem | 224 |
| 8.2 | QUERYING ABOUT PERFORMANCE | 226 |
| 8.2.1 | Using sp_monitor | 226 |
| 8.2.2 | sp_sysmon | 228 |
| 8.2.3 | Using sp_who | 228 |
| 8.2.4 | Using sp_lock | 231 |
| 8.2.5 | Useful Global Variables | 234 |
| 8.3 | OPERATING SYSTEM TOOLS | 235 |
| 8.3.1 | CPU Utilization | 236 |
| 8.3.2 | Disk Activity | 240 |
| 8.3.3 | Memory Use | 241 |
| 8.3.4 | Network Statistics | 242 |
| 8.4 | SQL MONITOR | 244 |
| 8.4.1 | Main Window | 246 |
| 8.4.2 | Cache Window | 248 |
| 8.4.3 | Device I/O | 250 |
| 8.4.4 | Performance Summary | 251 |
| 8.4.5 | Performance Trends | 253 |
| 8.4.6 | Process Windows | 254 |
| 8.4.7 | Transaction Activity | 257 |

Chapter 9: Instrumenting SQL Code 259

| | | |
|------------|------------------------------|------------|
| 9.1 | SHOWPLAN | 259 |
| 9.1.1 | 4.9 and System 10 SQL Server | 260 |
| 9.1.2 | System 11 SQL Server | 262 |
| 9.1.3 | Select Query | 265 |
| 9.1.4 | Join Query | 267 |
| 9.1.5 | Subquery | 272 |

| | | |
|---|--|------------|
| 9.1.6 | Correlated Subquery | 272 |
| 9.1.7 | Update Statement | 275 |
| 9.1.8 | Stored Procedure | 281 |
| 9.1.9 | Update with Trigger | 283 |
| 9.1.10 | Constraints | 286 |
| 9.1.11 | Cursor | 287 |
| 9.2 | SET STATISTICS COMMANDS | 291 |
| 9.2.1 | Time Statistics | 291 |
| 9.2.2 | IO Statistics | 294 |
| 9.2.3 | Subquery Cache Statistics | 296 |
| 9.3 | STORED PROCEDURES | 297 |
| Chapter 10: Transaction Processing Performance | | 301 |
| 10.1 | CONFIGURING FOR OLTP | 303 |
| 10.1.1 | Checkpoint and Recovery | 303 |
| 10.1.2 | Memory Management | 305 |
| 10.1.3 | Memory Management with System 11 SQL Server | 310 |
| 10.1.4 | Other Server Configuration Parameters | 313 |
| 10.2 | DATABASE DESIGN FOR OLTP | 317 |
| 10.2.1 | Constraints and Triggers | 317 |
| 10.2.2 | Designing Tables for Transaction Processing | 324 |
| 10.2.3 | Designing Indexes for Transaction Processing | 329 |
| 10.2.4 | Logging and Group Commits | 338 |
| 10.3 | APPLICATION DESIGN | 343 |
| 10.3.1 | Managing a Large Number of Users | 343 |
| 10.3.2 | Access Methods for Transaction Processing | 347 |
| 10.3.3 | Distributed Transactions | 353 |
| 10.4 | LOCKING AND CONCURRENCY | 355 |
| 10.4.1 | Granularity of Locks | 355 |
| 10.4.2 | Eliminating Single Points of Contention | 358 |
| 10.4.3 | Duration of Locks and Isolation Levels | 359 |
| 10.4.4 | Other Issues in Locking | 365 |

| | |
|--|------------|
| Chapter 11: Query Processing Performance | 367 |
| 11.1 MYTH VERSUS REALITY | 368 |
| 11.1.1 Creating a Stored Procedure Generates the Plan | 368 |
| 11.1.2 Query Processing is Read-Only | 371 |
| 11.1.3 Plan Depends on the Order of Tables or Clauses | 372 |
| 11.1.4 Plan Depends on Schema, Statistics, and Phrasing | 373 |
| 11.1.5 There is One Optimal Plan | 374 |
| 11.2 HOW THE OPTIMIZER VIEWS QUERIES | 375 |
| 11.2.1 An Overview of the Process | 376 |
| 11.2.2 Statistics and Cost Estimation | 380 |
| 11.2.3 Accuracy of Estimates | 385 |
| 11.2.4 Worktables | 389 |
| 11.3 INDEX SELECTION | 391 |
| 11.3.1 Clustered and Nonclustered Index Scans | 393 |
| 11.3.2 Joins and Composite Indexes | 394 |
| 11.3.3 Like Optimization | 395 |
| 11.3.4 OR and IN Optimization Strategy | 396 |
| 11.3.5 Cursors and Index Selection | 398 |
| 11.3.6 MAX and MIN Strategies | 399 |
| 11.3.7 Index Covering | 400 |
| 11.4 JOIN ORDER SELECTION | 401 |
| 11.4.1 Join Processing Strategies | 402 |
| 11.4.2 Join Order Heuristics | 406 |
| 11.4.3 Join Order Enumeration | 410 |
| 11.5 SUBQUERIES AND OUTER JOINS | 414 |
| 11.5.1 Introduction to Subqueries | 414 |
| 11.5.2 Subquery Processing: 4.9 and System 10 SQL Server | 417 |
| 11.5.3 Subquery Processing: System 11 SQL Server | 422 |
| 11.5.4 Outer Joins | 428 |
| 11.6 VIEWS AND AGGREGATES | 429 |
| 11.6.1 View Materialization | 429 |
| 11.6.2 Common Subexpression | 429 |
| 11.6.3 Scalar Aggregates | 430 |
| 11.6.4 Group By and Having | 431 |
| 11.6.5 Order By and Distinct | 432 |

| | | |
|---|--|------------|
| 11.7 | OVERRIDING OPTIMIZER DECISIONS | 434 |
| 11.7.1 | Forcing the Use of Indexes and Join Orders | 434 |
| 11.7.2 | What You Cannot Force | 435 |
| 11.7.3 | Subtle Ways to Influence the Optimizer's Decisions | 436 |
| Chapter 12: Batch Processing Performance | | 441 |
| 12.1 | GENERAL GUIDELINES | 442 |
| 12.1.1 | Concurrent Batch Operations | 442 |
| 12.1.2 | Host-Based Communication | 443 |
| 12.1.3 | Divide and Conquer | 445 |
| 12.1.4 | Data Loading and Indexes | 446 |
| 12.1.5 | Server Configuration | 446 |
| 12.2 | DUMP, LOAD, AND RECOVERY | 449 |
| 12.2.1 | Dump | 450 |
| 12.2.2 | Load | 453 |
| 12.2.3 | Using Dump and Load to Copy a Database | 454 |
| 12.2.4 | Recovery | 455 |
| 12.2.5 | Dump and Load with 4.9 SQL Server | 458 |
| 12.2.6 | Dump and Load Using OS Tools | 459 |
| 12.3 | PERFORMANCE OF BACKUP SERVER | 459 |
| 12.3.1 | Architecture of Backup Server | 460 |
| 12.3.2 | Performance of Dump | 464 |
| 12.3.3 | Mitigating Impact on OLTP Performance | 465 |
| 12.4 | CREATE AND ALTER DATABASE | 465 |
| 12.4.1 | Improvements in System 10 SQL Server | 466 |
| 12.4.2 | Ordering of Devices | 466 |
| 12.4.3 | Create Database for Load | 467 |
| 12.5 | POPULATING A DATABASE | 467 |
| 12.5.1 | General Issues | 468 |
| 12.5.2 | select into | 472 |
| 12.5.3 | Fast BCP | 472 |
| 12.5.4 | Normal BCP | 473 |
| 12.6 | CREATING INDEXES | 474 |
| 12.6.1 | Nonclustered Indexes | 475 |
| 12.6.2 | Clustered Indexes | 476 |
| 12.6.3 | Creating Indexes Concurrently | 477 |

| | | |
|--|---|------------|
| 12.7 | LOADING DATA INTO INDEXED TABLES | 478 |
| 12.7.1 | Clustered Indexes | 479 |
| 12.7.2 | Nonclustered Indexes | 481 |
| 12.7.3 | Recommendations | 481 |
| 12.8 | MISCELLANEOUS BATCH OPERATIONS | 482 |
| 12.8.1 | Checkpoints | 482 |
| 12.8.2 | DBCC Commands | 483 |
| 12.8.3 | Long-Running Reports | 486 |
| 12.8.4 | Defragmentation | 488 |
| Chapter 13: Advanced Topics | | 491 |
| 13.1 | HIGH PERFORMANCE I/O SUBSYSTEMS | 491 |
| 13.1.1 | Implications of I/O Performance | 491 |
| 13.1.2 | Caching by the I/O Subsystems | 494 |
| 13.1.3 | Representative I/O Subsystems | 496 |
| 13.2 | NAMED CACHES AND BUFFER POOLS | 500 |
| 13.2.1 | Managing Named Caches and Buffer Pools | 503 |
| 13.2.2 | Why Partition Memory? | 509 |
| 13.2.3 | Impact of Buffer Sizes | 512 |
| 13.2.4 | Buffer Replacement Policies | 518 |
| 13.3 | OTHER ENHANCEMENTS | 519 |
| 13.3.1 | Table Partitioning | 519 |
| 13.3.2 | Log and Transaction Management | 521 |
| 13.3.3 | Lock Management | 523 |
| 13.3.4 | Index Management | 525 |
| 13.3.5 | Page Allocation and Deallocation | 527 |
| 13.3.6 | Housekeeper | 528 |
| Appendix A: General Configuration Practices | | 531 |
| Appendix B: SQL Server Configuration | | 533 |
| B.1 | CONFIGURATION PROPERTIES | 533 |
| B.1.1 | Scope | 533 |
| B.1.2 | SQL Server Version | 534 |
| B.1.3 | Persistence | 534 |

| | | |
|---|--|------------|
| B.1.4 | Dynamic versus Static | 534 |
| B.1.5 | Config Value versus Run Value | 534 |
| B.2 | CHANGING SQL SERVER CONFIGURATION | 535 |
| B.2.1 | sp_configure | 535 |
| B.2.2 | Configuration File | 536 |
| B.2.3 | buildmaster -y | 537 |
| B.2.4 | sp_dboption | 538 |
| B.2.5 | Trace Flags | 539 |
| B.2.6 | dbcc tune | 539 |
| B.2.7 | sp_cacheconfig, etc. | 540 |
| B.3 | MANAGING CONFIGURATION ERRORS | 540 |
| B.4 | TABLE AND INDEX CONFIGURATION | 542 |
| B.5 | DATABASE CONFIGURATION | 543 |
| B.6 | SERVER-WIDE CONFIGURATION | 545 |
| Appendix C: Network Configuration | | 561 |
| C.1 | INTERFACES FILE | 561 |
| C.2 | SQL SERVER'S INTERFACES FILE ENTRY | 563 |
| C.3 | CLIENTS' INTERFACES FILE ENTRIES | 564 |
| C.4 | USING LOOPBACK | 565 |
| Appendix D: DBCC Commands | | 567 |
| Appendix E: Troubleshooting | | 571 |
| Appendix F: Industry Standard Benchmarks | | 577 |
| F.1 | TRANSACTION PROCESSING COUNCIL | 578 |
| F.2 | TPC-A AND TPC-B | 578 |
| F.3 | TPC-C | 579 |
| F.4 | TPC-D | 580 |
| F.5 | TPC BENCHMARK RESULTS | 581 |

| | |
|--|------------|
| Appendix G: Reference and Scripts | 583 |
| G.1 USEFUL FUNCTIONS | 583 |
| G.2 SQL SCRIPTS | 585 |
| G.2.1 Locks | 588 |
| G.2.2 Row Size | 590 |
| G.2.3 Space Used | 591 |
| G.2.4 Columns | 592 |
| G.2.5 Devices View | 592 |
| G.2.6 Indexes | 593 |
| G.2.7 Statistics | 593 |
| G.2.8 Last Log Record | 594 |
| G.3 SQL SERVER DATATYPES | 595 |
| G.4 SYSTEM TABLES | 597 |
| | |
| Glossary | 605 |
| | |
| Index | 613 |