

SYBEX Book Index

MCSE: SQL Server™ 2000 Design Study Guide

Marc Israel
J. Steven Jones

Index

Copyright © 2001 SYBEX Inc., 1151 Marina Village Parkway, Alameda, CA 94501. World rights reserved. No part of this publication may be stored in a retrieval system, transmitted, or reproduced in any way, including but not limited to photocopy, photograph, magnetic or other record, without the prior agreement and written permission of the publisher.

ISBN: 0-7821-2942-0

SYBEX and the SYBEX logo are either registered trademarks or trademarks of SYBEX Inc. in the USA and other countries.

TRADEMARKS: Sybex has attempted throughout this book to distinguish proprietary trademarks from descriptive terms by following the capitalization style used by the manufacturer. Copyrights and trademarks of all products and services listed or described herein are property of their respective owners and companies. All rules and laws pertaining to said copyrights and trademarks are inferred.

This document may contain images, text, trademarks, logos, and/or other material owned by third parties. All rights reserved. Such material may not be copied, distributed, transmitted, or stored without the express, prior, written consent of the owner.

The author and publisher have made their best efforts to prepare this book, and the content is based upon final release software whenever possible. Portions of the manuscript may be based upon pre-release versions supplied by software manufacturers. The author and the publisher make no representation or warranties of any kind with regard to the completeness or accuracy of the contents herein and accept no liability of any kind including but not limited to performance, merchantability, fitness for any particular purpose, or any losses or damages of any kind caused or alleged to be caused directly or indirectly from this book.

SYBEX Inc.
1151 Marina Village Pkwy.
Alameda, CA 94501
USA
Phone: 510-523-8233
www.sybex.com

Index

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

Symbols

@@ERROR system function, 393–394
@@IDENTITY global variable, 158–160
.LDF log files, 67
.MDF primary data files, 67
.NDF secondary data files, 67

Numbers

1NF (First Normal Form), 32–35, 32, 33, 34
2NF (Second Normal Form), 35–36, 36
3NF (Third Normal Form), 36–39, 38
4NF (Fourth Normal Form), 39–40
5NF (Fifth Normal Form), 40

A

access to databases, **659–664**
 adding new users, 660–662, 660, 661, 662
 adding roles and users, 663–664
 basics, 659–660
 user-defined roles, 662–663
accessing data. *See* data, accessing
accessing servers. *See under* SQL Server 2000
ACID properties, 396
AFTER triggers, 407–408
aggregates
 aggregate operators, 464–468
 exercise using, 468–469
aliases
 ease in reading code and, 458
 specifying, 450–451

allocation
 data allocation, 99–103, 99
 extent allocation, 96–98
 tracking, 103–107, 104
alternate keys, 16
ANSI NULL default option, 155
applications
 application access and design, 129
 application roles, 664–666
 Decision Support System (DSS), 129
 OnLine Transaction Processing (OLTP), 129
ARITHABORT, 328
artificial keys, 17
atomic columns, 29
attributes, *See also* keys
 datatype attributes, 23–24
 defined, 6, 7
 defining, 7–9
autoparametization, 745

B

B, meaning of in SQL Server, 291
B-Trees, 199
back up issues, and filegroups, 127
Balanced Trees, 199
batch optimization, 702–703
batches defined, 382–383
BCNF (Boyce/Codd Normal Form), 39
BCP (bulk copy program), 588–607, *See also*
 bulk copy
 BCP syntax, 589–594
 BCP utility, 600–601
 format files, 594–600
binary attributes datatypes, 23–24

BINARY BASE64 option, 503
binary data, isolating, 126
Binary Large objects. *See* BLOBs (Binary Large Objects)
binary strings, 182
BLOBs (Binary Large Objects)
 filegroups and, 203
 table storage and, 182
 text in row OFF, 199–202, 200, 201, 202
 text in row ON, 202–203
blocking, 634–635
Boyce/Codd Normal Form (BCNF), 39
bulk copy
 BCP (bulk copy program), 588–607
 BCP syntax, 589–594
 BCP utility, 600–601
 format files, 594–600
 BULK INSERT, 602–607
 BULK INSERT syntax, 602–605
 using, 606–607
business rules, 28

C

candidate keys, 16
cascading deletes
 defined, 258
 enforcing, 260
 triggers and, 260
cascading updates
 defined, 258
 enforcing, 260
 triggers and, 260
CD-ROM that comes with this book,
 xxxiv–xxxvi, xxxv
character attributes datatypes, 23
character strings, 181
check rules, 234–240
 check constraints, 235–239
 rules, 239–240
checkpoint process, 72
Chen, Peter
 definition of entities and, 7
 Entity/Relationship model and, 5
clustered indexes
 accessing data and, 299–302, 300, 301

 basics of, 291–294, 292, 293
 creating, 321–324, 322
 optimizing queries and, 719
Codd, E. F., 28
collation
 changing in columns, 172–173
 creating databases and, 79–80
 table character columns and, 163–165, 165
 Unicode vs., 164
columns
 adding, 174–176
 altering, 171–174
 collation of character columns, 163–165, 165
 computed
 basics, 188–190
 creating indexes on, 326–328
 creating statistics on multiple, 331–332
 default values, 227–234
 basics, 227–228
 default constraints, 228–231, 230
 default objects, 232–234
 derived columns, 45–46, 45
 dropping, 174–176
 granting rights to, 674–675
 GUID columns, 162–163
 identity columns defined, 540
 inserting rows with selected, 538–539
 names, and creating tables, 153
 partitioning columns, 373–375
 redundant columns, 42, 43–45
 relational tables and, 29–30
 system datatypes in, 176–184
 date and time, 181
 listed, 177–180
 numeric, 180–181
 special, 182–183
 strings, 181–182
 synonyms, 183–184
 text, ntext and image column storage,
 198–203, 200, 201, 202
 updating multiple, 548–550
 user-defined datatypes in, 184–187
command line utilities defined, 589
compatibility levels
 basics, 118–120
 in SQL Server 2000, 383
compilation, stored procedures and, 387–389

- composite indexes
 - creating, 325–326
 - described, 296
- COMPUTE clause, 467–468
- conceptual design
 - aspects of, 28
 - importance of, 3
 - vs. relational model, 5
- connections, 609–610
- constraints, 659–664, *See also* foreign keys
 - default constraints, 228–231
 - defining at table creation, 228–230, 230
 - defining for existing tables, 230–231
 - unique constraints, 246–252
 - defining at table creation, 247–251, 249
 - defining at table modification, 251–252
 - vs. default objects, 233
 - vs. triggers, 25
- correlated subqueries, 473
- Create Database statement
 - options, 73–94
 - automatic shrinking, 80–81
 - collation, 79–80
 - filenames, 73–74
 - manual shrinking, 82–86, 82, 85
 - size and growth, 74–79, 77, 78
 - script, 68
- CREATE INDEX statement, 320–322
- creation scripts, and Generate SQL Scripts
 - utility, 69
- cross joins, 462
- crow’s feet, 12, 12
- CUBE and ROLLUP, 466–467
- cursors, 480–489
 - basics, 480
 - cursor declarations, 484–488, 555–556
 - modifying data inside, 555–558
 - retrieving data from, 488–489
 - using, 481–484
- extracting in XML format, 507
- functions for transforming, 476–479
- masking complexity with views, 361
- modifying using XML, 559–568
 - modifying rows with OPENXML, 564–567
 - using OPENXML, 559–563
- placement and database performance, 129–130
- retrieving using views, 368–369
- updating
 - using cursors, 556–558
 - using views, 368–369
- data access, analyzing and optimizing, 698–767
 - analyzing queries, 703–716
 - execution plan analysis, 709–714
 - exercise in, 714–716
 - SET FORCEPLAN, 704
 - SET NOEXEC, 704–705
 - SET SHOWPLAN_ALL, 705–707
 - SET statements basics, 703–704
 - SET STATISTICS IO, 708
 - SET STATISTICS PROFILE, 709
 - SET STATISTICS TIME, 709
 - SETSHOWPLAN_TEXT, 707–708
 - exam essentials, 757–758
 - key terms, 757
 - optimizing queries, 717–732
 - changing an index, 723–726, 724, 726
 - indexes and, 717–721
 - limiting the result set, 721–723, 722, 723
 - query hints and, 726–731
 - optimizing stored procedures and triggers, 744–746
- Profiler, 731–744
 - basics, 731–732, 732
 - creating traces using, 738–739, 739
 - functions and use, 733–738, 735, 736, 737
 - Index Tuning Wizard, 741–744
 - lock detection, 740–741
 - replaying saved trace files, 739–740
- query optimizer, 699–703
 - basics, 699–700
 - batch optimization, 702–703
 - single statement optimization, 700–702
- review questions and answers, 758–767
- stored procedure debugger, 746–755
 - exercise in debugging, 749–755, 752, 753, 754

D

- data, *See also* datatypes
 - allocation, 99–103, 99
 - allocation tracking, 103–107, 104
 - binary data, isolating, 126

- using, 746–749, 748
- summary, 755–756
- data, accessing, 442–532
 - cursors, 480–489
 - declarations, 484–488
 - definition and uses, 480
 - retrieving data from, 488–489
 - using, 481–484
 - data transforming functions, 476–479
 - distributed queries, 489–500
 - about, 489–490
 - distributed partitioned views, 496–500, 497
 - linked servers, 490–494
 - OPENROWSET function, 494–496
 - exam essentials, 516–517
 - heterogeneous, 495–496
 - indexes and, 297–306
 - clustered indexes, 299–302, 300, 301
 - heaps, 297–299, 298, 299
 - nonclustered indexes, 302–306, 304, 305
 - key terms, 518
 - review questions and answers, 519–532
- SELECT statement and, 443–475
 - aggregate operators, 464–465
 - COMPUTE clause, 467–468
 - cross joins, 462
 - CUBE and ROLLUP, 466–467
 - DISTINCT command, 470–471
 - exercise in joining tables, 463–464
 - exercise in using aggregates, 468–469
 - exercise in using queries, 453–455
 - GROUP BY clause, 465
 - HAVING clause, 466
 - inner joins, 456–458
 - limiting returned results, 451–453
 - ORDER BY clause, 446–447
 - outer joins, 458–461
 - scalars, 471–472
 - single table SELECT, 444–445, 445
 - specifying aliases, 450–451
 - subqueries, 472–475
 - UNION command, 470
 - WHERE clause, 447–449
- SQL Server 2000 and, 442–443
- summary, 515–516
- XML and, 500–514
 - access through URLs using template files, 508–511
 - access through URLs using XPATH, 514
 - access through URLs using XSL, 511–513, 513
 - description, 500–501
 - exercise in extracting data in XML format, 507
 - using SELECT with, 502–506
- data files, *See also* filegroups; files
 - databases and
 - adding to, 92, 124
 - creating databases and, 66–72, 66
 - removing files, 93
 - defragmenting, 339–341
 - filenames, and database management, 73–74, 73
 - placement
 - data placement, 129–130
 - default placement, 67
 - placement and performance, 127–131
 - shrinking, 80–90
 - automatically, 80–81
 - files directly, 87–90, 89
 - manually, 82–86, 82, 85
 - with Enterprise Manager, 86–87, 90
 - size and growth options, 74–75
 - space management and, 94–107, 95, 96, 99, 104
 - allocation tracking, 103–107, 104
 - data allocation, 99–103
 - extents, 95–98, 95, 96
 - pages, 98–99, 99
 - types, 67
 - vs. log files, 107
- data, importing and exporting, 588–624
 - bulk copy program (BCP), 588–601
 - BCP syntax, 589–594
 - BCP utility, 600–601
 - format files, 594–600
 - BULK INSERT, 602–607
 - BULK INSERT syntax, 602–605
 - using, 606–607
 - Data Transformation Services (DTS), 607–617
 - basics, 608–609, 608
 - connections, 609–610

- packages, 608, 612–613
 - tasks, 610–612
 - using, 613–617, 614, 615, 616
- exam essentials, 618
- key terms, 618
- review questions and answers, 619–624, 622
- summary, 617–618
- data integrity, 222–286
 - basics, 222
 - check rules, 234–240
 - check constraints, 235–239
 - rules, 239–240
 - default values, 227–234
 - basics, 227–228
 - default constraints, 228–231, 230
 - default objects, 232–234
 - exam essentials, 268
 - foreign key constraints, 252–267, 252
 - cascading foreign key constraints, 260–265, 264, 265
 - creating and using exercise, 257–258
 - defining at table creation, 253–260
 - defining at table modification, 266–267
 - real world scenario, 259
 - implementing, 223–227
 - key terms, 268
 - primary keys, 240–246
 - basics, 240–241
 - defining at table creation, 241–245, 244
 - defining at table modification, 245–246
 - review questions and answers, 269–286, 269, 271, 273, 278, 279
 - summary, 267
 - types, 22–28
 - domain integrity, 22–24, 222
 - enterprise integrity, 28, 223
 - entity integrity, 24, 223
 - referential integrity, 24–28, 223
 - unique constraints, 246–252
 - defining at table creation, 247–251, 249
 - defining at table modification, 251–252
- data, modifying, 534–586
 - DELETE statement, 551–554
 - exam essentials, 574
 - INSERT statement, 535–544
 - multiple rows, 542–544
 - single rows insert, 535–538
 - single rows with identity fields, 540–541
 - single rows with selected columns, 538–539
 - inside cursors, 555–558
 - key terms, 573
 - review questions and answers, 574–586
 - summary, 573
 - UPDATE statement, 544–551
 - all rows, 545–546
 - multiple columns, 548–550
 - sets of rows, 547–548
 - views, 550–551
 - using distributed queries, 568–572
 - using linked servers, 569
 - using OPENROWSET, 570–572, 571, 572
 - using OPENXML, 559–568
 - deleting rows, 566–567
 - inserting rows, 564–565
 - new rowset views, 561–564
 - SP_XML_PREPAREDDOCUMENT, 559–561
 - updating rows, 565–566
- Data Read and Write Log process, 70, 70
- Data Transformation Services (DTS), 607–617
 - basics, 608–609, 608
 - connections, 609–610
 - packages, 608, 612–613
 - tasks, 610–612
 - using, 613–617, 614, 615, 616
- database logical modeling, 2–61
 - denormalization process, 41–47, 41, 42
 - adding derived columns, 45–46, 45
 - adding redundant columns, 43–45, 44
 - partitioning tables, 46–47, 46
 - designing database systems, 2–5
 - Entity/Relationship (ER) database model, 5–28
 - basics, 5–7
 - defining entities and attributes, 7–10, 9
 - domain integrity, 22–24
 - enterprise integrity, 28
 - entity integrity, 24
 - keys, 16–21
 - referential integrity, 24–28
 - relationships, 10–16, 11, 12, 13, 14, 15, 16
 - exam essentials, 48–49
 - key terms, 49
 - relational models and normalization, 28–40
 - advanced normalization, 39–40

- First Normal Form (1NF), 32–35, 32, 33, 34
- relational tables, 29–31
- Second Normal Form (2NF), 35–36, 36
- Third Normal Form (3NF), 36–39, 38
- review questions and answers, 50–61, 50, 52, 55, 56, 58
- database objects, creating and maintaining, 357–440
- exam essentials, 421–422
- key terms, 421
- review questions and answers, 422–440, 423, 429, 435
- stored procedures, 378–395
 - creating and altering, 380–384
 - definition and advantages, 379–380, 380
 - error handling, 392–396
 - executing, 386–392
 - using parameters, 384–385
- summary, 420
- transactions, 396–398
- triggers, 406–420
 - AFTER, 407–408, 419
 - AFTER INSERT, 419
 - changing firing order, 419
 - definition and advantages, 406–407
 - DELETE, 410–412, 411
 - disabling, 419–420
 - INSERT, 409–410
 - INSTEAD OF, 413–418
 - multiple, 419
 - performance considerations, 418
 - UPDATE, 412–413, 412
 - Web, 419
- user-defined functions (UDFs), 399–405
 - creating and altering, 400–404
 - definition and advantages, 399–400
 - using, 404–405
- views, 359–378
 - creating and altering, 362–367
 - defined, 360, 360
 - indexed views, 369–372
 - partitioned views, 372–378, 372
 - uses, 361
 - using, 367–369
- database options, 113–120, 113
 - compatibility levels, 118–120
 - default ANSI NULL, 155
 - recovery model and, 116
 - table of options values, 114–115
- database physical modeling, 64–147
 - creating and managing databases. *See* databases, creating and managing
 - exam essentials, 133
 - file placement and performance, 127–131
 - filegroups, 120–127
 - basics, 120–121, 121
 - creating, 121–125, 122, 123
 - maintenance and performance, 125–127
 - key terms, 132–133
 - review questions and answers, 133–147
 - summary, 132
- databases
 - access, 659–664
 - adding new users, 660–662, 660, 661, 662
 - adding roles and users, 663–664
 - basics, 659–660
 - user-defined roles, 662–663
 - adding files and filegroups, 124
 - altering, 90–94
 - monitoring activity with Profiler, 733–734
 - names, changing, 93–94
 - options values, 114–115
 - shrinking, 80–90
 - automatically, 80–81
 - manually, 82–85, 82
 - using Enterprise Manager, 85–87, 85
 - size, increasing, 86–87
- databases, creating and managing, 64–120
 - creating simple databases, 65–72, 65
 - Create Database statement, 68
 - data and log files, 66–72, 66, 70
 - with Enterprise Manager, 65–66, 65, 76–79, 77, 78
 - data files and, 66–72, 66
 - database options, 113–120, 113
 - compatibility levels, 118–120
 - database options value, 114–115
 - recovery Model, 116
 - managing, 73–94
 - altering databases, 90–94
 - collation, 79–80
 - filenames, 73–74, 73
 - shrinking databases and files, 80–90, 82, 85, 89
 - size and growth options, 74–75

- space management, 94–112
 - data files, 94–107, 95, 96, 104
 - log files, 107–113, 109, 110, 111
- datatypes
 - attributes of, 23–24
 - changing in columns, 172
 - creating tables and, 154
 - system datatypes in columns, 176–184
 - date and time, 181
 - listed, 177–180
 - numeric, 180–181
 - special, 182–183
 - strings, 181–182
 - synonyms, 183–184
 - user-defined datatypes in columns, 184–187
- date and time, values in columns, 181
- DBCC LOG statement, 108–109
- DBCC SHOW CONTIG, 335–336
- DBCC SHRINKDATABASE, 88–89
- DBCC SHRINKFILE, options, 87–90
- deadlocks, 635–638
- Decision Support System (DSS) application, 129
- declarative integrity, 223
- default constraints, 228–231, 230
- default objects, 232–234
- default values, 227–234
 - basics, 227–228
 - default constraints, 228–231, 230
 - default objects, 232–234
- deferred name resolution defined, 383
- defragmenting data files, 339–341
- delete rules, 25
- DELETE statement, 551–554
- DELETE triggers, 410–412, 411
- deletes, and fragmentation, 338–339
- denormalization process, 41–47, 41, 42
 - derived columns, adding, 45–46, 45
 - redundant columns, adding, 43–45, 44
 - tables, partitioning, 46–47, 46
- derived columns, adding, 45–46, 45
- derived tables, 473
- design
 - conceptual design, *See also* Entity/Relationship (ER) database model
 - aspects of, 28
 - importance of, 3
 - vs. logical, 5
 - vs. relational model, 5

- logical design, *See also* relational models
 - and normalization
 - designing database systems, 2–5
- dirty pages defined, 72
- disks, and log files, 131
- DISTINCT command, 470–471
- distributed partitioned views, 496–500, 497
- distributed queries, 489–500
 - about, 489–490
 - distributed partitioned views, 496–500, 497
 - linked servers, 490–494
 - OPENROWSET function, 494–496
 - using to modify data, 568–572
 - linked servers, 569
 - OPENROWSET, 570–572, 571, 572
- distribution statistics, and indexes, 306–311
- domains
 - defined, 22
 - domain integrity, 22–24, 222
- DSS (Decision Support System) application, 129
- DYNAMIC cursors, 486

E

- edge tables, 563
- ELEMENTS option, 503
- ENCRYPTION option, 364
- enterprise integrity, 28, 223
- Enterprise Manager
 - altering files, 91–93
 - adding, 92
 - modifying, 91–92
 - removing, 93
 - cascading foreign key constraints, creating, 262–265, 263, 264, 265
 - collation and, 164, 165
 - columns
 - defining default values, 229–230, 230
 - warning on altering, 171
 - data and log files, defining, 67
 - databases
 - adding new users, 660–662, 660, 661, 662
 - creating, 65–66, 65, 76–79, 77, 78
 - options, 113–115, 113
 - shrinking, 85–87, 85
 - datatypes (user-defined), creating and managing, 186–187, 187

- filegroups
 - adding, 124
 - changing default, 123, 123
 - creating, 121, 123, 123, 168–169, 168
- files
 - adding, modifying and removing, 91–92, 93
 - shrinking, 88–90, 89
- locks, viewing, 639–641, 639, 640
- logins, adding, 657–659, 657, 658, 659
- rights, viewing, 670–671, 671
- tables
 - creating simple, 152, 152, 155–156, 155
 - defining placement, 169–170, 170
 - modifying, 171
 - unique constraints, defining, 248–249, 249
- entities
 - attributes and, 7–10
 - defined, 6
 - entity integrity, 24, 223
- entity integrity, 24, 223
- Entity/Relationship (ER) database model, 5–28, 21
 - basics, 5–7
 - defining entities and attributes, 7–10, 9
 - domain integrity, 22–24
 - elements of, 6
 - enterprise integrity, 28
 - entity integrity, 24
 - keys, 16–21
 - foreign keys, 19–21, 19
 - primary keys, 16–19, 18, 19
 - referential integrity, 24–28
 - relationships, 10–16
 - many-to-many relationships, 14–15, 15
 - one-to-many relationships, 13, 14, 14
 - one-to-one relationships, 13–14, 13
 - recursive relationships, 16, 16
- ER. *See* Entity/Relationship (ER) database model
- errors, handling with stored procedures, 392–396
- exam preparation
 - adaptive exam format, xxix
 - data
 - accessing, 516–532
 - accessing, analyzing and optimizing, 757–767
 - importing and exporting, 618–624
 - integrity, 268–286
 - modifying, 584–586
 - databases
 - logical modeling, 48–61
 - objects, creating and maintaining, 421–440
 - physical modeling, 133–147
 - indexes, 342–355
 - locking, 646–652
 - security plan, developing, 686–695
 - tables, creating and maintaining, 204–219
 - tips for taking the exam, xxxi–xxxii
 - types of questions, xxvi–xxxi, xxviii
 - Web sites for, xxxvii–xxxix
- exclusive locks, 632
- execution plans
 - analyzing query execution, 709–714
 - basics, 709–710
 - query operations, 710–714
 - query optimizer and, 701–702
 - stored procedures and, 386
- exercises, listed, xix–xx
- EXPLICIT mode, and XML, 505–506
- exporting. *See* data, importing and exporting
- extended properties
 - metadata and, 195
 - tables and, 190–195, 191, 194
- Extensible Markup Language(XML), *See also* XML Path Language
 - modifying data using OPENXML, 559–568
 - deleting rows, 566–567
 - inserting rows, 564–565
 - new rowset views, 561–564
 - SP_XML_PREPAREDDOCUMENT, 559–561
 - updating rows, 565–566
 - options, 503
 - SQL Server 2000 and, 500–514
 - description, 500–501
 - extracting data in XML format, 507
 - integrating using XPATH, 514
 - integrating using XSL, 511–513, 513
 - using SELECT with, 502–506
 - using template files to access SQL Server through a URL, 508–511
- Extensible Stylesheet Language (XSL)
 - defined, 501
 - extracting data in XML format and, 511–513, 513
- extents, 95–98, 95, 96
- external fragmentation, 333–336, 334

F

FAST_FORWARD cursors, 487

FETCH statements

retrieving data from cursors and, 488–489

using cursors and, 484

fields

field lengths, 598–600

field terminators, 598–600

filegroups

BLOBs and, 203

creating databases and, 120–127, 121

adding filegroups, 124

automatically creating filegroups, 121–124,
122, 123

backing up and, 127

creating tables and, 165–169

basics, 165–167, 166, 167

creating databases using new filegroup,
168–169, 168

defined, 120

maintenance and performance, 125–127

RAID and, 169

filenames, database management and, 73–74, 73

files, *See also* data files; filegroups; log files

database files

modifying, 91–92

placement and performance, 127–131

format files, 594–600

field lengths and field terminators, 598–600

prefix lengths, 598

using, 595–596

removing from databases, 93

fill factor, and indexes, 328–331, 329, 330

First Normal Form (1NF), 32–35, 32, 33, 34

fixed point numeric attributes, 23

floating point numeric attributes, 23

FOR REPLICATION option, 382

foreign keys, 252–267, 252

basics, 19–21, 19

cascading foreign key constraints, 260–265,
263, 264, 265

defined, 19, 19

defining at table creation, 253–260

cascading deletes and updates and, 258–260

column level constraints, 253–254

creating tables and, 257–258

table-level constraints, 254–256

defining at table modification, 266–267

Entity/Relationship model and, 19–21, 19, 21

exercise to create and use, 257–258

nonclustered indexes and, 720

real world scenario, 259

format files, 594–600

field lengths and field terminators, 598–600

prefix lengths, 598

using, 595–596

forms. *See* normal forms

FORWARD_ONLY, 486–487

fragmentation and index maintenance, 332–341

defragmenting data files, 339–341

deletes and, 338–339

fragmentation types, 333–336, 333, 334

inserts and, 336–337, 337

updates and, 337–338, 338

FROM clause, subqueries in, 473

full functional dependency defined, 31

full outer joins, 460–461

functions. *See* user-defined functions (UDFs)

G

GAMs. *See* Global Allocation Maps

Generate SQL Scripts utility, 69

Global Allocation Maps (GAMs)

defined, 98

table of usage, 105

global cursors, 485

globally unique identifiers. *See* GUIDs (globally
unique identifiers)

GROUP BY clause, 465

GUIDs (globally unique identifiers), *See also*

ROWGUIDCOL

creating and managing GUID columns,

162–163

defined, 161

H

hardware, importance of powerful, 128

HAVING clause, 466

heaps

- accessing data and, 297–299, 298, 299
- and indexes, 289–291, 289, 290
- defined, 289

hints

- index hints, 727
- join hints, 726–727
- lock hints, 728–731
- query hints, 726–731
- query processing, 727–728

horizontal partitions defined, 675

I

IAMs (Index Allocation Maps) defined, 98

IDENT_CURRENT, 158–160

identifiers, *See also* unique identifiers

- conforming column names to, 153
- identifier rules, 153

identifying relationship, 20–21

identity, *See also* unique identifiers

- defined, 156
- Identity, 163
- identity columns, 161, 540
- identity values, 158–160, 541

identity columns, 161, 540

identity fields, 540–541

identity values, 158–160, 541

IIS (Internet Information Server) defined, 501

image columns, storage, 198–203, 202

importing. *See* data, importing and exporting

Index Allocation Maps (IAMs) defined, 98

index hints, 727

Index Tuning Wizard, 741–744

indexed views, 369–372

indexes, 288–355

- accessing data and, 297–306
 - clustered indexes, 299–302, 300, 301
 - heaps, 297–299, 298, 299
 - nonclustered indexes, 302–306, 304, 305
- creating, 320–331
 - clustered and nonclustered, 321–324, 322
 - composite, 325–326
 - fill factor, 328–331, 329, 330
 - on computed columns, 326–328
 - unique, 324–325

definition and advantages, 288–296

- clustered, 291–294, 292, 293
- composite, 296
- heaps, 289–291, 289, 290
- nonclustered, 294–296, 294, 295
- unique, 296

exam essentials, 342–343

fragmentation and maintenance, 332–341

- defragmenting data files, 339–341
- deletes, 338–339
- fragmentation types, 333–336, 333, 334
- inserts, 336–337, 337
- updates, 337–338, 338

Index Tuning Wizard, 741–744

key terms, 342

optimizing queries and, 717–721

- basics, 717–718
- by changing, 723–729, 724, 726
- clustered indexes, 719
- nonclustered indexes, 719–721

physically separating from tables, 126

reindexing tables, 341

review questions and answers, 343–355, 343, 344, 349, 350

statistics and, 306–320

- creating statistics, 331–332
- distribution statistics, 307–311
- index choice, 312–315, 312, 315
- statistics maintenance, 315–320

summary, 342

inline table-valued UDFs, 402–403

inner joins, 456–458

insert rules, 24–25

INSERT statement, 535–538

- multiple rows, 542–544
- single rows, 535–538
- single rows with identity fields, 540–541
- single rows with selected columns, 538–539

INSERT triggers, 409–410, 409

inserts, and leaf level fragmentation, 336–337, 337

INSTEAD OF triggers

- INSTEAD OF DELETE, 415
- INSTEAD OF INSERT, 414–415
- INSTEAD OF UPDATE, 416

integrity. *See* data integrity

intent locks, 631–632

intermediate levels defined, 291
 internal fragmentation
 defined, 333, 333
 determining, 334–336
 negative effect of deletes on, 333, 338–339
 Internet Explorer 5, and viewing XML
 documents, 506
 Internet Information Server (IIS) defined, 501

J

join hints, 726–727
 joins
 cross joins, 462
 exercise in joining tables, 463–464
 inner joins, 456–458
 outer joins, 458–461
 full outer joins, 460–461
 left and right outer joins, 459–460

K

key range locks, 634
 keys, *See also* foreign keys; primary keys
 alternate keys, 16
 artificial keys, 17
 candidate keys, 16
 KEYSET cursors, 486

L

latches, 633
 lazy writer process, 72
 leaf levels
 defined, 291
 fragmentation and inserts, 336–337, 337
 inserts as cause of fragmentation, 336
 storage of computed columns and, 327
 left and right outer joins, 459–460
 legacy syntax, 460

linked servers
 accessing data and, 490–494
 modifying data with, 569
 literals. *See* scalars
 local cursors, 485
 lock hints
 queries and, 728–731
 servers and, 643–644
 locking, 626–652
 described, 626
 exam essentials, 647
 key terms, 646
 lock manager, 627–634
 basics, 627–629
 levels of locking, 633–634
 lock modes, 629–633
 options, 638–644
 lock hints and servers, 643–644
 lock isolation levels, 642
 viewing locks, 638–642, 639, 640, 641
 review questions and answers, 647–652
 summary, 645–646
 transactions and locking, 634–638
 locks
 exclusive locks, 632
 lock detection, 740–741
 lock hints and queries, 728–731
 SP_LOCK result set, 641–642
 types of, 631–633
 log files, *See also* transaction logs
 creating databases and, 66–72, 66
 default placement, 67
 functioning of, 69–70, 70
 placement and performance, 130–131
 placement, 130–131
 shrinking, 87–90, 89
 size and growth options, 74–75
 space management and, 107–113, 109, 110, 111
 log records and, 110–111
 size of, 109–110
 transaction logs, 107–108
 vs. data files, 107
 logical design. *See* database logical modeling
 login
 databases and, 660
 SQL Sever 2000 and, 655–659, 657, 658, 659

M

mainframes, real world scenario, 601–602
 managing databases. *See* databases, creating and managing
 manual shrinking, 82–86, 82, 85
 many-to-many relationships, 14–15, 14, 15
 materialized views In Oracle, 370
 MCSE (Microsoft Certified System Engineer)
 described, xxi
 exam requirements, xxiv–xxvi
 types of questions, xxvi–xxxi
 metadata
 defined, 365
 extended properties and, 195
 Microsoft Certified System Engineer. *See* MCSE (Microsoft Certified System Engineer)
 min LSN, 110–111, 110, 111
 mixed extents, 95, 96
 multistatement table-valued UDFs, 403–404

N

natural order defined, 444
 non-key attributes defined, 7, 9
 non-leaf levels defined, 291
 nonclustered indexes
 accessing data and, 302–306, 304, 305
 creating, 321–324, 322
 definition and advantages, 294–296, 294, 295
 optimizing queries and, 719–721
 normal forms, 31–39
 Boyce/Codd Normal Form (BCNF), 39
 First Normal Form (1NF), 32–35, 32, 33, 34
 Second Normal Form (2NF), 35–36, 36
 Third Normal Form (3NF), 36–39, 38
 Fourth Normal Form (4NF), 39–40
 Fifth Normal Form (5NF), 40
 normalization process
 advanced normalization, 39–40
 Boyce/Codd Normal Form (BCNF), 39
 Fourth Normal Form (4NF), 39–40
 Fifth Normal Form (5NF), 40
 defined, 30
 normal forms, 31–39

 First Normal Form (1NF), 32–35, 32, 33, 34
 Second Normal Form (2NF), 35–36, 36
 Third Normal Form (3NF), 36–39, 38
 normalized logical models, 38, 38
 NOT FOR REPLICATION
 constraints and, 265
 when defining identities, 161
 ntext columns, and table storage, 198–203, 202
 NULL values
 changing in columns, 173
 defining columns and, 154–155
 nullability, changing columns, 173–174
 numeric datatypes
 attributes datatypes, 23
 in columns, 180–181

O

Object Browser, 747
 object rights, 667–685
 basics, 667–671, 671
 ownership chains, 684–685
 stored procedures, 677–680
 tables, 671–675
 triggers, 680–681
 user-defined functions (UDFs), 681–684
 views, 675–677
 Object Search, 747
 objects, *See also* database objects
 default objects, 232–234
 OLTP (OnLine Transaction Processing)
 application, 129
 one-to-many relationships, 13, 14, 14
 one-to-one relationships, 13–14, 13
 OnLine Transaction Processing (OLTP)
 application, 129
 OPENROWSET function
 data accessing and, 494–496
 modifying data using, 570–572, 571, 572
 OPENXML for modifying data, 559–568
 basics, 561–564
 deleting rows, 566–567
 inserting rows, 564–565
 new rowset views, 561–564

SP_XML_PREPAREDDOCUMENT, 559–561
 updating rows, 565–566
 OPTIMISTIC option, and cursors, 487
 options. *See* database options; various options
 ORDER BY clause, 446–447
 OSQL, 79
 outer joins, 458–461
 ownership chains, 684–685

P

packages in DTS, 608, 612–617
 creating, 613–617, 614, 615, 616
 defined, 612
 pages
 dirty pages defined, 72
 types of, 98–99, 99
 parameters, and stored procedures, 384–385
 partitioned views, 372–378, 372
 advantages and disadvantages, 376–377
 conditions for, 375–376
 defined, 372
 distributed partitioned views, 496–500, 497
 horizontal partitions defined, 675
 partitioning columns, 373–375
 partitioning columns, 373–375
 partitioning tables, 46–47
 horizontal, 47
 vertical, 46–47, 46
 partitions, vertical partitions defined, 445, 445
 performance
 data placement and, 129–130
 disk performance, 112
 files
 data files placement and, 127–131
 filegroups maintenance and, 125–127
 log file placement and, 130–131
 performance enhancers
 SELECT command and, 449
 stored procedures as, 380, 380
 views as, 361
 triggers and performance loss, 418
 Performance Monitor, 112
 permissions. *See* object rights
 physical design. *See* database physical modeling
 physical model defined, 28

placement
 data files
 default placement, 67
 performance and, 127–131
 log file placement, 130–131
 tables, defining placement, 169–170, 170
 positional updates defined, 556
 primary data files (extension .MDF), 67
 primary filegroups defined, 120
 primary keys
 clustered indexes and, 321, 719
 defined, 16, 18, 19, 240
 defining, 241–246
 at table creation, 241–245, 245
 at table modification, 245–246
 relational model and, 16–19
 procedural integrity, 224
 procedure caches defined, 386–387
 Profiler, 731–744
 basics, 731–732, 732
 creating traces using, 738–739, 739
 functions and use, 733–738, 735, 736, 737
 Index Tuning Wizard, 741–744
 lock detection, 740–741
 replaying saved trace files, 739–740
 properties
 ACID, 396
 extended
 metadata and, 195
 tables and, 190–195, 191, 194

Q

queries
 analyzing, 703–716
 execution plan analysis, 709–714
 exercise in, 714–716
 SET FORCEPLAN, 704
 SET NOEXEC, 704–705
 SET SHOWPLAN_ALL, 705–707
 SET statements basics, 703–704
 SET STATISTICS IO, 708
 SET STATISTICS PROFILE, 709
 SET STATISTICS TIME, 709
 SETSHOWPLAN_TEXT, 707–708

- distributed queries, 489–500
 - about, 489–490
 - distributed partitioned views, 496–500, 497
 - linked servers, 490–494
 - OPENROWSET function, 494–496
 - using to modify data, 568–572, 571, 572
- exercise using, 453–455
- importance of understanding processing, 71
- inserting multiple rows with, 542–543
- optimizing, 717–732
 - changing an index, 723–726, 724, 726
 - indexes, 717–721
 - limiting the result set, 721–723, 722, 723
 - query hints, 726–731
- Query Analyzer procedures, *See also* Transact-SQL procedures
- aggregates, using, 468–469
- application roles
 - creating and assigning rights, 665–666
 - using, 665–666
- columns
 - creating computed columns, 189–190
 - granting rights to columns in tables, 674–675
 - GUID columns, creating and managing, 162–163
 - inserting selected, 539
 - updating multiple, 549–550
- constraints
 - cascading foreign key constraints, creating and using, 260–262
 - check constraints, creating and using, 238–239
 - foreign key constraints, creating and using, 257–258
 - primary key constraints, creating and using, 244–245
 - unique constraints, creating and using, 250–251
- data
 - deleting with OPENXML, 566–567
 - inserting with OPENXML, 564–565
 - modification with cursors, 557–558
 - storage and, 100–102
- deadlock conditions, creating, 636–637
- default values, defining and testing, 233–234
- indexes, creating, 322, 322

- rows
 - deleting sets of, 553–554
 - inserting single, 537–538
 - removing from tables, 552
 - updating, 545–548
 - updating with OPENXML, 565–566
- scripts and queries, opening, 79
- single table queries, creating, 453–455
- stored procedures, granting rights to, 678–680
- tables
 - granting rights to, 672–673
 - granting rights to columns in, 674–675
 - joining, 463–464
- UDFs, granting rights to, 682–684
- views
 - creating distributed partitioned views, 497–499
 - granting rights to, 676–677
- query hints, 726–731
- query optimizer, 699–703
 - basics, 699–700
 - batch optimization, 702–703
 - single statement optimization, 700–702
- query plans, 386–387
- query processing hints, 727–728
- query trees, 386–387

R

- RAID
 - data placement and, 129–130
 - log file placement and, 131
- real world scenarios
 - delete dilemma, 259
 - new database analysis, 4–5
 - operation order issue, 26–27
 - using Index Tuning Wizard, 744
 - using stored procedures, 391–392
 - using triggers to capture an audit trail, 681
 - working with mainframes, 601–602
 - write ahead paradigm, 112
- records
 - inserting, and GUIDS, 162
 - storage, 196–198, 196
- recovery model, and database options, 116
- recursive relationships, 16, 16

redundant columns, adding, 43–45, 44
 referential integrity, 24–28, 223
 relational models and normalization, 28–40
 advanced normalization, 39–40
 Boyce/Codd Normal form (BCNF), 39
 Fourth Normal Form (4NF), 39
 Fifth Normal Form (5NF), 40
 basics, 28–29, 41
 normal forms, 31–39
 First Normal Form, 32–35, 32, 33, 34
 Second Normal Form, 35–36, 36
 Third Normal Form, 36–39, 38
 normalization defined, 30
 relational tables, 29–31
 relational tables, 29–31
 relationships
 characteristics of, 10–13, 11, 12
 defined, 6
 identifying, 20–21
 many-to-many, 14–15, 15
 one-to-many, 13, 14, 14
 one-to-one, 13–14, 13
 recursive, 16, 16
 result sets, limiting, 721–723, 722, 723
 right and left outer joins, 459–460
 rights, viewing, 670–671, 671
 roles
 application, 664–666
 user-defined, 662–664
 roots of indexes defined, 291
ROWGUIDCOL
 dropping from columns, 174
 GUIDs and, 161–162
 rows, *See also* unique constraints; unique identifiers
 deleting, 551–554
 inserting, 535–544
 multiple, 542–544
 single rows, 535–538
 single rows with identity fields, 540–541
 single rows with selected columns, 538–539
OPENXML, 564–567
 deleting with, 566–567
 inserting with, 564–565
 updating with, 565–566
 relational tables and, 30

 updating, 545–548
 all, 545–546
 sets of rows, 547–548
 rowset views, 561–564
 rules
 data integrity and, 239–240
 referential integrity and, 24–26

S

scalar subqueries, 472
 scalars
 scalar UDFs, 399, 400–401
 SELECT statement and, 471–472
 schema locks, 632–633
 SCHEMABINDING option, 364–365
 SCOPE_IDENTITY, 158–160
 scopes defined, 158
 scripts
 Create Database statement, 68
 Generate SQL Scripts utility, 69
 Query Analyzer and opening, 79
 SCROLL cursors, 486
 SCROLL_LOCKS cursors, 487
 Second Normal Form (2NF), 35–36, 36
 secondary data files (.NDF), 67
 security, 654–695
 assigning object rights, 667–685
 object rights, 667–671, 671
 ownership chains, 684–685
 stored procedures, 677–680
 tables, 671–675
 triggers, 680–681
 user-defined functions (UDFs), 681–684
 views, 675–677
 exam essentials, 686–687
 key terms, 686
 linked servers and, 493
 overview, 654–666
 adding new database users, 660–662, 660, 661, 662
 adding roles and users, 663–664
 application roles, 664–666
 basics on database access, 659–660

784 SELECT list, subqueries in – statistics

- server access, 655–659, 657, 658, 659
- user defined roles, 662–663
- review questions and answers, 687–695
- stored procedures and, 379
- summary, 685
- views and, 361
- SELECT list, subqueries in, 472–473
- SELECT statement, 443–475
 - aggregate operators, 464–468
 - aliases, specifying, 450–451
 - basics, 443–444
 - cross joins and, 462
 - DISTINCT command, 470–471
 - formal syntax for, 474–475
 - in views, 362–363, 367–368
 - inner joins and, 456–458
 - joining tables exercise, 463–464
 - ORDER BY clause, 446–447
 - outer joins and, 458–461
 - returned results, limiting, 451–453
 - scalars and, 471–472
 - single table SELECT, 444–445, 445
 - subqueries and, 472–473
 - UNION command, 470
 - using aggregates exercise, 468–469
 - using queries exercise, 453–455
 - using with XML, 502–506
 - WHERE clause, 447–449
- servers, *See also* SQL Server 2000
 - accessing, and security, 655–659, 657, 658, 659
 - extracting template files from, 508–511
 - linked servers
 - accessing data and, 490–494
 - modifying data with, 569
 - lock hints and, 643–644
- SET FORCEPLAN, 704
- SET NOEXEC, 704–705
- SET ROWCOUNT, 451–453
- SET SHOWPLAN_ALL, 705–707
- SET statements, 703–709
- SET STATISTICS IO, 708
- SET STATISTICS TIME, 709
- SETSHOWPLAN_TEXT, 707–708
- shared locks, 631
- shrinking databases, 80–90
 - automatic shrinking, 80–81
 - manual shrinking, 82–85, 82
 - shrinking in Enterprise Manager, 85–87, 85
- shrinking files, 87–90, 89, *See also* shrinking databases
- single statement optimization, 700–702
- single table SELECT, 444–445, 445
- size
 - of databases
 - Create Database statement options and, 74–79, 77, 78
 - increasing, 86–87
 - managing size and growth options, 74–75
 - of log files, and space management, 109–110
 - space management, 94–112
 - data files, 94–107, 95, 96, 104
 - allocation tracking, 103–107, 104
 - data allocation, 99–103
 - extents, 95–98, 95, 96
 - pages, 98–99, 99
 - log files, 107–113, 109, 110, 111
 - log records, 110–111
 - size of, 109–110
- sp_createstats system stored procedure, 317
- special attributes datatypes, 23
- SP_LOCK result set, 641–642
- SP_XML_PREPAREDDOCUMENT, 559–561
- SQL Enterprise Manager. *See* Enterprise Manager
- SQL Profiler. *See* Profiler
- SQL Query Analyzer. *See* Query Analyzer procedures
- SQL Server 2000
 - accessing, 655–659
 - authentication, 655
 - login, 655–659, 657, 658, 659
 - accessing data and, 442–443
 - importance of powerful hardware, 128
 - integrating with XML, 507–514
 - uses of, 514
 - using template files, 508–511
 - XPATH (XML Path Language), 514
 - XSL, 511–513, 513
 - OSQL, 79
- STATIC cursors, 486
- statistics
 - creating on multiple columns, 331–332
 - index choice and, 306–320
 - distribution statistics, 307–311

- index choice, 312–314, 312, 315
 - statistics maintenance, 315–320
- options defined, 318–319
- storage
 - computed columns storage, 327
 - data storage, 100–102
 - structured storage files, 613
 - table storage, 195–203
 - BLOBs and, 182
 - record storage, 196–198, 196
 - text columns, 198–203, 202
 - text, ntext and image storage, 198–203, 202
- stored procedures
 - advantages, 379–380, 380, 699
 - creating and altering, 380–384
 - defined, 379, 380, 380, 678
 - error handling, 392–396
 - executing, 386–392
 - compiling, 387–389
 - query tree and query plans, 386–387
 - running procedures, 389–392
 - object rights and, 677–680
 - optimizing, 744–746
 - stored procedure debugger, 746–755
 - exercise in debugging, 749–755, 752, 753, 754
 - using, 746–749, 748
 - using parameters, 384–385
 - using to insert multiple rows, 543–544
- striped disks, and data placement, 129–130
- structured storage files, 613
- subqueries, 472–475
 - correlated subqueries, 473
 - SELECT statement and, 472–475
- summaries
 - data
 - data access, analyzing and optimizing, 755–756
 - accessing, 515–516
 - importing and exporting, 617–618
 - data integrity, 267
 - modifying, 573
 - databases
 - database logical modeling, 48
 - database objects, 420
 - database physical modeling, 132
 - indexes, 342
 - locking, 645–646

- security, 685
- tables, creating and maintaining, 204
- synonyms, system datatypes, 183–184
- system datatypes in columns, 176–184
 - date and time, 181
 - listed, 177–180
 - numeric, 180–181
 - special, 182–183
 - strings, 181–182
 - synonyms, 183–184

T

- T-SQL. *See* Transact-SQL procedures
- table-valued UDFs defined, 400
- tables, *See also* columns; data integrity; rows
 - defined, 151
 - edge tables, 563
 - exercise in joining, 463–464
 - granting rights to, 672–673
 - granting rights to columns in, 674–675
 - isolating, 26
 - object rights and, 671–675
 - partitioning
 - horizontal, 47
 - vertical, 46–47, 46
 - physically separating from indexes, 126
 - reindexing, 341
 - relational, 29–31
- tables, creating and maintaining, 150–219
 - altering, 170–176
 - adding or dropping columns, 174–176
 - altering columns, 171–174
 - computed columns, 188–189
 - constraints, defining
 - check constraints, 235–239
 - default constraints, 228–231, 230
 - creating, 150–170
 - collation and, 163–165, 165
 - filegroups and, 165–170, 166, 167, 168, 170
 - GUID columns and, 162–163
 - identity, 156–158
 - identity comparisons, 158–161
 - simple, 151–155, 151, 152
 - simple with Enterprise Manager, 155–156
 - UniqueIdentifier, 161–162

- exam essentials, 205
- extended properties, 190–195, 191, 194
- foreign key constraints, defining at creation, 253–257
- key terms, 204
- placement, defining with Enterprise Manager, 169–170, 170
- primary keys, 241–246
 - defining at creation, 241–245, 244
 - defining at table modification, 245–246
- review questions and answers, 205–219, 207, 211, 214
- storage, 195–203
 - record storage, 196–198, 196
 - text, ntext and image storage, 198–203, 202
- summary, 204
- system datatypes, 176–184
 - date and time, 181
 - listed, 177–180
 - numeric, 180–181
 - special, 182–183
 - strings, 181–182
 - synonyms, 183–184
- unique constraints, 247–252
 - defining at table creation, 247–251, 249
 - defining at table modification, 251–252
- user-defined datatypes and, 184–187
- tasks, and DTS, 610–612
- template files, 508–511
- templates defined, 501
- text columns, and storage, 198–203, 202
- Third Normal Form (3NF), 36–39, 38
- time and date values, 181
- TOP, 451–452
- traces, 734–741
 - basics, 734–738, 735, 736, 737
 - creating using Profiler, 738–739, 739
 - replaying with Profiler, 739–740
 - trace flags with lock detection, 741
- Transact-SQL procedures, *See also* Query Analyzer procedures; various statements
 - data manipulation, 534
 - DELETE statement, 551–554
 - INSERT statement, 535–544
 - UPDATE statement, 544–551

- databases
 - adding new users, 660–662, 660, 661, 662
 - adding roles and users, 663–664
 - altering, 91–93
 - creating, 68
 - options, 113–114
 - files, defining data and log files, 68
 - inexperienced programmers and, 449
 - tables, creating simple, 151–152
- transaction isolation levels, 627–628
- transaction logs
 - importance of, 71
 - space management and, 107–108
 - storing, 131
- transactions
 - defining, 69
 - locking and, 634–638
- triggers, 406–420
 - advantages of, 406–407
 - AFTER, 407–408, 419
 - AFTER INSERT, 419
 - cascading updates and deletes and, 260
 - defined, 406–407, 680
 - DELETE, 410–412, 411
 - disabling, 419–420
 - firing order, changing, 419
 - INSERT, 409–410, 409
 - INSTEAD of, 413–418
 - INSTEAD OF DELETE, 415
 - INSTEAD OF INSERT, 414–415
 - INSTEAD OF UPDATE, 416
 - multiple, 419
 - object rights and, 680–681
 - optimizing, 744–746
 - performance considerations, 418
 - UPDATE, 412–413, 412
 - vs. constraints, 225
 - Web, 419

U

- UDFs. *See* user-defined functions (UDFs)
- Unicode vs. collation, 164
- uniform extents, 95, 96
- uniform resource locators (URLs)

- accessing SQL Server through, 507–514
- defined, 501
- UNION command, 470
- unique constraints, 246–252
 - defining at table creation, 247–251, 249
 - defining at table modification, 251–252
- unique identifiers, 156–163, 337
 - basics, 156–158
 - collation, 163–165, 165
 - comparisons of identity, 158–161
 - globally unique identifiers (GUIDs)
 - creating tables with GUID columns, using Query Analyzer, 162–163
 - statement creating tables with GUID columns, 161–162
 - identity, defined, 156
 - UniqueIdentifier, 161–162
- unique indexes
 - creating, 324–325
 - described, 296
- update locks, 632
- update rules, 25–26
- UPDATE statement, 544–551
 - all rows, 545–546
 - multiple columns, 548–550
 - sets of rows, 547–548
 - views, 550–551
- UPDATE triggers, 412–413, 412
- updateable cursors, 555–556
- updates
 - fragmentation and, 337–338, 338
 - statistics examples, 310–320
- URLs (uniform resource locators)
 - accessing SQL Server through, 507–514
 - defined, 501
- user access to databases, 663–664
- user-defined datatypes in columns, 184–187
- user-defined filegroups defined, 120
- user-defined functions (UDFs), 399–405
 - advantages of, 400
 - creating and altering UDFs, 400–404
 - defined, 399–400, 682
 - granting rights to, 682–684
 - object rights and, 681–684
 - types, 399–400
 - using UDFs, 404–405
- user-defined roles, 662–664

V

- values
 - database options values, 114–115
 - default values, 227–234
 - basics, 227–228
 - default constraints, 228–231, 230
 - default objects, 232–234
 - identity values, 158–160, 541
 - NULL values
 - changing in columns, 173
 - defining columns and, 154–155
 - time and date values, 181
 - vertical partitions defined, 445, 445
 - VIEW_METADATA option, 365
 - views, 359–378
 - creating and altering, 362–367
 - ENCRYPTION option, 364
 - METADATA option, 365
 - SCHEMABINDING option, 364–365
 - SELECT statement, 362–363
 - VIEW_METADATA option, 365
 - WITH CHECK OPTION option, 366–367
 - defined, 360, 360, 675
 - distributed partitioned views, 496–500, 497
 - granting rights to, 676–677
 - indexed views, 369–372
 - modifying new rowset views with OPENXML, 561–564
 - object rights and, 675–677
 - partitioned views, 372–378, 372
 - SELECT statement and, 362–363, 367–368
 - triggers and, 416–418
 - updating, 550–551
 - uses, 361
 - using, 367–369
 - retrieving data, 367–368
 - updating data, 368–369
- virtual log files, 109, 109

W

- Web triggers, 419
- WHERE clause
 - basics, 447–449
 - subqueries in, 473–474

WITH CHECK OPTION option, 366–367
workflows, in DTS, 611–612

X

XML (Extensible Markup Language)
 modifying data using OPENXML, 559–568
 deleting rows, 566–567
 inserting rows, 564–565
 new rowset views, 561–564
 SP_XML_PREPAREDDOCUMENT, 559–561
 updating rows, 565–566
 options, 503
 SQL Server 2000 and, 500–514
 description, 500–501
 extracting data in XML format, 507
 integrating using XPATH, 514

 integrating using XSL, 511–513, 513
 using SELECT with, 502–506
 using template files to access SQL Server
 through a URL, 508–511
XML Path Language (XPath)
 defined, 501
 navigating XML documents and, 514
XMLDATA option, 503
XMLTemplate1.xml, 509, 511
XMLTemplate2.xml, 510
XPath (XML Path Language)
 defined, 501
 navigating XML documents and, 514
XSL (Extensible Stylesheet Language)
 defined, 501
 extracting data in XML format and, 511–513,
 513
XSLStyle.xml.512