

Index

Note to the Reader: Throughout this index **boldfaced** page numbers indicate primary discussions of a topic. *Italicized* page numbers indicate illustrations.

A

- abort errors, 193–194
- Abort method
 - explicit transactions, 279
 - implicit transactions, 281
- absolute XPath expressions, 100
- ACID properties, 242
- Active cursors counter, 23
- ActiveX Data Objects (ADO), 6
 - error handling, 216
 - explicit transactions, 279
 - MARS support, 71
 - with SQLNCLI, 11–12
- Activity—All Blocking Transactions report, 344
- activity data, 425
- Activity Monitor
 - locking problems, **342–344**, 343–344
 - performance, 453
- actual execution plans, 315
- ad hoc queries
 - for plan reuse, 88–89
 - vs. stored procedures, 179
- Add Reference dialog box, 50
- addition, XQuery, 102
- ADDLOCAL option, 45
- ADO. *See* ActiveX Data Objects (ADO)
- ADO.NET
 - caching, 28–31, 28–29
 - connection pooling, 27
 - DataReader object, 347
 - DataSet object, 347
 - explicit transactions, 279
 - MARS support, 71
 - XmlReader object, 348
- ADOMD.NET, 64
- after keyword, 114
- AFTER triggers, 499
- agent alerts, 221, **226–229**, 227–229
- aggregations
 - denormalization for, 26
 - indexed views for, 385–387
 - OLTP systems, 522
- aging in procedure cache, 93–94
- airline tickets, 434
- alerts
 - agent, 221, **226–229**, 227–229
 - performance, 453
- algebraizer, 87, 313
- alias data types, 517
- All Blocking Transactions report, 344, 344
- All Transactions report, 344
- ALLOW_PAGE_LOCKS option, 391
- ALLOW_ROW_LOCKS option, 391
- ALLOW_SNAPSHOT_ISOLATION option, 253, 267
- ALTER permission, 185
- ALTER ANY permission, 180
- ALTER DATABASE statement, 264
- ALTER INDEX statement, 276, 392–393
- ALTER TABLE statement, 536–537
- American National Standards Institute (ANSI) encoding, 514
- AMO (analysis management objects), 44

- AMO (Automation Management Objects) applications, **61–64**, 62–63
- ampersands (&) in XQuery, 100
- Analysis Services Scripting Language (ASSL), 61
- and operator, 102
- angle bracket symbols (<) in XQuery, 100
- ANSI (American National Standards Institute) encoding, 514
- ANSI_NULL_DEFAULT option, 535
- ANSI_NULL_DFLT_ON option, 535
- ANSI_NULLS option, 386
- ANSI_PADDING option, 386
- ANSI_WARNINGS option, 386
- APIs (application programming interfaces)
 - function, 4–5
 - server cursors, 20
- application blocks, 19
- application layer caching, 27–28
- application programming interfaces (APIs)
 - function, 4–5
 - server cursors, 20
- ApplicationName property, 30
- applications, **43–44**
 - client libraries, **44–46**
 - design simplification. *See* Multiple Active Result Sets (MARS)
 - error handling, **216–218**
 - events. *See* user-defined messages
 - exam essentials, 77
 - management object, **46**, 47
 - AMO, **61–64**, 62–63
 - RMO, **58–61**
 - SMO. *See* server management objects (SMO)
 - SNI, **64–67**, 65
 - review questions, **78–83**
 - summary, **76–77**
 - well-performing, **423–425**
- approximate numeric data type, **511–512**
- architecture
 - data access technologies, **5–6**
 - SODA, **434–435**, **440–441**
- ARITHABORT option, 386
- arithmetic operators, 102
- assemblies, referencing, **50–51**, **56–57**
- Assemblies folder, 46
- ASSL (Analysis Services Scripting Language), 61
- asterisks (*) in XQuery, 102
- asynchronous queries, **64–67**, 65
- atomicity, transaction, 241
- attentions, 197
- authentication, 51
- AUTO_CREATE_STATISTICS option, 334
- auto-parameterization, 89
- AUTO_UPDATE_STATISTICS option, 334
- autocommit transactions, 280
- automating tasks, **50–51**
- Automation Management Objects (AMO) applications, **61–64**, 62–63
- availability as design goal, 187
- Available Bytes counter, 485
- Available Mbytes counter, 484–485
- Average Disk Sec/Read counter, 485
- Average Disk Sec/Write counter, 486
- Average Page Life Expectancy counter, **483–484**
- Avg. Disk Queue Length counter, 487
- avg_fragmentation_in_percent column, 401
- avg_page_space_used_in_percent column, 401

B

B-tree structures, 370, 370
Backup class, 16
backward compatibility in SMO, 346
balanced tree structures, 370, 370
baselines, performance, 457
batch abort errors, 193–194
batch queries, 87–88
batch-scoped transactions, 282–286
batches
 cached, 93–94
 for plan reuse, 92
before keyword, 114
BEGIN CATCH statement, 191–192
BEGIN DISTRIBUTED
 TRANSACTION statement, 294
BEGIN TRANSACTION statement,
 279, 287, 291–292
BEGIN TRY statement, 191–192
BeginTrans method, 279
BeginTransaction method, 279
BETWEEN operators
 clustered indexes, 373
 SARGs, 350
bigint data type, 511
binary data types, 515
binary large objects (BLOBs), 14
binding phase in queries, 119–121, 313
bit data type, 515
bitmaps, NULL, 520–521
BLOBs (binary large objects), 14
Blocked Process Report event class,
 489–492, 490–491
blocking problems, 488
 DAC for, 494–498, 496
 locking, 489–492, 490–491
 row lock waits value, 492–493
 statistics, 488–489
 wait types, 493–494

blocks

 application, 19
 variable, 521
bodies in XQuery expressions, 102
bound transaction sessions, 286
branching in queries, 135
BufferManager Buffer cache hit ratio
 counter, 33
BufferManager Database pages
 counter, 33
BULK INSERT statement, MARS, 72
bulk update (BU) locks, 245
Bytes Total/sec counter, 499

C

Cache Hit Ratio counter, 23
cached batch queries, 93–94
Cached Cursor Counts counter, 23
CacheHit metric, 33
CacheInsert metric, 33
CacheMiss metric, 33
CacheRemove metric, 33
caching, 25
 ADO.NET, 28–31, 28–29
 aging in, 93–94
 application layer, 27
 custom functionality, 31–32
 database server layer, 26
 front-end application layer, 27–28
 queries, 87–88
 refresh strategy for, 32–34
 SODA support for, 442
 strategies, 25–28
 web services layer, 27
CALLER option, 184
capture execution, 48
cardinality, 380–381
CardSystems Solutions, 171

- carets (^) in regular expressions, 214
- CAST function, 514
- catalog views, 407
- catchable errors. *See* TRY...
 - CATCH construct
- central relational databases, 218
- centralized message management, 221
- char data type, 514
- CHAR function, 171
- character strings
 - data types, 514
 - filtering, 171
 - Unicode, 515
- CHARINDEX function, 536
- CHECKPOINT command, 325, 336
- child objects, 54
- circular buffers, 214
- Client Network Utility, 45
- clients and client-side technologies, 3
 - cursors with, 20–22, 140–141, 140–141, 160
 - in data access architecture, 5–6
 - libraries, 3, 44–46, 71
 - optimizing, 347–348
- clinical trials, 70
- Clustered Index Scan operator, 339, 349
- Clustered Index Seek operator, 317
- clustered indexes, 367–368
 - guidelines, 374–375
 - overview, 370–374, 370–371
 - scans, 333–334, 339, 349, 368
- CLUSTERED option, 391
- Codd, Edgar F., 424, 523
- code for HTTP endpoints, 14
- COL_LENGTH function, 536
- columns
 - caching, 32
 - computed, 380–383, 382–383
 - derived, 526–527
 - included, 367, 383–384
 - in partitioning, 427
 - redundant, 525–526
 - variable-length, 530–534
- comments, XQuery, 100
- Commit method
 - explicit transactions, 279
 - implicit transactions, 281
- commit phase in distributed transactions, 287
- COMMIT TRANSACTION statement, 279, 291–293
- COMMIT WORK statement, 279, 291
- CommitTrans method, 279
- Common Vulnerabilities and Exposures (CVE), 171
- comparison operators, 101
- compatibility
 - locks, 75, 245–247
 - SMO, 346
- compilations
 - CPU problems, 471–476, 474–475
 - errors, 194–195
 - plans, 87, 307
 - process, 312
 - stored procedures, 26
- component interaction, 4–5
- composite index keys, 384
- COMPUTE keyword, 20
- COMPUTE BY keyword, 20
- computed columns, 380–383, 382–383
- CONCAT_NULL_YIELDS_NULL
 - option, 386
- concurrency
 - MARS, 70
 - side effects
 - inconsistent analysis, 250
 - lost updates, 247–249
 - phantom reads, 251
 - uncommitted dependency, 249–250

- transactions, 240, 242
 - well-performing applications, 424
- conditional branches, 135
- conjunction operator, 102
- CONNECT permission, 67
- ConnectAsUser property, 30
- ConnectAsUserName property, 30
- ConnectAsUserPassword property, 30
- connection strings
 - MARS, 71
 - SQLNCLI, 10–11
- ConnectionInfo assembly, 51
- ConnectionInfo.dll file, 47
- connections
 - ADO.NET caching setting, 29–31
 - to other data stores, 14
 - pooling, 27, 31
 - replication servers, 58–59
 - SQLNCLI references, 10–11
- @@CONNECTIONS function, 453
- ConnectionSettings class, 29
- ConnectionString property, 30
- ConnectTimeout property, 30
- consistency, transaction, 240–241
- CONTAINS predicate, 376–377
- CONTAINSTABLE predicate, 376
- context
 - execution. *See* execution contexts
 - vs. plans, 307
 - XQuery, 100
- context switching, 181–182
- CONTROL permission, 180
- CONVERT function, 514
- converting
 - characters, 514
 - single-row statements into set-based queries, 351–353
 - XML data into rowset data, 109–113
- COOKIE INTO option, 182
- copies for transactions, 242
- cost-based optimizers, 313
- costs
 - cached batches, 93
 - online index operations, 393
- counters. *See* performance
- CPU problems
 - counters, 471
 - inefficient plans, 476–482
 - parallelism, 468–471
 - plan reuse for, 467–468
 - recompilations, 471–476, 474–475
 - scheduler, 464–465
 - signal waits, 466
- crashes from input, 171
- CREATE ASSEMBLY statement, 173
- CREATE ENDPOINT statement,
 - 13, 66
- CREATE FUNCTION statement, 173
- CREATE INDEX statement, 276,
 - 390–391, 539
- Create method, 54
- CREATE TYPE statement, 517
- CROSS APPLY operator, 109
- cross-site scripting (XSS)
 - vulnerability, 171
- currency values, 513
- Current Bandwidth counter, 499
- Cursor Cache Use Count/sec
 - counter, 23
- cursor data type, 515
- Cursor memory usage counter, 23
- Cursor Requests/sec counter, 24
- @@CURSOR_ROWS function, 24
- CURSOR_STATUS function, 24
- Cursor worktable usage counter, 24
- CursorClose event, 23
- CursorExecute event, 23
- CursorImplicitConversion event, 23
- CursorOpen event, 23
- CursorPrepare event, 23

CursorRecompile event, 23
cursors, 133–134
 appropriate use, 139
 client-side, 20–22, 140–141, 140–141, 160
 CPU problems, 471
 dynamic, 138, 153–156, 157
 exam essentials, 161
 fast_forward, 137
 joins without, 157–159, 159
 keyset, 138, 149–152, 153
 metadata, 24
 ODBC, 136
 operation, 136
 opinions on, 143
 performance, 20–24, 351–352
 reasons for, 19
 result set-based operations, 135–136
 review questions, 162–167
 row-based operations, 136
 selecting, 25
 server, 20–21, 141, 142, 144–145
 vs. set-oriented operation, 143
 SQL Server Profiler for, 23
 static, 137, 146–148, 149
 stored procedures, 24
 strategies, 20
 summary, 160–161
 system functions, 24–25
 System Monitor for, 23–24
 tracking, 22–25
CursorUnprepare event, 23
custom log files, 218
custom permissions, 185–186
CVE (Common Vulnerabilities and Exposures), 171
CXPACKET wait type, 468–469

D

DAC (Dedicated Administrator Connection)
 for blocking problems, 494–498, 496
 for performance, 454
data access design, 1–4, 4
 architecture, 5–6
 caching. *See* caching
 component interaction, 4–5
 connections to other data stores, 14
 cursors. *See* cursors
 exam essentials, 35
 HTTP endpoints and SOAP, 12–14, 12
 legacy technologies, 6–8
 object models, 15–19
 review questions, 36–41
 SQLNCLI, 9–12
 summary, 34–35
Data Definition Language (DDL)
 statements, 15
 ASSL for, 61
 MARS, 73
 schema locks for, 245
 triggers, 184
data-dependent routing (DDR), 439, 440
data isolation, scaleout from, 426
data loss from SQL injection attacks, 174
Data Modification Language (DML), 106, 113
 delete keyword, 117–118
 insert keyword, 114–115
 LOB support, 536
 MARS statements, 73
 replace value of keyword, 116–117
 triggers, 184
data pages, 366, 369, 372

- data records, 520
- data source name (DSN) connection
 - definitions, 7
- data stewardship, 437
- data storage
 - database objects, 520–521
 - denormalization, 525–529
 - exam essentials, 549–550
 - LOBs, 541
 - normalization, 522–525, 522–524
 - NULLs, 534–535
 - OLTP vs. OLAP, 521–522
 - review questions, 551–559
 - row-overflow data, 545–549
 - row size limits, 535
 - summary, 549
 - types. *See* data types
 - variable-length columns, 530–534
- data types, 510–511
 - large value, 541–545
 - MAX specifier, 536–541
 - options, 511–516
 - synonyms, 516–518
 - using, 518–520
- Database Console commands, 453
- database engine isolation levels. *See* isolation levels
- Database Engine Tuning Advisor (DTA)
 - indexes, 393–394, 396–399, 397–398
 - performance, 453
 - query plans, 334
- database management system (DBMS)-specific drivers, 7
- database server layer caching, 26
- DatabaseName property, 30
- databases
 - concepts, 366
 - in data access architecture, 5
 - properties, 52–53, 57
 - shared, 435–436, 436
- Databases property, 53
- DATALENGTH function, 518, 536
- DataReader object, 29, 347
- DataSet object, 28–29, 29, 347
- date data types, 513
- DATETIME columns, 374
- datetime data type, 513–514
- DB-Library
 - considerations, 6
 - errors and messages, 216
 - explicit transactions, 279
- DBCC DROPCLEANBUFFERS
 - command, 325, 336
- DBCC FREEPROCCACHE command, 325, 336, 456, 484
- DBCC FREESYSTEMCACHE
 - command, 484
- DBMS (database management system)-specific drivers, 7
- DDL (Data Definition Language)
 - statements, 15
 - ASSL for, 61
 - MARS, 73
 - schema locks for, 245
 - triggers, 184
- DDR (data-dependent routing), 439, 440
- deadlocks, 209–213
- decimal data type, 511
- Decision Support, 521
- DECLARE CURSOR statement
 - fast_forward cursors, 137
 - Transact-SQL cursors, 20
- Dedicated Administrator Connection (DAC)
 - for blocking problems, 494–498, 496
 - for performance, 454
- default result sets, 67
- default trace tool, 454–456
- defense-in-depth security model, 179, 179

delete keyword, 117–118
 DELETE statement
 for fragmentation, 402
 with server cursors, 20
 denormalization, 525
 advantages, 26
 derived columns, 526–527
 duplicate tables, 528
 indexed views, 529
 prejoined tables, 527–528
 redundant columns, 525–526
 row width, 529–530
 splitting tables, 528–529
 tables, 458
 deprecated features, 6
 derived columns, 526–527
 DETAILED scan mode, 401
 dirty reads, 249–250
 DISABLE option, 393
 disjunction operator, 102
 disk activity statistics, 323–324
 % Disk Time counter, 487
 distributed partitioned views, 438–439
 distributed queries, 438, 438
 distributed transactions
 phases, 286–287
 working with, 294–296, 295
 div operator, 102
 division, 102
 DLLs (dynamic-link libraries), 8
 dm_exec_cached_plans view, 33
 dm_exec_query_plan view, 33
 dm_exec_query_stats view, 33
 dm_exec_requests view, 33
 DML. *See* Data Modification Language (DML)
 DMO object model, 18
 dollar signs (\$)
 regular expressions, 214
 XQuery, 100
 drivers, ODBC, 7

DROP ENDPOINT statement, 66
 DROP_EXISTING option, 391, 393
 DSN (data source name) connection
 definitions, 7
 DTA (Database Engine Tuning Advisor)
 indexes, 393–394, 396–399, 397–398
 performance, 453
 query plans, 334
 duplicate tables, 528
 durability in transactions, 241
 dynamic cursors, 136, 138,
 153–156, 157
 dynamic-link libraries (DLLs), 8
 dynamic management views and
 functions
 caching, 33–34
 index metadata, 408
 locking problems, 344–345
 performance, 453, 457
 dynamic SQL, 178
 dynamic strings, 186

E

EIF (Enterprise Instrumentation Framework), 218
 elevation of privileges, 174
 encapsulation
 EXECUTE AS, 185
 indexed views, 385
 stored procedures, 179
 EncryptConnection property, 30
 encrypting passwords, 178
 END CATCH statement, 191
 END TRY statement, 191
 endpoints
 HTTP, 12–14, 12
 SODA support, 441
 TDS, 64–66, 65
 Transact-SQL, 3, 66–67

- Enterprise Instrumentation Framework (EIF), 218
- entity references, 100
- EnumLocks method, 345
- equal signs (=)
 - clustered indexes, 373
 - SARGs, 350
 - XQuery, 101
- @@ERROR function, 189–191
- error handling
 - applications, 216–218
 - best practices, 220–221
 - design for, 187
 - @@ERROR, 189–191
 - error functions, 192, 200–201
 - error structure, 187–188
 - event notifications, 197–200
 - exam essentials, 229
 - input validation, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
 - logging, 218–220, 453
 - principle of least privilege, 180
 - EXECUTE AS, 181–186
 - permissions, 180–181
 - SQL injection, 178
 - PRINT, 214–216
 - RAISERROR, 201–206
 - review questions, 230–238
 - severity levels, 188–189
 - summary, 229
 - transactions, 206–209, 294
 - TRY...CATCH construct. *See* TRY...CATCH construct
 - user-defined messages. *See* user-defined messages
- ERROR_LINE function, 201
- error logs, 218–220, 453
- ERROR_MESSAGE function, 201
- ERROR_NUMBER function, 200
- error numbers, 188
- ERROR_PROCEDURE function, 200
- ERROR_SEVERITY function, 200
- ERROR_STATE function, 200
- estimated execution plans, 315
- EstimateExecution attribute, 334
- EstimateRows attribute, 334, 336
- Event Notifications
 - blocking problems, 490
 - warning handling, 197–200
- events
 - application. *See* user-defined messages
 - cursor, 23
- ever-increasing patterns, 374
- exact numeric data types, 511
- exception handling, 178, 187. *See also* error handling
- exclamation points (!) in XQuery, 101
- exclusive (X) locks, 244, 246
- EXECUTE AS clause, 181
 - custom permissions, 185–186
 - execution context
 - dynamic strings, 186
 - modules, 183–185
 - sessions, 181–183
- EXECUTE AS CALLER option, 182
- execution contexts
 - dynamic strings, 186
 - modules, 183–185
 - plan cache for, 307
 - vs. plans, 307
 - sessions, 181–183
- execution plans. *See* plans
- exist method, 106–109, 122, 380
- EXPAND VIEWS hint, 387
- expanding data types, 517–518
- explicit memory errors, 484

explicit transactions, 279–280
 marked, 287–289
 nesting, 291–293
 savepoints, 289–291
 statements not allowed in, 293
 syntax, 287

expressions
 regular, 173, 213–214
 XQuery, 100–102

extents, 366, 520

external fragmentation, 399

external memory pressure, 483

F

fast_forward cursors, 137

FAST hint, 476

federated database servers, 439

FETCH statement, 72

@@FETCH_STATUS function, 24–25

File DSNs, 7

FILLFACTOR option
 fragmentation, 402, 404–405
 indexes, 391

filters
 input data, 171
 plans, 336

FINALLY blocks, 192

first normal form, 424, 523–524,
 523–524

fixed-length data, 520–521

flexibility
 OLE DB data sources, 8
 scaleout factor, 426

float data type, 512–513

FLWOR statements, 102–105

fn_trace_getinfo function, 454

fn_trace_gettable function, 456

FOR BROWSE keyword, 20

FOR statement, 102

FOR XML function, 99

FOREIGN KEY constraint, 376

FORMATMESSAGE function,
 222, 226

formatting user-defined messages, 226

forward-only cursors, 136
 DYNAMIC READ_ONLY,
 153–156, 157
 KEYSER READ_ONLY,
 149–152, 153
 STATIC READ_ONLY,
 146–148, 149

fragment_count column, 401

fragmentation, 369, 399–400
 detecting, 401
 fixing, 402
 preventing, 402–405

FREETEXT predicate, 376

FREETEXTTABLE predicate, 376

front-end application layer caching,
 27–28

full scans as problem indicator, 486

full-text indexes, 367, 376–377

fully qualified object references, 88

functions
 computed columns, 381
 cursor, 24–25
 error, 200–201
 index metadata, 408
 for performance, 453
 XQuery, 101

G

Global Allocation Map (GAM)
 pages, 367

globally unique identifiers
 (GUIDs), 374

GOTO statements, 192

granularity, locking
 hints, 270–276
 overview, 243–244

graphical plan display, 325–326, 326
 greater than signs (>)
 clustered indexes, 373
 SARGs, 350
 XQuery, 101
 GROUP BY clause, 373
 GUIDs (globally unique identifiers), 374

H

hackers, 171
 Hash Match operator, 349
 hash operators, 336, 349
 headers, record, 520
 heaps, 243, 368–370, 369
 hierarchical operator tree, 317
 high-level query life cycle view,
 306–307, 307
 hints
 evaluating, 336
 locking, 270–276
 HOLDLOCK hint, 256, 261, 271–272
 horizontal table splits, 528–529
 HTTP endpoints, 12–14, 12

I

IAM (Index Allocation Map) pages,
 367, 369–370
 IDENTITY columns, 374
 @@IDENTITY function, 220
 IDENTITY property, 220
 identity theft, 174
 IF..GOTO statements, 190
 IGNORE_DUP_KEY option, 391
 image columns, 539
 image data type, 515, 535
 IMPERSONATE ON permission, 180
 implementation
 MARS, 67–69
 queries, 314

implicit transactions, 280–282
 in-process data providers, 15
 IN_ROW_DATA allocation
 unit, 545
 in-row large values, 541–545
 Include Actual Execution Plan
 command, 326
 included columns, indexes with, 367,
 383–384
 inconsistent analysis, 250
 Index Allocation Map (IAM) pages,
 367, 369–370
 INDEX_COL function, 408
 index pages, 366
 Index Physical Statistics report, 406
 Index Property window, 406
 index records, 520
 Index Scan operator, 349
 index seeks, 373
 Index Tuning Wizard, 393
 Index Usage Statistics report, 406
 indexed views, 385–389, 529
 indexes, 365–366
 benefits, 136
 catalog views, 407
 clustered, 367–368
 guidelines, 374–375
 overview, 370–374, 370–371
 scans, 333–334, 339, 349, 368
 computed columns, 380–383,
 382–383
 concepts, 366
 creating, 390–391
 exam essentials, 409
 and fragmentation, 399–405
 full-text, 376–377
 vs. heaps, 368–370, 369
 with included columns, 383–384
 key size, 384
 limits, 389
 locking behavior, 276
 metadata, 406–408

614 INDEXKEY_PROPERTY function – joins

- missing, 331
 - identifying, 393–399, 397–399, 462
 - memory problems from, 483
 - modifying, 392–393
 - nonclustered, 375–376, 375, 391
 - online operations, 393
 - performance, 459–460
 - review questions, 410–420
 - row versioning framework
 - operations, 252
 - summary, 409
 - text, ntext, and image columns
 - for, 539
 - types, 367
 - unused, 463–464
 - for well-performing
 - applications, 424
 - XML, 367, 377–379
 - INDEXKEY_PROPERTY
 - function, 408
 - INDEXPROPERTY function, 408
 - inefficient plans
 - improving, 476–482
 - uncovering, 334–336
 - InfoMessage events, 217
 - informational messages, 192
 - Inner Join logical operator, 317
 - input validation, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
 - insert keyword, 114–115
 - instantiation
 - local variables, 53–54
 - server objects, 51–52
 - SMO, 48
 - int data type, 511
 - integer data types, 511
 - intent locks, 245
 - intent exclusive (IX) locks, 245–246
 - intent shared (IS) locks, 245–246
 - intent update (IU) locks, 245
 - interaction
 - component, 4–5
 - service, 426
 - interleaving, 67–68
 - intermediary index levels, 373
 - intermediate B-tree nodes, 370, 370
 - internal fragmentation, 399
 - internal memory pressure, 483
 - INTO keyword
 - DML, 114
 - server cursors, 20
 - I/O activity
 - performance problems, 485–487
 - statistics, 323–324
 - table scans and range scans, 461–463
 - IS (intent shared) locks, 245–246
 - IS NOT NULL option, 535
 - IS NULL option, 535
 - ISNULL option, 535
 - isolation levels, 241–242, 252–253
 - Read Committed, 256–259
 - Read Committed Snapshot, 262–266
 - Read Uncommitted, 253–256
 - Repeatable Read, 259–260
 - Serializable, 260–261
 - Snapshot, 266–269
 - summary, 270
 - iterators, 317
 - IU (intent update) locks, 245
 - IX (intent exclusive) locks, 245–246
-
- J**
- Jet technology, 6
 - joins
 - client-side technologies, 140–141, 140–141
 - clustered indexes, 373
 - without cursors, 157–159, 159

indexed views for, 385
 nonclustered indexes, 376
 performance, 335, 458–459
 prejoined tables, 527–528
 rewriting subqueries to, 348–350
 server-side technologies, 141, 142

K

Keep All Existing PDS option, 398
 KEEP PLAN hint, 476
 KEEPFIXED PLAN hint, 476
 key-range locks, 245
 keys
 clustered indexes, 374
 scaleout factor, 426
 size, 384
 keyset cursors, 136, 138, 149–152, 153
 KILL statement, 196

L

large object (LOB) data types
 DML support, 536
 in-row, 541–545
 in key columns, 384
 limitations, 535
 text/image pages for, 367
 LastName columns, 375
 layering process in caching, 26–28
 LCK_ wait types, 493
 leaf levels and nodes, 373
 B-tree structures, 370, 370
 nonclustered indexes, 376
 legacy technologies, 6–8
 LEN function, 536
 less than signs (<)
 clustered indexes, 373
 SARGs, 350
 XQuery, 100–101

LET statement, 102
 level abort errors, 194
 levels
 isolation. *See* isolation levels
 severity, 188–189
 leveraging applications
 data considerations, 425–426
 foundation, 422–423, 423
 partitioning, 426–433
 scaleout factors, 426
 scaleup vs. scaleout, 425
 well-performing applications, 423–425
 libraries
 client, 3, 44–46, 71
 DLL, 8
 LIKE operator, 350
 Limit Tuning Time option, 398
 LIMITED scan mode, 401
 line numbers for error
 messages, 188
 linked servers, 438, 438
 literals, 100
 LOB (large object) data types
 DML support, 536
 in-row, 541–545
 in key columns, 384
 limitations, 535
 text/image pages for, 367
 LOB_DATA allocation unit, 545
 @local_variable argument, 203
 local variable instantiation, 53–54
 Lock Manager, 74–75
 locks and locking, 242
 compatibility, 245–247
 database engine, 242–243
 granularity
 hints, 270–276
 overview, 243–244
 indexes, 276
 modes, 244–246

problems, 340–341
 Blocked Process Report event
 class, 489–492, 490–491
 dynamic management objects for,
 344–345
 row lock waits value, 492–493
 SMO for, 345–346
 SQL Profiler for, 341
 SQL Server Management Studio
 for, 342–344, 343–344
 System Monitor for, 341–342
 Locks by Object page, 342
 Locks by Process page, 342
 Log File Viewer, 453
 LOG option, 203–204
 logical operators, 102
 logical query processing, 308–311
 Login property, 30
 LoginSecure property, 30
 logs, 453
 errors, 218–220
 performance, 453
 RAISERROR, 203–204
 user-defined messages, 225
 loops, nested, 317, 319
 loosely coupled modules, 179
 lost updates, 247–249

M

manageability
 as design goal, 187
 partitioning for, 433
 managed code, errors from, 213–214
 ManagedComputer class, 16
 management object applications, 46, 47
 AMO, 61–64, 62–63
 RMO, 58–61
 SMO. *See* server management
 objects (SMO)
 SNI, 64–67, 65

marked transactions, 287–289
 MARS. *See* Multiple Active Result
 Sets (MARS)
 matching text for plan reuse, 88
 max degree of parallelism option, 476
 MAX specifier, 536–541
 MAXDOP option, 391, 476
 MaxPoolSize property, 30
 MDAC (Microsoft Data Access
 Components)
 for clients, 12
 vs. SQLNCLI, 9
 memory
 cursors, 21
 troubleshooting, 483–485
 Memory Grants Pending counter, 484
 messages
 error, 188
 user-defined. *See* user-defined
 messages
 metadata
 cursors, 24
 indexes, 406–408
 Microsoft ActiveX Data Objects
 (ADO), 6
 error handling, 216
 explicit transactions, 279
 MARS support, 71
 with SQLNCLI, 11–12
 Microsoft Data Access
 Components (MDAC)
 for clients, 12
 vs. SQLNCLI, 9
 Microsoft Distributed Transaction
 Coordinator (MSDTC), 286,
 294, 295
 Microsoft Full-Text Engine for SQL
 Server (MSFTESQL), 376
 Microsoft SQL Server Development
 Customer Advisory Team, 457
 middle layer caching, 27

MinPoolSize property, 30
minus signs (-) in XQuery, 102
mirrors
 SQLNCLI, 9
 table, 528
missing indexes, 331
 identifying, 393–399, 397–399, 462
 memory problems from, 483
MissingIndexes section, 320–321
mod operator, 102
modes, locks, 244–246
modify method, 380
modifying
 indexes, 392–393
 for query performance, 97–98, 98
 SMO objects, 55–56
 user-defined messages, 224
modules
 execution context, 183–185
 loosely coupled, 179
monetary data types, 513
monitoring
 buffer caches, 26
 performance, 456–457
 recompilations, 473–476, 474–475
MSDTC (Microsoft Distributed
 Transaction Coordinator), 286,
 294, 295
MSFTESQL (Microsoft Full-Text
 Engine for SQL Server), 376
msg_id argument, 202–203
msg_str argument, 203
multiple active cursor statements, 21
Multiple Active Result Sets (MARS), 44
 batch-scoped transactions, 282
 in clinical trials, 70
 evaluating, 73–74
 features, 69
 implementation, 67–69
 queries, 71–72, 347

 reads and writes mixing on, 73, 74
 resources, 75–76, 252
 row versioning, 69–70
 and SQL Server Lock Manager,
 74–75
 statements, 72–73
multiple language message support, 221
MultipleActiveResultSets property, 30
multiplication, XQuery, 102
My Computer Properties window,
 295, 295
myConn_InfoMessage function, 217

N

named pipes, 66
names
 savepoints, 289
 tables, 334
namespaces, SMO, 16–17, 48
narrow keys, 374
native web service access, 441
nchar data type, 515, 517
nested loops, 317, 319
nesting explicit transactions, 291–293
network adapters, 499
Network Monitor Agent, 453
NetworkProtocol property, 30
NEWID function, 375
NEWSEQUENTIALID function, 375
NEXT statement, 137
NO REVERT option, 182
NO_WAIT option, 265
nodes, B-tree structures, 370, 370
nodes method, 107, 109–113, 122
NOEXPAND hint, 387
NOLOCK hint, 244, 250, 272–273
nonclustered indexes, 367–368,
 375–376, 375, 391

NONCLUSTERED option, 391
 nonparameterized SQL, 32
 NonPooledConnection property, 31
 nonrepeatable reads, 250
 nontransparent scaleout, 435
 normalization, 424, 522–523, 522
 advantages, 525
 first normal form, 523–524,
 523–524
 second normal form, 524, 524
 third normal form, 524, 524
 notifications, queries, 34
 NOWAIT locking hint, 273–274
 NOWAIT option, 204
 ntext columns, 539
 ntext data type, 515, 517, 535
 NULLs
 bitmaps, 520–521
 handling, 534–535
 Number of active cursor plans
 counter, 24
 numeric data type, 511
 NUMERIC_ROUND_ABORT
 option, 386
 nvarchar data type, 515, 517
 row-overflow data, 545
 row size, 535
 using, 536–541

O

Object Explorer, 343, 343, 401
 object models
 data access design, 15–19
 OLE DB data providers, 8
 objects, SMO, 53–56
 ODBC. *See* Open Database
 Connectivity (ODBC)
 off-row large values, 541

OLAP (Online Analytical Processing)
 characteristics, 522
 FILLFACTOR for, 402
 vs. OLTP, 521–522

OLE DB
 client-side cursor support, 22
 connection pooling, 31
 considerations, 7–8
 error handling in, 216
 explicit transactions, 279
 implicit transactions, 281
 MARS support, 71
 SQLNCLI as replacement for, 6
 SQLNCLI connections, 11

OLTP. *See* Online Transaction
 Processing (OLTP)

Online Analytical Processing (OLAP)
 characteristics, 522
 FILLFACTOR for, 402
 vs. OLTP, 521–522

online index operations, 252, 391, 393

ONLINE option, 391

Online Transaction Processing (OLTP)
 blocking problem indicators, 492
 characteristics, 522
 FILLFACTOR for, 402
 with HTTP endpoints, 14
 vs. OLAP, 521–522

Open Data Services, 64

Open Database Connectivity (ODBC)
 connection pooling, 31
 considerations, 6–7
 cursors, 136
 errors and messages, 216
 explicit transactions, 279
 implicit transactions, 281
 MARS support, 71
 SQLNCLI as replacement for, 6
 SQLNCLI connections, 10

OPEN statement, 20

OPENROWSET function, 536, 540
 OPENXML function, 99, 113
 operating system commands, 176
 operators
 computed columns, 381
 XQuery, 101–102
 optimistic concurrency mode, 242
 optimization. *See* performance
 OPTIMIZE FOR option, 476, 479, 481
 option argument, 204
 OPTION (MAXDOP) hint, 469
 or operator, 102
 ORDER BY statement
 clustered indexes, 373
 XQuery, 102
 ORIGINAL_LOGIN function, 182
 OSTRESS tool, 494–495
 outdated cached data, 27
 overflow, row-overflow data, 545–549
 overhead, monitoring tools, 457
 OWNER option, 184

P

PacketSize property, 31
 PAD_INDEX option, 391, 402
 page_count column, 401
 Page Free Space (PFS) pages, 367
 page level locks, 276
 Page life expectancy counter, 483
 Page/sec counter, 485
 pages, 366–367, 520
 PAGLOCK locking hint, 274
 parallel execution plans, 334
 parallelism, CPU, 468–471
 parameterized user-defined
 messages, 221
 parameterizing XML queries,
 119–121

parameters
 plan reuse, 89
 queries, 135
 sniffing, 91
 parsing phase, 312, 314, 314
 partial results, caching, 27
 \$partition function, 427
 partitioning, 426–427
 example, 428–433
 for manageability, 433
 operation, 427
 for performance, 427–428
 scaleout using, 438–439
 tables, 529
 Password property, 31
 passwords
 ADO.NET, 31
 SQL injection attacks, 177–178
 path expressions, 100
 PATH secondary XML indexes,
 378, 380
 PATINDEX function, 535–536
 Patterns and Practices team, 19
 peer-to-peer replication, scaleout using,
 436–438, 437
 performance, 97, 305–306,
 451–452
 blocking problems. *See*
 blocking problems
 CPU problems. *See* CPU problems
 cursors, 20–24, 351–352
 default trace tool, 454–456
 as design goal, 187
 exam essentials, 353–354, 500
 fragmentation, 369, 399–400
 detecting, 401
 fixing, 402
 preventing, 402–405
 indexes. *See* indexes
 I/O activity, 461–463, 485–487

- join operations, 458–459
- locking. *See* locks and locking
- memory problems, 483–485
- miscellaneous factors, 499
- modifying queries for, 347–353
- monitoring, 456–457
- partitioning for, 427–428
- plans. *See* plans
- problem causes, 457–458
- query life cycle
 - high-level view, 306–307, 307
 - logical processing, 308–311
 - physical processing, 312–314, 312, 314
- review questions, 355–363, 501–507
- stored procedures, 179, 454
- summary, 353, 499–500
- testing, 336–340, 339
- tools, 452–454
- Performance Monitor, 453
 - for caching, 33
 - System Monitor setup, 144–145
- permissions
 - custom, 185–186
 - setting, 180–181
- PERSISTED keyword, 527
- pessimistic concurrency mode, 242
- PFS (Page Free Space) pages, 367
- phantom reads, 251
- phone number validation, 171
- physical query processing, 312–314, 312, 314
- plan cache
 - aging in, 93–94
 - compiled plans and execution contexts, 307
 - queries, 87–88
- Plan Cache\Cache Hit Ratio counter, 467
- plan_generation_num column, 472
- plans
 - CPU problems, 476–482
 - vs. execution contexts, 307
 - graphical display, 325–326, 326
 - guides, 481
 - index scans, 333–334
 - inefficient
 - improving, 476–482
 - uncovering, 334–336
 - parallel, 334
 - red items in, 334
 - reuse, 88
 - ad hoc queries, 88–89
 - batches, 92
 - best practices, 92
 - for CPU problems, 467–468
 - simple parameters, 89
 - sp_executesql, 89–91
 - stored procedures, 91–92
 - SET statements, 315–316
 - SET SHOWPLAN_ALL, 318–319
 - SET SHOWPLAN_TEXT, 316–317
 - SET SHOWPLAN_XML, 319–322, 321
 - SET STATISTICS IO, 323–324
 - SET STATISTICS_PROFILE, 322–323
 - SET STATISTICS TIME, 324–325
 - SET STATISTICS XML, 322
 - Showplan for, 94–97
 - sorts, filters, and hash operators in, 336
 - SQL Profiler for, 326–330, 328–330
 - table scans, 331–333
 - terminology, 315
 - thick arrows in, 334
- plus signs (+)
 - regular expressions, 214
 - XQuery, 102

PooledConnectionLifetime property, 31

positioned cursor updates, 21

precision, vardecimal, 530

predicates in plan performance, 335

prejoined tables, 527–528

prepare phase for distributed transactions, 286–287

PRIMARY KEY constraint, 391

primary XML indexes, 367, 377–378

principle of least privilege, 180

EXECUTE AS, 181–186

permissions, 180–181

SQL injection attacks, 178

PRINT statement, 214–216

privileges, SQL injection attacks, 174

problematic queries, identifying, 462–463

procedure cache

- aging in, 93–94
- compiled plans and execution contexts, 307
- queries, 87–88

procedures. *See* stored procedures

Process Info page, 342

processor problems. *See* CPU problems

Processor Queue Length counter, 471

% Processor Time counter, 471

programmatic identifiers (PROGIDs), 8

prologs in XQuery expressions, 102

properties, object, 54–55

PROPERTY secondary XML indexes, 378, 380

providers in data access architecture, 5

Q

queries, 85–86

- ad hoc, 88–89
- cached batches, 93–94

- exam essentials, 123
- life cycle
 - high-level view, 306–307, 307
 - logical processing, 308–311
 - physical processing, 312–314, 312, 314
- MARS, 71–72, 347
- notifications, 34
- performance. *See* performance plans. *See* plans
- procedure cache, 87–88
- review questions, 124–132
- summary, 122–123
- well-performing applications, 424
- XQuery. *See* XQuery language

Query Editor, 105–106

query governor configuration

- option, 97

query method, 106

Query Notifications, 442

Query Optimizer, 271–272

Query Processor, 308–311

queues, runnable, 465

QUOTED_IDENTIFIER option, 386

quotes (") in XQuery, 100

R

RAISERROR function, 188

- vs. PRINT, 214
- syntax, 201–206
- with TRY...CATCH, 192

range scans

- I/O activity from, 461–463
- problem indicators, 486

Read Committed isolation level

- examples, 257–259
- facts, 256–257
- usage, 257

- Read Committed Snapshot
 - isolation level
 - example, 262–263
 - facts, 264–265
 - usage, 265–266
- READ_COMMITTED_SNAPSHOT
 - option, 256–257, 274
- Read Uncommitted isolation level
 - examples, 254–256
 - facts, 253
 - usage, 254
- READCOMMITTED hint, 274
- READCOMMITTEDLOCK hint, 257, 274
- READPAST hint, 274–275
- READTEXT statement, 72, 538
- READUNCOMMITTED hint, 272–273
- real data type, 512–513
- real-world applications, OLTP and OLAP, 522
- REBUILD option, 393
- rebuilding indexes, 393, 402
- REBUILT option, 393
- RECEIVE statement, 72
- recompilations
 - CPU problems, 471–476, 474–475
 - statement-level, 468
- RECOMPILE hint, 476
- RECOMPILE option, 92, 468, 476, 479
- recompiling
 - best practices, 92
 - errors in, 195–196
 - stored procedures, 26, 91–92
- records, 520
- recordsets, 22
- recreating indexes, 402
- red plan items, 334
- redundant columns, 525–526
- reference data, 425
- references
 - assemblies, 50–51, 56–57
 - connection, 10–11
 - SQL management objects, 48–50, 49
- refresh strategy for cached data, 32–34
- RegExMatch function, 172
- regular expressions, 173, 213–214
- relational data systems, 134–135, 511
 - result set-based operations, 135–136
 - row-based operations, 136
 - XML queries, 119–121
- relational operator trees, 313
- relative XPath expressions, 100
- reliability
 - as design goal, 187
 - input validation for, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
 - SODA messaging, 441–442
- remainder operator, 102
- REMOTE_PROC_TRANSACTION
 - option, 294
- REORGANIZE option, 393
- reorganizing indexes, 393, 402
- Repeatable Read isolation level
 - example, 260
 - facts, 259
 - usage, 259
- REPEATABLE_READ locking hint, 275
- replace value of keyword, 116–117
- REPLICATE function, 537
- replication management objects (RMO), 18–19, 44, 58
 - replication management tasks, 59–61
 - replication server connections, 58–59
- replication server connections, 58–59
- ReplicationServer objects, 59

required characters in input validation, 171
 resource data, 425
 Resource Locking Statistics by Logins report, 344
 Resource Locking Statistics By Objects report, 344
 resources, MARS, 75–76
 Restore class, 16
 result set-based operations, 135–136
 result sets, 135
 client servers, 22
 clustered indexes, 373
 RETURN statement, 102
 reusing plans. *See* plans
 REVERT statement, 182
 ring buffers, 214
 RMO (replication management objects), 18–19, 44, 58
 replication management tasks, 59–61
 replication server connections, 58–59
 Rollback method, 279
 ROLLBACK TRANSACTION statement, 279, 291–293
 ROLLBACK WORK statement, 279, 291
 RollbackTrans method, 279
 rolling updates, 436
 root nodes, 370, 370
 root pages, 372–373
 row-based operations, 136, 138–139
 row-by-row processing, 135
 row lock waits value, 492–493
 row-overflow data, 535, 545–549
 ROW_OVERFLOW_DATA allocation unit, 545
 row width in denormalization, 529–530
 @@ROWCOUNT function, 190–191
 ROWLOCK locking hint, 275

rows, 520
 caching, 32
 size limits, 535
 versioning, 69–70, 252
 rowsets
 client servers, 22
 converting XML data into, 109–113
 runnable queues, 465
 running operating system commands, 176

S

S (shared) locks, 244, 246
 SAC (Surface Area Configuration) tool, 45, 496
 same-session data caching, 27
 SAMPLED scan mode, 401
 SANs (storage area networks), 435
 SARGs (search arguments)
 nonclustered indexes, 376
 queries with, 350–351
 SAVE TRANSACTION statement, 289
 savepoints, 289–291
 saving
 @@ERROR value, 189
 @@ROWCOUNT value, 191
 SHOWPLAN_XML pLANs, 96–97
 scalability as design goal, 187
 scalable shared databases, 435–436, 436
 scaleout, 421–422, 434–435
 data-dependent routing, 439, 440
 distributed partitioned views, 438–439
 exam essentials, 443
 factors, 426
 leverage. *See* leveraging applications

- linked servers and distributed queries, 438, 438
- partitioning, 426–433
- peer-to-peer replication, 436–438, 437
- review questions, 444–449
- scalable shared databases, 435–436, 436
- service-oriented data architecture, 440–443
- summary, 443
- scans
 - index, 333–334, 368
 - I/O activity from, 461–463
 - problem indicators, 486
 - table, 331–333, 368
- Sch-M (schema modification) locks, 245–246
- Sch-S (schema stability) locks, 245–246
- schedulers, 464–465
- schema for well-performing applications, 424
- schema locks, 245
- schema modification (Sch-M) locks, 245–246
- schema stability (Sch-S) locks, 245–246
- SCHEMABINDING option, 386
- scientific notation, 512
- scope abort errors, 194
- SCOPE_IDENTITY function, 220
- scopes, transaction, 278–279
 - autocommit, 280
 - batch-scoped, 282–286
 - bound, 286
 - distributed, 286–287
 - explicit, 279–280
 - implicit, 280–282
- scripting, SMO, 48
- SDK\Assemblies folder, 46
- search arguments (SARGs)
 - nonclustered indexes, 376
 - queries with, 350–351
- second normal form (2NF), 424, 524, 524
- secondary XML indexes, 367, 378–379
- securability as design goal, 187
- Secure Sockets Layer (SSL), 14
- SecurePassword property, 31
- SecureWorks provider, 172
- security
 - HTTP endpoints, 14
 - input validation, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
 - stored procedures, 179
- Security Configuration window, 295, 295
- SELECT statement, 69, 72
- SELF option, 184
- SERIALIZABLE hint, 271–272
- Serializable isolation level, 251, 260–261
- server management objects (SMO), 44, 47
 - assembly references, 56–57
 - automating tasks using, 50–51
 - features, 48
 - locking problems, 345–346
 - object model, 48–50, 49
 - for objects
 - creating, 53–55
 - modifying, 55–56
 - overview, 16–18
 - server information, 51–53, 57
- Server Networking Interface (SNI), 64–67, 65

- server-side components
 - cursors, 20–21, 140, 141, 144–145
 - dynamic, 153–156, 157
 - keyset, 149–152, 153
 - static, 146–148, 149
 - overview, 3
- ServerConnection class, 29, 51, 58–59
- ServerInstance property, 31
- servers
 - cursors, 20–21
 - object instantiation, 51–52
 - optimizing, 347–348
 - properties, 52–53, 57
 - replication, 58–59
 - SMO information, 51–53, 57
- Service Broker, 441–442
- service-centric applications, 442–443
- service interaction data, 426
- Service Manager, 45
- service-oriented architecture (SOA), 423
- service-oriented data architecture (SODA), 434–435, 440–441
 - features, 441–442
 - service-centric applications, 442–443
- sessions
 - bound, 286
 - execution context, 181–183
- sessions ID, 465
- set-based operations, 135–136
- set-based queries from single-row statements, 351–353
- SET IMPLICIT_TRANSACTIONS statement, 281
- SET LOCK_TIMEOUT statement, 273
- SET options
 - indexed views, 386
 - for plan reuse, 92
- set-oriented operations
 - appropriate use, 139
 - vs. cursors, 143
- SET QUERY_GOVERNOR_COST_LIMIT statement, 98
- SET SHOWPLAN_ statements, 94–95
- SET SHOWPLAN_ALL statement, 315, 318–319
- SET SHOWPLAN_TEXT statement, 315–317
- SET SHOWPLAN_XML statement, 315, 319–322, 321
- SET STATISTICS IO statement, 316, 323–324
- SET STATISTICS_PROFILE statement, 316, 322–323
- SET STATISTICS TIME statement, 316, 324–325
- SET STATISTICS XML statement, 315, 322
- SET TRANSACTION ISOLATION LEVEL statement, 253
- SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED statement, 257
- SET TRANSACTION ISOLATION LEVEL REPEATABLE READ statement, 259
- SET TRANSACTION ISOLATION LEVEL SERIALIZABLE statement, 261
- SET TRANSACTION ISOLATION LEVEL SNAPSHOT statement, 267
- SETERROR option, 204
- Setup executable, 46
- severity levels
 - error handling, 188–189
 - RAISERROR, 203
- shared databases, 435–436, 436
- Shared Global Allocation Map (SGAM) pages, 367
- shared (S) locks, 244, 246
- shared intent update (SIU) locks, 245
- shared with intent exclusive (SIX) locks, 245–246

- Showplan tool, 94–97
- SHOWPLAN_XML plans, 94–97
- signal_wait_time_ms column, 465–466
- signal waits, 465–466
- simple message management, 221
- Simple Object Access Protocol (SOAP), 12–14, 12
- simple plan use parameters, 89
- simplification of queries, 314
- SINGLE_BLOB option, 540
- SINGLE_CLOB option, 540
- SINGLE_NCLOB option, 540
- single-row statements, set-based queries from, 351–353
- SIU (shared intent update) locks, 245
- SIX (shared with intent exclusive) locks, 245–246
- slashes (/) in XPath expressions, 100
- small LOBs (SLOBs), 545
- smalldatetime data type, 513–514
- smallint data type, 511
- smallmoney data type, 513
- SMO. *See* server management objects (SMO)
- Smo.dll file, 47
- SmoEnum.dll file, 47
- Snapshot isolation level, 266
 - example, 268–269
 - facts, 266–268
 - usage, 268
- SNI (SQL Server Network Interface), 64–67, 65
- SOA (service-oriented architecture), 423
- SOAP (Simple Object Access Protocol), 12–14, 12
- SODA (service-oriented data architecture), 434–435, 440–441
 - features, 441–442
 - service-centric applications, 442–443
- SORT_IN_TEMPDB option, 391, 393
- Sorts operator, 336
- sp_addmessage procedure, 222–224
- sp_altermessage procedure, 222, 224
- sp_configure procedure, 97
- sp_control_plan_guide procedure, 482
- sp_create_plan_guide procedure, 481
- sp_cursor_list procedure, 24
- sp_db_vardecimal_storage_format procedure, 531
- sp_describe_cursor procedure, 24
- sp_describe_cursor_columns procedure, 24
- sp_describe_cursor_tables procedure, 24
- sp_doption procedure, 293
- sp_dropmessage procedure, 222, 225
- sp_estimated_rowsize_reduction_for_vardecimal procedure, 533
- sp_executesql procedure, 89–91, 139, 143
- sp_help procedure, 406, 517
- sp_helpindex procedure, 406
- sp_indexoption procedure, 276
- sp_lock procedure, 346
- sp_recompile procedure, 92
- sp_tableoption procedure, 531, 542–543
- sp_trace_ procedures, 453
- sp_updatestats procedure, 334
- spaghetti code, 190
- splitting tables, 528–529
- SQL-92 data type synonyms, 516–518
- sql:column function, 119–120
- SQL command (SQLCMD) utility, 3, 494
- SQL common language runtime (SQLCLR), 15, 442
- SQL Distributed Management Objects (SQL-DMO) API, 15, 346
- SQL injection, 171–172
 - preventing, 178–179
 - susceptibility to, 174–178, 175–176

- SQL management object model. *See*
 - server management objects (SMO)
- SQL Native Client (SQLNCLI), 3, 6, 9
 - ADO with, 11–12
 - connection references, 10–11
 - new feature support, 9–10
- SQL OS (SQL Server Operating System), 464–465
- SQL Server Agent alerts, 221, 226–229, 227–229
- SQL Server Configuration Manager, 45
- SQL Server Lock Manager, 74–75
- SQL Server Management Studio
 - graphical plan display, 325–326, 326
 - index metadata from, 406
 - locking problems, 342–344, 343–344
 - performance, 453
- SQL Server Network Utility, 45
- SQL Server Networking Interface (SNI), 64–67, 65
- SQL Server Operating System (SQL OS), 464–465
- SQL Server Profiler
 - blocking problems, 490–491, 491
 - caching, 33
 - cursors, 23
 - execution plans, 326–330, 328–330
 - locking problems, 341
 - monitoring recompilations, 473–476, 474–475
 - overhead, 457
 - performance, 453
- SQL Trace tool, 197, 453
- sql:variable function, 120–121
- sql_variant data type, 515, 545
- sqlcli.lib file, 45
- SqlClient API functions, 284–285
- SQLCLR (SQL common language runtime), 15, 442
- sqlcmd utility, 3, 494
- SQLEndTran function, 281
- SqlException class, 217
- SqlInfoMessageEventHandler
 - delegate, 217
- SQLNCLI (SQL Native Client), 3, 6, 9
 - ADO with, 11–12
 - connection references, 10–11
 - new feature support, 9–10
- sqlncli.h file, 45
- SQLOLEDB provider, 11
- SQLSetConnectAttr function, 281
- SQLXML, 8
- SSL (Secure Sockets Layer), 14
- StartTransaction method
 - explicit transactions, 279
 - implicit transactions, 281
- state argument, 203
- statement abort errors, 193
- statement-level recompilation, 468
- states, error messages, 188, 203
- static clustered indexes, 374
- static cursors, 136–137, 146–148, 149
- STATISTICS_NORECOMPUTE
 - option, 391
- StmtRecompile metric, 33
- storage. *See* data storage
- storage area networks (SANs), 435
- stored procedures
 - vs. ad hoc queries, 179
 - cursors, 24
 - error messages, 188
 - explicit transactions, 293
 - for index metadata, 406
 - MARS, 73
 - for performance, 179, 454
 - for plan reuse, 91–92
 - recompiling, 26, 91–92
- string functions vs. full-text indexes, 376

- strings
 - data types, 514
 - execution context, 186
 - input validation, 171
 - SARGs summary statistics, 350
 - XQuery, 100
- strong passwords, 178
- subqueries, joins from, 348–350
- subscriptions, RMO, 18, 60
- SUBSTRING function, 535–538
- subtraction, XQuery, 102
- Summary Reports, 343–344, 453
- Surface Area Configuration (SAC) tool, 45, 496
- synchronized updates, 436
- synchronizing subscriptions, 18
- synonyms, data type, 516–518
- syntax, XQuery, 99–100
- sys.dm_as_wait_stats view, 465
- sys.dm_db_index_operational_stats view, 345, 408, 460–461, 488, 492
- sys.dm_db_index_physical_stats function, 401, 408, 478
- sys.dm_db_index_usage_stats function, 408, 460–461, 463
- sys.dm_db_missing_index_columns function, 394, 408
- sys.dm_db_missing_index_details function, 394, 408
- sys.dm_db_missing_index_group_stats view, 394, 408
- sys.dm_db_missing_index_groups view, 394, 408
- sys.dm_exec_cached_plans function, 330, 467
- sys.dm_exec_connections view, 66
- sys.dm_exec_plan_attributes function, 330
- sys.dm_exec_query_memory_grants view, 484
- sys.dm_exec_query_optimizer_info view, 313
- sys.dm_exec_query_plan function, 315, 330–331, 470
- sys.dm_exec_query_resource_semaphore view, 484
- sys.dm_exec_query_stats function, 330, 466, 472–473
- sys.dm_exec_requests view, 331
- sys.dm_exec_sql_text function, 330, 466, 470
- sys.dm_os_buffer_descriptors view, 457
- sys.dm_os_memory_clerks view, 483
- sys.dm_os_memory_objects view, 483
- sys.dm_os_ring_buffers view, 214
- sys.dm_os_wait_stats view, 344, 469, 486, 493
- sys.dm_os_waiting_stats view, 466
- sys.dm_os_waiting_tasks view, 345
- sys.dm_tran_locks view, 243, 246, 344, 346
- sys.indexes view, 407, 463
- sys.indexes_columns view, 407
- sys.messages catalog view, 222–223
- sys.stats_columns view, 407
- sys.syslockinfo view, 346
- sys.system_internals_allocation_units view, 370
- sys.xml_indexes view, 407
- System DSNs, 7
- system functions
 - cursors, 24–25
 - index metadata for, 408
 - performance, 453
- System Monitor, 453
 - cursors, 23–24
 - I/O performance problems, 485
 - locking problems, 341–342
 - memory problems, 483–484
 - overhead, 457
 - setting up, 144–145

T

T-SQL. *See* Transact-SQL (T-SQL)

table data type, 516

table level locks, 276

table names, red, 334

table scans, 331–333, 368

- vs. index seeks, 373
- I/O activity from, 461–463
- problem indicators, 486

tables, 366

- defined, 511
- duplicate, 528
- elements, 520
- prejoined, 527–528
- splitting, 528–529

TABLOCK hint, 275

TABLOCKX hint, 276

Tabular Data Stream (TDS), 12

- endpoints, 64–66, 65
- packets, 44, 64

tags, versioning, 521

Task Manager, 453

tempdb database, 499

testing performance, 336–340, 339

text, error message, 188

text columns for indexes, 539

text data type, 514, 535

text/image pages, 367

text in row option, 543

thick arrows in plans, 334

third normal form (3NF), 424, 524, 524

threads, worker, 464–465

time

- data types, 513
- statistics, 324–325

time stamps in caching, 27

timeouts, 484

timestamp data type, 516–517

tinyint data type, 511

Top Transactions by Blocked Transactions Count report, 344

Top Transactions by Locks Count report, 344

@@TOTAL_READ function, 453

Trace Properties dialog box, 145, 328–329, 328–329, 474, 474

Trace tool, 197, 453

traces for performance, 453–456

tracking cursors, 22–25

@@TRANCOUNT function, 207, 291–293

Transact-SQL (T-SQL)

- cursors, 20
- endpoints, 3, 66–67
- explicit transactions, 279
- FLWOR, 102
- implicit transactions, 281
- statements, 293

transaction abort errors, 194

transaction sequence number (XSN), 265

transactional replication

- relationships, 437

transactions, 239–240

- best practices, 296–297
- characteristics, 241–242
- concurrency

 - control modes, 242
 - side effects, 247–252

- distributed

 - phases, 286–287
 - working with, 294–296, 295

- error handling, 206–209, 294
- exam essentials, 297
- explicit. *See* explicit transactions
- isolation levels. *See* isolation levels
- locking. *See* locks and locking
- review questions, 298–304
- row versioning framework, 252

- scopes, 278–279
 - autocommit, 280
 - batch-scoped, 282–286
 - bound, 286
 - distributed, 286–287
 - explicit, 279–280
 - implicit, 280–282
 - summary, 297
- Transfer class, 16
- transparent scaleout, 434–435
- tree structures, 370, 370
- triggers
 - DDL and DML, 184
 - row versioning framework, 252
- trivial plans, 313
- TRUNCATE TABLE statement, 185
- TruncateOrders procedure, 185–186
- TRY...CATCH construct,
 - 191–193, 294
 - catchable errors, 193–194
 - deadlocks, 209–213
 - managed code, 213–214
 - RAISERROR, 205–206
 - transactions, 206–209
 - unhandled errors, 194–197

U

- U (update) locks, 244, 246
- UIX (update intent exclusive) locks, 245
- uncommitted dependency, 249–250
- unhandled errors, 194–197
- Unicode characters, 514
 - character strings, 515
 - XQuery, 100
- UNIQUE constraint, 391
- unique indexes, 367, 374, 391
- uniqueidentifier data type, 516
- unused indexes, 463–464
- update (U) locks, 244, 246

- update intent exclusive (UIX) locks, 245
- UPDATE statement
 - fragmentation, 402
 - MARS, 69
 - with server cursors, 20
 - WRITE clause, 538
- updates
 - lost, 247–249
 - scaleout factor, 426
- UPDLOCK locking hint, 276
- USE PLAN hint, 480–481
- user-defined messages, 188, 221–222
 - agent alerts, 226–229, 227–229
 - FORMATMESSAGE, 226
 - sp_addmessage, 223–224
 - sp_altermessage, 224
 - sp_dropmessage, 225
 - sys.messages catalog view, 222–223
 - xp_logevent procedure, 225
- user DSNs, 7
- user input validation, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
- uspGetAverageQuantity
 - procedure, 481
- uspGetAvgQtyNonModifiable
 - procedure, 482
- uspLogError procedure, 220

V

- ValidateFunctions class, 172
- validation, input, 170
 - principles, 170–171
 - process, 171–174
 - SQL injection, 174–179, 175–176, 179
- value method, 106–107, 109, 122, 380

VALUE secondary XML indexes, 378–380

varbinary data type, 515, 535
 row-overflow data, 545
 using, 536–541

varchar data type, 514, 534–535
 row-overflow data, 545
 using, 536–541

vardecimal storage format, 530–534

variable-length columns, 521, 530–534

variables
 instantiating, 53–54
 XQuery, 100

versioning tags, 521

vertical bars (|), operator trees, 317

vertical table splits, 528–529

VIEW DATABASE STATE
 permission, 33

VIEW SERVER STATE permission, 33

views
 catalog, 407
 indexed, 367, 385–389, 529
 partitioned, 438–439

W

wait types
 blocking problems, 493–494
 CXPACKET, 468–469
 I/O, 486–487

waits, signal, 465–466

warnings
 event notifications for, 197–200
 with TRY...CATCH, 192

weak passwords, 177

web-based access, 12, 12

Web Service Definition Language (WSDL) data, 441

web services
 SODA, 441
 SQL Server as, 13

web services layer caching, 27

well-performing applications, 423–425

WHERE clause
 data types in, 513
 LOBs, 535
 server cursors, 20
 XQuery, 102

WHERE CURRENT OF clause, 20

Windows authentication, 51

Windows Event Log, 218

Windows Management
 Instrumentation (WMI), 16, 46, 48

WITH LOG option, 203–204

WITH MARK option, 287–288

WITH NO_WAIT option, 265

WITH NORECOVERY option, 287

WITH NOWAIT option, 204

WITH RECOMPILE option, 92, 468, 476, 479

WITH SCHEMABINDING option, 386

WITH SETERROR option, 204

WITH STOPATMARK option, 287

WMI (Windows Management
 Instrumentation), 16, 46, 48

worker threads, 464–465

workspace memory grants, 484

WorkstationId property, 31

WSDL (Web Service Definition
 Language) data, 441

X

X (exclusive) locks, 244, 246

XACT_ABORT option, 207–208, 294, 296

XACT_STATE function, 206–207, 294

XLOCK locking hint, 276

XML
 in caching, 32
 Data Modification Language, 113–119

632 xml data type – yield points

indexes, 367, 377
 primary, 377–378
 secondary, 378–379
plans, 319–322, 321
queries, 105–106
 binding and parameterizing,
 119–121
 guidelines, 121–122
 XQuery. *See* XQuery language
 rowset data from, 109–113
xml data type, 516
XMLA (XML for Analysis)
 messages, 61
XmlReader object, 348
xp_cmdshell procedure, 176
xp_logevent procedure, 222, 225
XPath expressions, 100

XQuery language, 86, 99
 expressions, 100–102
 FLWOR statements, 102–105
 functions, 101
 operators, 101–102
 syntax, 99–100
 with XML data, 106–113
XSN (transaction sequence
 number), 265
XSS (cross-site scripting)
 vulnerability, 171

Y

yield points, 68