



International
Edition

Database Processing

FUNDAMENTALS, DESIGN, AND IMPLEMENTATION

Eleventh Edition

David M. Kroenke • David J. Auer

PEARSON

Preface 17

PART 1 ● GETTING STARTED**31****Chapter 1: Introduction 32**

Chapter Objectives	32
The Characteristics of Databases	33
A Note on Naming Conventions	33
A Database Has Data and Relationships	34
Databases Create Information	35
Database Examples	36
Single-User Database Applications	36
Multiuser Database Applications	36
E-Commerce Database Applications	37
Reporting and Data Mining Database Applications	37
The Components of a Database System	38
Database Applications and SQL	39
The DBMS	41
The Database	42
Personal Versus Enterprise-Class Database Systems	43
What Is Microsoft Access?	43
What Is an Enterprise-Class Database System?	45
Database Design	46
Database Design from Existing Data	46
Database Design for New Systems Development	47
Database Redesign	48
What You Need to Learn	49
A Brief History of Database Processing	50
The Early Years	50
The Emergence and Dominance of the Relational Model	52
Post-Relational Developments	53
<i>Summary</i>	54
<i>Key Terms</i>	55
<i>Review Questions</i>	56
<i>Project Questions</i>	58

Chapter 2: Introduction to Structured Query Language 61

Chapter Objectives	61
Cape Codd Outdoor Sports	62
The Retail Sales Data Extraction	63
RETAIL_ORDER Data	63
ORDER_ITEM Data	64
SKU_DATA Table	65
Data Extracts Are Common	65

SQL Background	65
The SQL SELECT/FROM/WHERE Framework	66
Reading Specified Columns from a Single Table	66
Reading Specified Rows from a Single Table	68
Reading Specified Columns and Rows from a Single Table	69
Submitting SQL Statements to the DBMS	70
Using SQL in Microsoft Access 2007	70
Using SQL in Microsoft SQL Server 2008	76
Using SQL in Oracle Database 11g	78
Using SQL in Sun Microsystems MySQL 5.1	80
SQL Enhancements for Querying a Single Table	82
Sorting the Results	83
SQL WHERE Clause Options	85
Combing the SQL WHERE Clause and the SQL ORDER BY Clause	89
Performing Calculations in SQL Queries	89
Using SQL Built-in Functions	89
SQL Expressions in SQL SELECT Statements	92
Grouping in SQL SELECT Statements	94
Looking for Patterns in NASDAQ Trading	97
Investigating the Characteristics of the Data	98
Searching for Patterns in Trading by Day of Week	99
Querying Two or More Tables with SQL	100
Querying Multiple Tables with Subqueries	100
Querying Multiple Tables with Joins	103
Comparing Subqueries and Joins	106
<i>Summary</i>	107
<i>Key Terms</i>	107
<i>Review Questions</i>	108
<i>Project Questions</i>	112
<i>Marcia's Dry Cleaning</i>	116
<i>Morgan Importing</i>	119

PART 2 ● DATABASE DESIGN

123

Chapter 3: The Relational Model and Normalization 124

Chapter Objectives	124
Relational Model Terminology	126
Relations	126
Characteristics of Relations	127
Alternative Terminology	128
Functional Dependencies	130
Finding Functional Dependencies	131
Keys	134
Normal Forms	136
<i>Modification Anomalies</i>	136
A Short History of Normal Forms	137
Normalization Categories	138
From First Normal Form to Boyce-Codd Normal Form	138
Eliminating Anomalies from Functional Dependencies	140
Eliminating Anomalies from Multivalued Dependencies	147
Fifth Normal Form	150
Domain/Key Normal Form	150
<i>Summary</i>	151
<i>Key Terms</i>	151
<i>Review Questions</i>	152

<i>Project Questions</i>	154
<i>Marcia's Dry Cleaning</i>	155
<i>Morgan Importing</i>	156

Chapter 4: Database Design Using Normalization 157

Chapter Objectives	157
Assess Table Structure	158
Designing Updatable Databases	159
Advantages and Disadvantages of Normalization	159
Functional Dependencies	159
Normalizing with SQL	160
Choosing Not to Use BCNF	161
Multivalued Dependencies	162
Designing Read-Only Databases	162
Denormalization	162
Customized Duplicated Tables	164
Common Design Problems	165
The Multivalue, Multicolumn Problem	165
Inconsistent Values	166
Missing Values	167
The General-Purpose Remarks Column	168
<i>Summary</i>	169
<i>Key Terms</i>	169
<i>Review Questions</i>	169
<i>Project Questions</i>	171
<i>Marcia's Dry Cleaning</i>	172
<i>Morgan Importing</i>	173

Chapter 5: Data Modeling with the Entity-Relationship Model 174

Chapter Objectives	174
The Purpose of a Data Model	175
The Entity-Relationship Model	175
Entities	175
Attributes	176
Identifiers	176
Relationships	177
Maximum Cardinality	179
Minimum Cardinality	180
Entity-Relationship Diagrams and Their Versions	181
Variations of the E-R Model	181
E-R Diagrams Using the IE Crow's Foot Model	182
Strong Entities and Weak Entities	183
ID-Dependent Entities	183
Non-ID-Dependent Weak Entities	184
Subtype Entities	186
Patterns in Forms, Reports, and E-R Models	187
Strong Entity Patterns	188
ID-Dependent Relationships	192
Mixed Identifying and Nonidentifying Patterns	198
The For-Use-By Pattern	201
Recursive Patterns	202
The Data Modeling Process	204
The College Report	205
The Department Report	205
The Department/Major Report	206
The Student Acceptance Letter	208

<i>Summary</i>	210
<i>Key Terms</i>	211
<i>Review Questions</i>	212
<i>Project Questions</i>	214
<i>Marcia's Dry Cleaning</i>	221
<i>Morgan Importing</i>	221

Chapter 6: Transforming Data Models into Database Designs 222

Chapter Objectives	222
Create a Table for Each Entity	223
Selecting the Primary Key	224
Specifying Candidate (Alternate) Keys	225
Specify Column Properties	225
Verify Normalization	227
Create Relationships	228
Relationships Between Strong Entities	228
Relationships Using ID-Dependent Entities	231
Relationships with a Weak Non-ID-Dependent Entity	236
Relationships in Mixed Entity Designs	236
Relationships Between Supertype and Subtype Entities	238
Recursive Relationships	238
Representing Ternary and Higher-Order Relationships	240
Relational Representation of the Highline University Data Model	243
Design for Minimum Cardinality	244
Actions When the Parent Is Required	246
Actions When the Child Is Required	247
Implementing Actions for M-O Relationships	248
Implementing Actions for O-M Relationships	249
Implementing Actions for M-M Relationships	249
Designing Special Case M-M Relationships	249
Documenting the Minimum Cardinality Design	250
An Additional Complication	252
Summary of Minimum Cardinality Design	252
The View Ridge Gallery Database	252
Summary of Requirements	252
The View Ridge Data Model	253
Database Design with Data Keys	254
Minimum Cardinality Enforcement for Required Parents	255
Minimum Cardinality Enforcement for the Required Child	257
Column Properties for the View Ridge Database Design Tables	258
<i>Summary</i>	260
<i>Key Terms</i>	261
<i>Review Questions</i>	261
<i>Project Questions</i>	263
<i>Marcia's Dry Cleaning</i>	264
<i>Morgan Importing</i>	264

Chapter 7: SQL for Database Construction and Application Processing 266

Chapter Objectives	266
The View Ridge Gallery Database	267
SQL, DDL, DML, and a New Type of Join	267

Managing Table Structure with SQL DDL	268
Creating the View Ridge Database	268
Using the SQL CREATE TABLE Statement	269
Variations in SQL Data Types	269
Creating the ARTIST Table	270
Creating the WORK Table and the 1:N ARTIST-to-WORK Relationship	274
Implementing Required Parent Rows	275
Implementing 1:1 Relationships	275
Casual Relationships	275
Creating Default Values and Data Constraints with SQL	276
Creating the View Ridge Database Tables	278
The SQL ALTER Statement	281
The SQL DROP TABLE Statement	282
SQL DML Statements	282
The SQL INSERT Statement	282
Populating the View Ridge Database Tables	283
The SQL UPDATE Statement	284
The SQL DELETE Statement	290
New Forms of Join	290
The SQL JOIN ON Syntax	290
Outer Joins	292
Using SQL Views	296
Using SQL Views to Hide Columns and Rows	298
Using SQL Views to Display Results of Computed Columns	299
Using SQL Views to Hide Complicated SQL Syntax	300
Layering Built-in Functions	300
Using SQL Views for Isolation, Multiple Permissions, and Multiple Triggers	303
Updating SQL Views	303
Embedding SQL in Program Code	304
Using SQL Triggers	305
Using Triggers to Provide Default Values	306
Using Triggers to Enforce Data Constraints	308
Using Triggers to Update Views	309
Using Triggers to Implement Referential Integrity Actions	309
Using Stored Procedures	311
Advantages of Stored Procedures	311
The WORK_AddWorkTransaction Stored Procedure	313
<i>Summary</i>	313
<i>Key Terms</i>	315
<i>Review Questions</i>	316
<i>Project Questions</i>	320
<i>Marcia's Dry Cleaning</i>	323
<i>Morgan Importing</i>	324

Chapter 8: Database Redesign 325

Chapter Objectives	325
The Need for Database Redesign	326
SQL Statements for Checking Functional Dependencies	326
Correlated Subqueries	327
Analyzing the Existing Database	332
Reverse Engineering	332
Dependency Graphs	333
Database Backup and Test Databases	334
Changing Table Names and Table Columns	335
Changing Table Names	335
Adding and Dropping Columns	336

Changing a Column Data Type or Column Constraints	337
Adding and Dropping Constraints	338
Changing Relationship Cardinalities and Properties	338
Changing Minimum Cardinalities	338
Changing Maximum Cardinalities	339
Adding and Deleting Tables and Relationships	342
Forward Engineering(?)	342
<i>Summary</i>	343
<i>Key Terms</i>	344
<i>Review Questions</i>	344
<i>Project Questions</i>	346
<i>Marcia's Dry Cleaning</i>	347
<i>Morgan Importing</i>	348

PART 4 ● MULTIUSER DATABASE PROCESSING

349

Chapter 9: Managing Multiuser Databases 350

Chapter Objectives	350
Database Administration	351
Managing the Database Structure	352
Concurrency Control	353
The Need for Atomic Transactions	354
Resource Locking	358
Optimistic Versus Pessimistic Locking	360
Declaring Lock Characteristics	361
Consistent Transactions	362
Transaction Isolation Level	363
Cursor Type	364
Database Security	365
Processing Rights and Responsibilities	365
DBMS Security	367
DBMS Security Guidelines	367
Application Security	369
The SQL Injection Attack	370
Database Backup and Recovery	370
Recovery via Reprocessing	371
Recovery via Rollback/Rollforward	371
Managing the DBMS	373
Maintaining the Data Repository	374
<i>Summary</i>	375
<i>Key Terms</i>	376
<i>Review Questions</i>	377
<i>Project Questions</i>	378
<i>Marcia's Dry Cleaning</i>	379
<i>Morgan Importing</i>	380

Chapter 10: Managing Databases with SQL Server 2008 381

Chapter Objectives	381
Installing SQL Server 2008	382
The Microsoft SQL Server 2008 Management Studio	383
Creating an SQL Server 2008 Database	384
SQL Server 2008 Utilities	386
SQL CMD and Microsoft PowerShell	386
Microsoft SQL CLR	386

SQL Server 2008 GUI Displays	387
SQL Server 2008 SQL Statements and SQL Scripts	387
Creating and Populating the View Ridge Database Tables	390
Creating the View Ridge Database Table Structure	390
Reviewing Database Structures in the SQL Server GUI Display	394
Indexes	398
Populating the VRG Tables with Data	399
Creating Views	409
SQL Server 2008 Application Logic	411
Transact-SQL	412
Stored Procedures	415
Triggers	423
Concurrency Control	437
Transaction Isolation Level	438
Cursor Concurrency	438
Locking Hints	439
SQL Server 2008 Security	440
SQL Server Database Security Settings	442
SQL Server 2008 Backup and Recovery	444
Backing Up a Database	444
SQL Server Recovery Models	445
Restoring a Database	446
Database Maintenance Plans	446
Topics Not Discussed in This Chapter	447
<i>Summary</i>	447
<i>Key Terms</i>	447
<i>Review Questions</i>	448
<i>Project Questions</i>	450
<i>Marcia's Dry Cleaning</i>	451
<i>Morgan Importing</i>	451

ONLINE CHAPTER: SEE PAGE 453 FOR INSTRUCTIONS

Chapter 10A: Managing Databases with Oracle Database 11g 10A-1

Chapter Objectives	10A-1
Installing Oracle Database 11g	10A-2
Oracle Database 11g Documentation	10A-4
Oracle and Java	10A-5
Oracle Database 11g Administration and Development Tools	10A-5
The Oracle Database Configuration Assistant	10A-6
The Oracle Enterprise Manager 11g Database Control	10A-6
Oracle Tablespace	10A-8
Oracle Security	10A-11
User Privileges	10A-11
Creating a User Account	10A-11
Creating a Role	10A-13
Oracle Application Development Tools	10A-15
Oracle SQL*Plus	10A-15
Oracle SQL Developer	10A-16
Oracle Schemas	10A-17
Oracle Database 11g 2008 SQL Statements and SQL Scripts	10A-18
Creating and Populating the View Ridge Database Tables	10A-20
Creating the View Ridge Database Table Structure	10A-20
Reviewing Database Structures in the SQL Developer GUI Display	10A-23

Indexes	10A-26
Populating the VRG Tables	10A-28
Creating Views	10A-36
Application Logic	10A-39
Oracle PL/SQL	10A-40
Stored Procedures	10A-42
Triggers	10A-50
Concurrency Control	10A-61
Read Committed Transaction Isolation Level	10A-62
Serializable Transaction Isolation Level	10A-63
Read-Only Transaction Isolation	10A-63
Additional Locking Comments	10A-63
Oracle Backup and Recovery	10A-64
Oracle Recovery Facilities	10A-64
Types of Failure	10A-64
Topics Not Discussed in This Chapter	10A-65
<i>Summary</i>	10A-65
<i>Key Terms</i>	10A-66
<i>Review Questions</i>	10A-67
<i>Project Questions</i>	10A-69
<i>Marcia's Dry Cleaning</i>	10A-70
<i>Morgan Importing</i>	10A-70

ONLINE CHAPTER: SEE PAGE 455 FOR INSTRUCTIONS

Chapter 10B: Managing Databases with MySQL 5.1 10B-1

Chapter Objectives	10B-1
The MySQL 5.1 DBMS	10B-2
Installing and Updating MySQL	10B-2
Configuring MySQL	10B-2
MySQL Storage Engines	10B-3
The MySQL GUI Utilities	10B-4
Creating and Using a MySQL Database	10B-6
MySQL Utilities	10B-7
MySQL Command Line Client	10B-8
MySQL GUI Displays	10B-8
MySQL SQL Statements and SQL Scripts	10B-10
Creating and Populating the View Ridge Database Tables	10B-13
Creating the View Ridge Database Table Structure	10B-13
Reviewing Database Structures in the MySQL GUI Display	10B-16
Indexes	10B-17
Populating the VRG Tables with Data	10B-22
Creating Views	10B-30
MySQL Application Logic	10B-33
MySQL Procedural Statements	10B-34
Stored Procedures	10B-36
Triggers	10B-43
A Last Word on MySQL Stored Procedures and Triggers	10B-53
Concurrency Control	10B-54
MySQL 5.1 Security	10B-55
MySQL Database Security Settings	10B-56
MySQL 5.1 Server Backup and Recovery	10B-58
Backing Up a MySQL Database	10B-58
Restoring a MySQL Database	10B-59
Topics Not Discussed in This Chapter	10B-61

Summary 10B-61
Key Terms 10B-62
Review Questions 10B-62
Project Questions 10B-63
Marcia's Dry Cleaning 10B-65
Morgan Importing 10B-65

PART 5 • DATABASE ACCESS STANDARDS

457

Chapter 11: The Web Server Environment 458

Chapter Objectives 458
The Web Database Processing Environment 459
The ODBC Standard 461
 ODBC Architecture 461
 Conformance Levels 462
 Creating an ODBC Data Source Name 464
The Microsoft .NET Framework and ADO.NET 467
 OLE DB 468
 ADO and ADO.NET 472
 The ADO.NET Object Model 472
The JAVA Platform 476
 JDBC 476
 Java Server Pages (JSP) and Servlets 478
 Apache Tomcat 479
Web Database Processing with PHP 479
 Web Database Processing with PHP and Eclipse 481
 Getting Started with HTML Web Pages 483
 The index.html Web Page 483
 Creating the index.html Web Page 483
 Using PHP 486
 Challenges for Web Database Processing 493
Web Page Examples with PHP 493
 Example 1: Updating a Table 494
 Example 2: Using PHP Data Objects (PDO) 500
 Example 3: Invoking a Stored Procedure 500
Summary 508
Key Terms 510
Review Questions 511
Project Questions 513
Marcia's Dry Cleaning 515
Morgan Importing 515

Chapter 12: Database Processing with XML 516

Chapter Objectives 516
The Importance of XML 517
XML as a Markup Language 518
 XML Document Type Declarations 519
 Materializing XML Documents with XSLT 519
XML Schema 524
 XML Schema Validation 524
 Elements and Attributes 525
 Flat Versus Structured Schemas 526
 Global Elements 530

Creating XML Documents from Database Data	532
Using the SQL SELECT . . . FOR XML Statement	532
Multitable SELECT with FOR XML	536
An XML Schema for All CUSTOMER Purchases	540
A Schema with Two Multivalue Paths	544
Why Is XML Important?	544
Additional XML Standards	550
<i>Summary</i>	552
<i>Key Terms</i>	553
<i>Review Questions</i>	553
<i>Project Questions</i>	554
<i>Marcia's Dry Cleaning</i>	555
<i>Morgan Importing</i>	555

Chapter 13: Database Processing for Business Intelligence Systems 556

Chapter Objectives	556
Business Intelligence Systems	556
The Relationship Between Operational and BI Systems	556
Reporting Systems and Data Mining Applications	557
Reporting Systems	558
Data Mining Applications	558
Data Warehouses and Data Marts	558
Components of a Data Warehouse	558
Data Warehouses Versus Data Marts	561
Dimensional Databases	562
Reporting Systems	569
RFM Analysis	570
Producing the RFM Report	571
Reporting System Components	574
Report Types	575
Report Media	575
Report Modes	576
Report System Functions	576
OLAP	579
Data Mining	582
Unsupervised Data Mining	584
Supervised Data Mining	585
Three Popular Data Mining Techniques	585
Market Basket Analysis	585
Using SQL for Market Basket Analysis	586
<i>Summary</i>	587
<i>Key Terms</i>	588
<i>Review Questions</i>	589
<i>Project Questions</i>	591
<i>Marcia's Dry Cleaning</i>	593
<i>Morgan Importing</i>	594

APPENDICES

ONLINE APPENDICES: SEE PAGE 595 FOR INSTRUCTIONS

Appendix A: Getting Started with Microsoft Access 2007 A-1

Chapter Objectives	A-1
Table Keys	A-2
Relationships Among Tables	A-2

Creating an Access Database	A-3
The Microsoft Office Fluent User Interface	A-5
The Ribbon and Command Tabs	A-5
Contextual Command Tabs	A-5
Database Objects and the Navigation Pane	A-6
Closing a Database and Exiting Microsoft Access	A-7
Opening an Existing Access Database	A-8
Creating Access Database Tables	A-9
Inserting Data into Tables—The Datasheet View	A-17
Modifying and Deleting Data in Tables in the Datasheet View	A-19
Creating Relationships Between Tables	A-20
Working with Microsoft Access Queries	A-23
Microsoft Access Forms and Reports	A-26
Creating Access 2007 Forms	A-26
Creating Access 2007 Reports	A-27
Closing a Database and Exiting Access 2007	A-27
<i>Key Terms</i>	A-28
<i>Review Questions</i>	A-28

Appendix B: The IDEF1X Standard B-1

Chapter Objectives	B-1
IDEF1X Entities	B-1
IDEF1X Relationships	B-2
Nonidentifying Connection Relationships	B-2
Identifying Connection Relationships	B-3
Nonspecific Relationships	B-4
Categorization Relationships	B-4
Domains	B-6
Domains Reduce Ambiguity	B-6
Domains Are Useful	B-7
Base Domains and Typed Domains	B-7
<i>Key Terms</i>	B-8
<i>Review Questions</i>	B-8

Appendix C: UML-Style Entity-Relationship Diagrams C-1

Chapter Objectives	C-1
UML Entities and Relationships	C-1
Representation of Weak Entities	C-2
Representation of Subtypes	C-2
OOP Constructs Introduced by UML	C-3
The Role of UML in Database Processing Today	C-4
<i>Key Terms</i>	C-5
<i>Review Questions</i>	C-5

Appendix D: Data Structures for Database Processing D-1

Chapter Objectives	D-1
Flat Files	D-1
Processing Flat Files in Multiple Orders	D-2
A Note on Record Addressing	D-2
Maintaining Order with Linked Lists	D-3
Maintaining Order with Indexes	D-5
B-Trees	D-5
Summary of Data Structures	D-7
Representing Binary Relationships	D-8
Review of Record Relationships	D-8
Representing Trees	D-10

Representing Simple Networks	D-12
Representing Complex Networks	D-13
Summary of Relationship Representations	D-14
Secondary-Key Representations	D-16
Linked-List Representation of Secondary Keys	D-17
Index Representation of Secondary Keys	D-18

Key Terms D-18

Review Questions D-19

Appendix E: The Semantic Object Model E-1

Chapter Objectives E-1

Semantic Objects E-2

Defining Semantic Objects E-2

Attributes E-3

Object Identifiers E-5

Attribute Domains E-6

Semantic Object Views E-6

Types of Objects E-7

Simple Objects E-7

Composite Objects E-8

Compound Objects E-11

Representing One-to-One Compound Objects E-12

Representing One-to-Many and Many-to-One Relationships E-14

Representing Many-to-Many Relationships E-14

Hybrid Objects E-17

Representing Hybrid Relationships E-20

Association Objects E-22

Parent/Subtype Objects E-24

Archetype/Version Objects E-26

Comparing the Semantic Object and the E-R Models E-29

Key Terms E-31

Review Questions E-32

***Bibliography* 597**

***Glossary* 599**

***Index* 621**