

# Database Management Systems

Designing and Building  
Business Applications

**Version 6.0.0**

Gerald V. Post

*University of the Pacific*

Database Management Systems  
Designing and Building Business Applications

Copyright © 2014 by Gerald V. Post

All rights reserved. No part of this publication may be reproduced or distributed in any form or stored in any database or retrieval system without the prior written consent of Gerald V. Post.

Students:

Your honesty is critical to your reputation. No company wants to hire a thief—particularly for jobs as critical as application development and database administration. If someone is willing to steal something as inexpensive as an e-book, how can that person be trusted with billions of dollars in corporate accounts?

You are not allowed to “share” this book in any form with anyone else. You cannot give or sell any information from this publication in any form to anyone else.

To purchase this book or other books: <http://JerryPost.com/Books>



## Goals and Philosophy

---

Working with business information systems is exciting. You get to work with people to find ways to improve their jobs. You get to use your creativity to build new applications. You often work on teams to share ideas and techniques. In the end, you get to see an application that you built help companies and people. Database management systems are key components in almost every business application. They organize and store the data so it can be retrieved and analyzed.

The goal of this text is straightforward: At the end of the text students should be able to evaluate a business situation and build a database application.

Building an application requires four basic steps. The book is organized by these four steps: (1) Database Design, (2) SQL Queries, (3) Application forms and reports, and (4) Database administration.

The first two steps (design and SQL) are standard for almost every database management system. Normalization shows how to carefully design databases to gain the strengths of the database approach. SQL is a standard query language that is used for virtually every step of application development. These two topics must be covered carefully and thoroughly, particularly because they are both difficult topics for students.

The concepts of forms and reports are relatively standard, but every database system has different tools to build applications. Similarly, database administrators perform relatively standard tasks at every company, but the tools are different for every system. The main textbook explains the basic tasks in neutral terms that apply to every database management system.

## Workbooks

---

Students need to do more than learn the basic concepts and theory. They need to actually build an application. A database management system (DBMS) is a complex tool with thousands of options and individual quirks. The accompanying workbooks are designed to show students how to build an application using a particular DBMS. The workbooks follow the main text chapters but focus on creating an application for a specific company.

The first step in using the Workbooks is to choose a DBMS. Be sure all of the components of the DBMS are available and installed. In particular, Oracle requires several components and Microsoft SQL Server needs Visual Studio. A nice feature of the Workbooks is that they all use the same case. So students can learn one DBMS and then later use a different workbook to learn the details of a new DBMS.

## Learning Assessment

---

Learning assessment is important to students as well as faculty and employers. Students need to determine what aspects they are strong in and which ones need additional work. Students need to understand that if they can successfully learn this material, they will have acquired several skills that will get them jobs and help them contribute to businesses by being able to quickly build and maintain business applications.

Learning assessment in this book is straightforward: At the end of the book, students should be able to analyze a business situation and develop a database

application. The complexity of the application and tools used will depend on the specific class and the background of the students. Students should develop a term-project as part of a course. Several sample projects are included as a separate project book and several more are in the workbooks. The project provides an excellent opportunity to assess overall learning. The final project can be evaluated in terms of (1) correctly meeting the business needs, (2) an efficient database structure, and (3) usability.

It is also useful to assess individual skills independently—particularly if groups are used to create the final project. In this case, assessment consists of individual exams for (1) database design and normalization, (2) SQL and creating queries from business questions, and (3) selected topics including database programming, security, data mining, and distributed systems.

## Organization

The organization of the text follows the basic steps of application development: design, queries, applications, administration, and advanced topics. Some instructors might prefer to teach queries before database design, so the initial chapters are written with that flexibility.

The introduction explains the importance of databases and relates database applications to topics the students have likely seen in other classes.

The section on database design has two chapters: Chapter 2 on general design techniques (systems techniques, diagramming, and control) and Chapter 3, which details data normalization. Chapter 2 leans towards an object and graphical approach, while Chapter 3 emphasizes normalization rules. The objective is to cover design early in the term so that students can get started on their end-of-term projects. Students should use the online Database Design system to work on exercises for both chapters to obtain feedback.

Queries are covered in two chapters. Chapter 4 introduces queries and focuses on the fundamentals of converting business questions to SQL queries. Chapter 5 discusses more complex queries, including subqueries and outer joins.

Part 3 describes the development of database applications, beginning with the essentials of building forms and reports in Chapter 6. Chapter 7 examines the common problems created in a multiuser environment. It explains the techniques used to handle data integrity and transactions. Chapter 8 shows how to put ev-

### Chapter 1: Introduction

#### Part 1: Systems Design

Chapter 2: Database Design

Appendix: Database Design System

Chapter 3: Data Normalization

Appendix: Formal Definitions of Normalization

#### Part 2: Queries

Chapter 4: Data Queries

Appendix: SQL Syntax

Chapter 5: Advanced Queries and Subqueries

Appendix: Introduction to Programming

#### Part 3: Applications

Chapter 6: Forms and Reports

Chapter 7: Database Integrity and Transactions

Chapter 8: Applications

Chapter 9: Data Warehouses and Data Mining

#### Part 4: Database Administration and New Systems

Chapter 10: Database Administration

Chapter 11: Distributed Databases

Chapter 12: Physical Data Storage

Chapter 13: Non-Relational Databases

everything together to build a complete application, including navigation and help files. Chapter 9 explains why analytical processing requires a different database configuration than transaction processing. It covers the main tools for analysis and data mining in a nonstatistical context.

Part 4 examines various topics in database administration and new tools. Chapter 10 examines management issues emphasizing planning, implementation, performance, and security. It explains the major tasks and controls needed by an administrator. Chapter 11 investigates the growing importance of providing distributed access to databases. It examines the impact of various network configurations. Chapter 12 leans toward computer science when it looks at how the DBMS physically stores data. The last two chapters can be difficult to cover in a single-term course, but are presented at an introductory level. Chapter 13 is new with the sixth edition and introduces contemporary non-relational databases typically used for high-scalability, massively parallel Web-based tasks. The design, queries, and tradeoffs are illustrated with the Cassandra DBMS.

Additionally, four chapters have appendixes that discuss programming concepts that are more technical. The appendix to Chapter 2 describes the online database design system that is available to instructors and students. It provides immediate feedback on database designs, making it easier for students to understand the problems and explore different designs. The appendix to Chapter 3 presents the formal definitions of normalization. They are provided for instructors and students who want to see the more formal set-theory definitions. The appendix to Chapter 4 is a convenient list of the primary SQL statements. The appendix to Chapter 5 provides an introduction to programming. It is designed as a summary or simple reminder notes.

## Pedagogy

---

The educational goal of the text is straightforward and emphasized in every chapter: By the end of the text, students should be able to build business applications using a DBMS. Throughout the text, many examples are used to apply and illustrate the concepts. The Web site also provides several databases so students can work with data, queries, forms, and applications. Students should be encouraged to apply the knowledge from each chapter by solving the exercises and working on their final projects.

Each chapter contains several sections to assist in understanding the material and in applying it to the design and creation of business applications:

- **What You Will Learn in This Chapter.** A list of questions that are answered within the chapter. Each question is echoed at the start of a section.
- **A Developer's View.** A student's perspective of the chapter contents.
- **Chapter Summary.** A brief review of the chapter topics.
- **A Developer's View.** A short summary of how the material in the chapter applies to building applications.
- **Getting Started.** A short statement of the main goal of the chapter focused on how the chapter contributes to designing and building applications.
- **Key Words.** A list of words introduced in the chapter. A full glossary is provided at the end of the text.

- **Additional Reading.** References for more detailed investigation of the topics.
- **Website References.** Some sites provide detailed information on the topic. Some are newsgroups where developers share questions and tips.
- **Review Questions.** Designed as a study guide for the exams, with a focus on the major topics within the chapter.
- **Exercises.** Problems that apply the concepts presented in the chapter. Most require the use of a DBMS.
- **Projects.** Several longer projects are available in a separate online document. They are suitable for an end-of-term project.
- **Workbooks.** Each workbook outlines the steps to build an application using a specific DBMS. Each workbook chapter illustrates tasks that match the discussion in the textbook. The workbook also provides exercises to build six other databases for different companies.
- **Sample Databases.** Three sample databases are provided to illustrate the concepts. Sally's Pet Store illustrates a database in the early design stages, whereas Rolling Thunder Bicycles presents a more finished application, complete with realistic data. Corner Med is a database for a neighborhood medical facility that tracks patient and physician interactions. Exercises for all three databases are provided in the chapters. The sample databases can be installed in several DBMS formats.

## Features of the Text

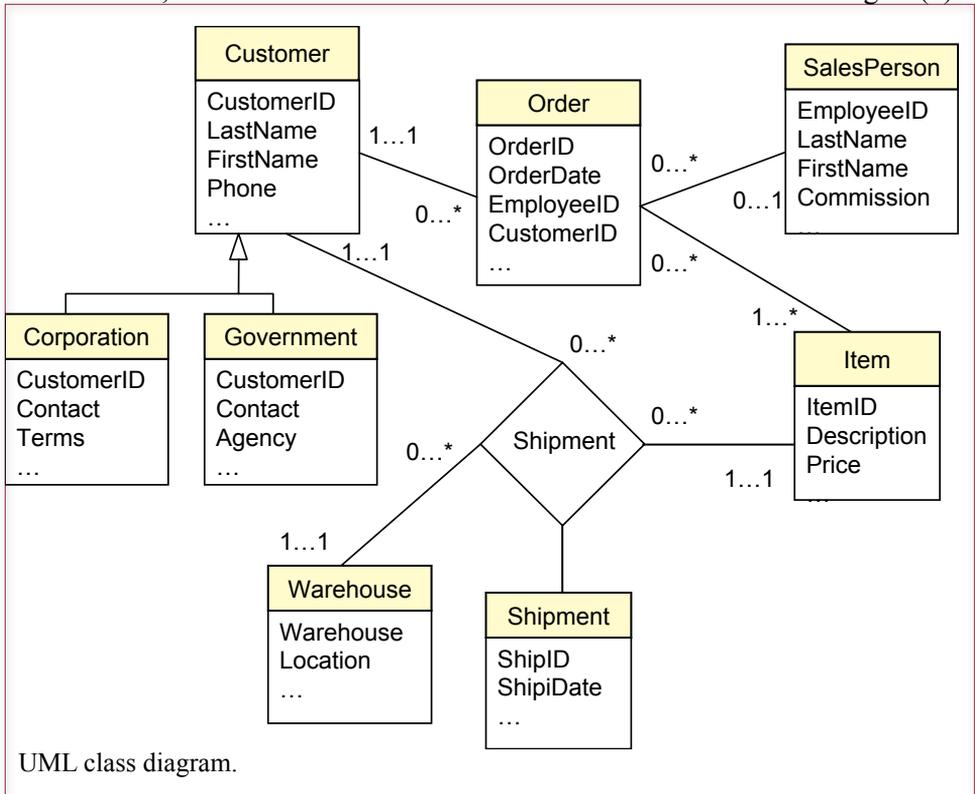
---

1. Focus on modern business application development.
  - Database design explained in terms of business modeling.
  - Application hands-on emphasis with many examples and exercises.
  - Emphasis on modern graphical user interface applications.
  - Chapters on database programming and application development.
  - Introduction to data mining.
  - Some answers to exercises are provided to students through the online system. The exercises with answers are highlighted in the text with a check mark icon. Sometimes students need to see a solved problem.
2. Hot topics.
  - Description and use of the unified modeling language (UML) for modeling and system diagrams.
  - In-depth discussion of security topics in a database environment.
  - Development of databases for the Internet and intranets.
  - Emphasis on SQL 92, with an introduction to SQL 99 and the XML features of SQL 2003 and SQL 2008.
  - Integrated applications and objects in databases.
  - Introduction to non-relational (NoSQL) systems such as Cassandra.

3. Applied business exercises and cases.
  - Many database design problems.
  - Exercises covering all aspects of application development.
  - Sample cases suitable for end-of-term projects.
4. A complete sample database application (Rolling Thunder Bicycles).
  - Fully functional business database.
  - Sample data and data generator routines.
  - Program code to illustrate common database operations.
5. Two additional databases (Sally's Pet Store and Corner Med) for comparison and additional assignments.
6. Lecture notes as PowerPoint slide show.
7. Hundreds of database exercises and problems for students to work on.
8. Workbooks built for specific database technologies that illustrate the hands-on steps needed to build an actual application. Check the online site for versions of the workbooks for additional systems.

## End-of-Term Projects

Several projects are described in the project document available online. These cases are suitable for end-of-term projects. Students should be able to build a complete application in one term. The grading focus should be on the final project. However, the instructor should evaluate at least two intermediate stages: (1) a



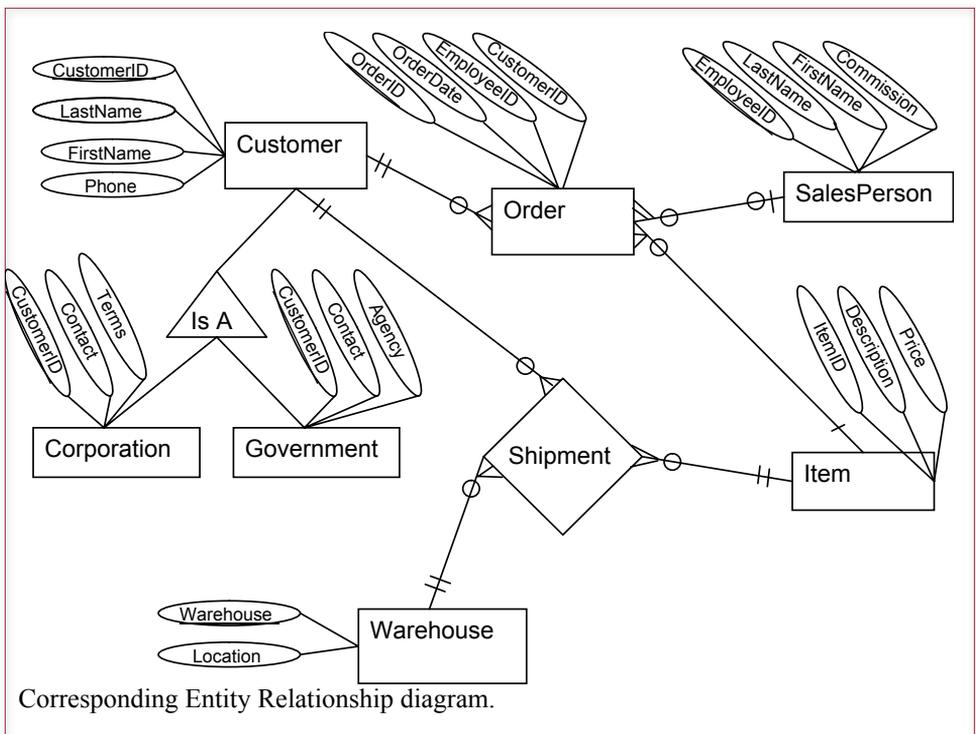
list of the normalized tables collected shortly after Chapter 3 is completed and (2) a design preview consisting of at least two major forms and two reports collected shortly after Chapter 6. The six additional cases in the workbook can also be used as an end-of-term project.

Some instructors may choose to assign the projects as group assignments. However, it is often wiser to avoid this approach and require individual work. The project is a key learning tool. If some members of the group avoid working on the project, they will lose an important learning opportunity.

## Database Design and the Unified Modeling Language

For several years, entity-relationship diagrams were the predominant modeling technique for database design. However, this approach causes problems for instructors (and students) because there are several different diagramming techniques. This edition continues to help solve these problems by incorporating the Unified Modeling Language (UML) method, instead of traditional entity-relationship (ER) diagramming, as the modeling technique for database design. This change will be most apparent in the replacement of the ER diagram notation and terminology with the parallel concepts in UML class diagrams.

UML class diagrams, although very similar to ER diagrams, are superior in several ways. First, they are standardized, so students (and instructors) need learn only one set of notations. Second, they are “cleaner” in the sense that they are easier to read without the bubbles and cryptic notations of traditional ER diagrams. Third, they provide an introduction to object-oriented design, so students will be better prepared for future development issues. Fourth, with the rapid adoption of UML as a standard design methodology, students will be better prepared to move



into future jobs. Many of the systems design organizations have adopted UML as a standard method for designing systems. UML has the support of major authors in systems design (e.g., Booch, Rumbaugh, and Jacobsen) as well as being supported by the major software development firms including IBM, Microsoft, and Oracle. Note that Microsoft Access and SQL Server both use a diagramming tool that is similar to UML. In addition, students should have little difficulty transferring their knowledge of the UML method if they need to work with older ER methods.

The basic similarities between ER and class diagrams are (1) entities (classes) are drawn as boxes, (2) binary relationships (associations) are drawn as connecting lines, and (3) N-ary associations (relationships) are drawn as diamonds. Hence, the overall structures are similar. The main differences between UML and ER diagrams occur in the details. In UML the multiplicity of an association is shown as simple numerical notation instead of as a cryptic icon. An example is shown in the accompanying figures.

UML also has provisions for n-ary associations and allows associations to be defined as classes. There are provisions for naming all associations, including directional names to assist in reading the diagram. Several situations have defined icons for the association ends, such as composition (rarely handled by ER) and subtypes (poorly handled by ER).

More details of the UML approach are shown in Chapters 2 and 3. Only a small fraction of the UML diagrams, notation, and terminology will be used in the database text. You can find the full specification on the Web at <http://www.omg.org/spec/UML/> with some introductory descriptions at <http://www.ibm.com/developerworks/rational/library/769.html>.

## Instructional Support

---

- **The online Database Design expert system.** It has been statistically proven to help students learn database design. It provides immediate feedback to students. It saves hours of instructor grading time. It contains over 100 design problems that can be used for teaching or testing.
- **A test bank** with multiple choice and short answer questions.
- **Lecture notes** and overheads are available as slide shows in Microsoft PowerPoint format. The slides contain all the figures and additional notes. The slides are organized into lectures and can be rearranged to suit individual preferences.
- **The sample databases** and solutions can be downloaded from the online site. The instructor can add new data, modify the exercises, or use them to expand the discussion in the text. The databases are provided for several DBMSs.
- **The Instructor's Manual** contains answers to the exercises.

## The Online System

---

All of the instructional material is available online. The main reason for this move is cost. The main textbook and the workbooks were rewritten and expanded. The costs and student prices for print books are out of line. Electronically, it is possible to make the entire set available for an almost trivial fee. It is impossible to

illustrate multiple DBMSs within a single textbook. Splitting the workbooks from the main text makes it possible to cover a variety of DBMSs—without confusing the reader. The DBMS market is becoming increasingly fragmented, and students need (1) a relatively agnostic main textbook to describe the common features, and (2) a workbook that provides the hands-on steps to actually build a database application. The online e-book method provides the additional benefit of showing the students how to accomplish the same tasks with multiple DBMSs. Even if students begin by learning one DBMS, they can download a second or third workbook and transfer their knowledge to a new DBMS.

E-books provide additional benefits, including advanced search capabilities. Students can also set bookmarks and highlight sections. More importantly, they get to keep the books, instead of being forced to sell them back at the end of the term. Database application development is an important topic, and the examples, comments, and tips in the books will be valuable to students throughout their careers.

## Major Changes with the Sixth Edition

---

The overall goals and structure remain the same with the sixth edition. Several sections were rewritten to improve clarity and incorporate some newer concepts. But the biggest change in each chapter was the addition and rewriting of most of the exercises.

Every chapter has a new Two-Minute Chapter section. This section summarizes the most important topics and overall goal of the chapter. It is useful for review to ensure students understand the key topics.

The new text also has a stronger bias towards Web-based applications. Almost any new application built today is either Web-based or possibly phone-based. Web (and mobile) applications heavily use centralized databases. Database design and queries remain similar, but interface, usability, and management issues are different from traditional in-house applications. As part of this emphasis, a new chapter (13) has been added to discuss features and limitations of the new crop of non-relational (sometimes called noSQL) database systems. These tools are designed for specific tasks: writing and retrieving specific data for millions of users. The open-source project Cassandra is used as an example to demonstrate the design and query tradeoffs in a highly-parallel environment.

# Brief Contents

---

1 Introduction

## **PART ONE: SYSTEMS DESIGN**

2 Database Design

3 Data Normalization

## **PART TWO: QUERIES**

4 Data Queries

5 Advanced Queries and Subqueries

## **PART THREE: APPLICATIONS**

6 Forms and Reports

7 Database Integrity and Transactions

8 Applications

9 Data Warehouses and Data Mining

## **PART FOUR: DATABASE ADMINISTRATION AND NEW SYSTEMS**

10 Database Administration

11 Distributed Databases

12 Physical Data Storage

13 Non-Relational Databases

- Goals and Philosophy, v
- Workbooks, v
- Learning Assessment, v
- Organization, vi
- Pedagogy, vii
- Features of the Text, viii
- End-of-Term Projects, ix
- Database Design and the Unified Modeling Language, x
- Instructional Support, xi
- The Online System, xi
- Major Changes with the Sixth Edition, xii
- Introduction, 1**
  - Introduction, 2
  - Two-Minute Chapter, 3
  - A Small Sample Database Application, 4
  - Databases and Application Development, 6
  - Components of a Database Management System, 10
    - Database Engine, 10*
    - Data Dictionary, 11*
    - Query Processor, 12*
    - Report Service, 13*
    - Forms Development, 14*
    - Management Utilities and Security, 15*
  - Advantages of the Database Management System Approach, 16
    - Focus on Data, 17*
    - Data Independence, 18*
    - Data Independence and Web Applications, 19*
  - Leading Commercial Database Systems, 20
  - The Evolution of Database Management Systems, 21
    - Hierarchical Databases, 21*
    - Network Databases, 23*
    - Relational Databases, 23*
    - Object-Oriented Databases, 24*
  - Key-Value Pairs: Cassandra, 28
  - Drawbacks to Database Management Systems, 29
    - Application Development, 30
    - Introduction to this Book's Databases, 31
      - Sally's Pet Store, 31*
      - Corner Med, 32*
      - Rolling Thunder Bicycles, 33*
    - Starting a Project: The Feasibility Study, 33
      - Costs, 33*
      - Benefits, 35*
- Summary, 36
- Key Terms, 37
- Review Questions, 37
- Exercises, 38
- Web Site References, 41
- Additional Reading, 41
- Systems Design, 42**
- Database Design, 43**
  - Introduction, 44
  - Two-Minute Chapter, 45
  - Models, 46
  - Getting Started, 47
  - Designing Databases, 48
    - Identifying User Requirements, 48*
    - Business Objects, 48*
    - Tables and Relationships, 50*
    - Definitions, 50*
    - Primary Key, 51*
  - Class Diagrams: Introduction, 51
    - Classes and Entities, 52*
    - Associations and Relationships, 53*
    - Class Diagram Details, 53*
  - Quick Start, 54
    - Creating a Class Diagram, 55*
    - Primary Keys and Relationships, 57*
  - Class Diagrams: Details, 59
    - Association Details: N-ary Associations, 60*
    - Association Details: Aggregation, 61*
    - Association Details: Composition, 62*
    - Association Details: Generalization, 63*
    - Association Details: Reflexive Association, 66*
  - Sally's Pet Store Class Diagram, 66
  - Data Types (Domains), 69
    - Text, 69*
    - Numbers, 69*
    - Dates and Times, 72*
    - Binary Objects, 72*
    - Computed Values, 73*
    - User-Defined Types (Domains/Objects), 73*
  - Events, 73
  - Large Projects, 75
  - Rolling Thunder Bicycles, 77
  - Application Design, 81
  - Corner Med, 82
  - Summary, 87
  - Key Terms, 88
  - Review Questions, 88

- Exercises, 89
- Web Site References, 98
- Additional Reading, 98
- Appendix: DBDesign, 99
- Getting Started: Identifying Columns, 100
- Creating a Table and Adding Columns, 101
- Relationships: Connecting Tables, 102
- Saving and Opening Solutions, 104
- Grading: Detecting and Solving Problems, 105
- Specifying Data Types, 107
- Generating Tables, 108
- Data Normalization, 111**
  - Introduction, 112
  - Two-Minute Chapter, 113
  - Tables, Classes, and Keys, 113
    - Composite Keys, 114*
    - Surrogate Keys, 115*
    - Notation, 116*
  - Database Normalization: Atomic Values and Dependency, 117
    - Atomic Data Values, 117*
    - Dependency, 119*
  - Sample Database for Typical Sales, 121
    - Initial Objects, 122*
    - Initial Form Evaluation, 123*
    - Problems with Repeating Sections, 125*
  - First Normal Form, 125
    - Repeating Groups, 125*
    - Multiple Repeating Groups, 127*
    - Nested Repeating Groups, 127*
  - Second Normal Form, 128
    - Problems with First Normal Form, 128*
    - Second Normal Form Definition, 129*
  - Third Normal Form, 132
    - Problems with Second Normal Form, 132*
    - Third Normal Form Definition, 132*
    - Checking Your Work, 135*
  - Beyond Third Normal Form, 135
    - Boyce-Codd Normal Form, 136*
    - Fourth Normal Form, 137*
    - Domain-Key Normal Form, 137*
    - Summary, 139*
  - Data Rules and Integrity, 139
  - The Effects of Business Rules, 141
  - Converting a Class Diagram to Normalized Tables, 143
    - One-to-Many Relationships, 144*
    - Many-to-Many Relationships, 146*
    - N-ary Associations, 147*
    - Generalization or Subtypes, 149*
  - Composition, 150*
  - Recursive (Reflexive) Associations, 151*
  - The Pet Store Example, 151
  - View Integration, 153
    - The Pet Store Example, 154*
    - Rolling Thunder Sample Integration Problem, 156*
  - Data Dictionary, 162
    - DBMS Table Definition, 163*
    - Data Volume and Usage, 166*
  - Summary, 168
  - Key Terms, 170
  - Review Questions, 170
  - Exercises, 171
  - Web Site References, 179
  - Additional Reading, 179
  - Appendix: Normal Form Definitions, 181
  - Queries, 185**
    - Data Queries, 186**
      - Introduction, 187
      - Two-Minute Chapter, 188
      - Three Tasks of a Query Language, 189
      - SQL SELECT Overview, 190
      - Four Questions to Retrieve Data, 191
        - What Output Do You Want to See?, 191*
        - What Do You Already Know?, 192*
        - What Tables Are Involved?, 192*
        - How Are the Tables Joined?, 193*
      - Sally's Pet Store, 195
      - Vendor Differences, 196
    - Query Basics, 196
      - Single Tables, 197*
      - Introduction to SQL, 198*
      - Sorting the Output, 200*
      - Distinct, 200*
      - Criteria, 201*
      - Pattern Matching, 202*
      - Boolean Algebra, 204*
      - DeMorgan's Law, 206*
      - Useful WHERE Clauses, 208*
    - Computations, 209
      - Basic Arithmetic Operators, 209*
      - Aggregation, 210*
      - Functions, 212*
    - Subtotals and GROUP BY, 214
      - Conditions on Totals (HAVING), 216*
      - WHERE versus HAVING, 216*
      - The Best and the Worst, 217*
    - Multiple Tables, 218
      - Joining Tables, 219*

- Identifying Columns in Different Tables*, 220
- Joining Many Tables*, 220
- Hints on Joining Tables*, 222
- Table Alias*, 223
- Create View*, 224
- Newer Searches and Patterns, 226
  - XQuery*, 227
  - Regular Expressions (RegEx) Patterns*, 233
- Summary, 239
- Key Terms, 240
- Review Questions, 240
- Exercises, 241
- Web Site References, 247
- Additional Reading, 247
- Advanced Queries and Subqueries, 251**
  - Introduction, 252
  - Two-Minute Chapter, 253
  - Sally's Pet Store, 254
  - Outer Joins (LEFT JOIN), 255
  - Subqueries: IN and NOT IN, 258
  - Subqueries, 261
    - Calculations or Simple Lookup*, 262
    - Calculations for Percentages*, 262
    - Subqueries and Sets of Data*, 264
    - Subquery with ANY, ALL, and EXISTS*, 266
  - Correlated Subqueries, 268
  - More Features and Tricks with SQL SELECT, 270
    - UNION, INTERSECT, EXCEPT*, 270
    - Multiple JOIN Columns*, 272
    - Reflexive Join*, 273
    - CASE Function*, 275
    - Inequality Joins*, 276
    - Exists and Crosstabs*, 277
    - SQL SELECT Summary*, 280
  - SQL Data Definition Commands, 280
  - SQL Data Manipulation Commands, 283
    - INSERT and DELETE*, 283
    - UPDATE*, 284
  - Quality: Testing Queries, 285
  - Summary, 287
  - Key Terms, 288
  - Review Questions, 289
  - Exercises, 290
  - Web Site References, 295
  - Additional Reading, 295
  - Variables and Data, 296
  - Variable Scope, 297
  - Computations, 298
  - Standard Internal Functions, 300
  - Input and Output, 300
    - Conditions, 301
    - Loops, 303
    - Subroutines, 304
    - Summary, 305
- Applications, 306**
- Forms and Reports, 307**
  - Introduction, 308
  - Two-Minute Chapter, 310
  - Effective Design of Reports and Forms, 310
    - Human Factors Design*, 311
    - Standard Form Controls*, 313
    - User Interface—Events*, 316
    - User Interface—Accessibility Issues*, 316
    - User Interface—International Environment*, 317
    - Style Sheets and Templates*, 320
  - Form Layout, 320
    - Tabular Forms*, 321
    - Single-Row or Columnar-Forms*, 322
    - Subform Forms*, 322
    - Startup Forms*, 324
  - Creating Forms, 325
    - Updateable Queries*, 326
    - Linked Forms*, 327
    - Properties and Controls*, 327
    - Controls on Forms*, 328
    - Multiple Forms*, 332
  - Direct Manipulation of Graphical Objects, 333
    - Sally's Pet Store Example*, 334
    - The Internet*, 335
    - Complications and Limitations of a Graphical Approach*, 335
  - Database Design Revisited, 336
  - Reports, 337
    - Report Design*, 338
    - Terminology*, 339
    - Basic Report Types*, 340
    - Charts*, 346
  - Summary, 346
  - Key Terms, 348
  - Review Questions, 348
  - Exercises, 349
  - Web Site References, 352
  - Additional Reading, 352
- Database Integrity and Transactions, 353**
  - Introduction, 354
  - Two-Minute Chapter, 355
  - Procedural Languages, 355

- Where Should Code Be Located?, 356
- User-Defined Functions, 357
- Looking Up Data, 358
- Programming Tools, 359
- Data Triggers, 360
  - Statement versus Row Triggers, 361
  - Canceling Data Changes in Triggers, 362
  - Cascading Triggers, 363
  - INSTEAD OF Triggers, 364
  - Trigger Summary, 365
- Transactions, 366
  - A Transaction Example, 366
  - Starting and Ending Transactions, 367
  - SAVEPOINT, 368
- Multiple Users and Concurrent Access, 369
  - Optimistic Locks, 370
  - Pessimistic Locks: Serialization, 373
  - Multuser Databases: Concurrent Access and Deadlock, 373
- ACID Transactions, 375
- Key Generation, 377
- Database Cursors, 378
  - Cursor Basics, 379
  - Scrollable Cursors, 380
  - Changing or Deleting Data with Cursors, 381
  - Cursors with Parameters, 383
- Merchandise Inventory at Sally's Pet Store, 384
- Summary, 388
- Key Terms, 389
- Review Questions, 389
- Exercises, 390
- Web Site References, 394
- Additional Reading, 394
- Application Development, 395**
  - Introduction, 396
  - Two-Minute Chapter, 397
  - Design Consistency, 398
    - Page Design Templates, 398
    - Usability, 399
    - Fonts and Customization, 400
    - Mobile Devices, 401
  - Application Structure, 402
    - Designing Applications, 403
    - The Startup Form, 403
    - Sally's Pet Store: Application Organization, 404
    - Administrative Tasks, 407
  - Menus and Toolbars, 407
    - Purpose of the Menu, 408
    - Toolbars, 409
    - Creating Menus and Toolbars, 409
  - Accessibility, 410
  - Custom Help, 412
    - Creating a Help File for Windows, 413
    - Context-Sensitive Help, 415
    - Windows Help 3/Help Viewer, 417
  - Handling Errors, 419
    - Catching Errors, 419
    - Logging Errors, 420
    - Debugging, 420
  - Testing, 420
    - Form and Module Testing, 421
    - Integrated Application Testing, 422
    - Stress or Performance Testing, 422
    - Usability Testing, 422
    - Security Testing, 423
  - Deploying an Application, 424
    - Packaging Files, 424
    - Installation Programs, 425
    - Server and Database Configuration, 425
  - Summary, 425
  - Key Terms, 426
  - Review Questions, 427
  - Exercises, 427
  - Web Site References, 429
  - Additional Reading, 429
- Data Warehouses and Data Mining, 430**
  - Introduction, 431
  - Two-Minute Chapter, 432
  - Indexes, 433
    - Binary Search, 434
    - Pointers and Indexes, 435
    - Creating Indexes, 436
    - Problems with Indexes, 437
  - Data Warehouses and Online Analytical Processing, 437
    - Data Warehouse Goals, 438
    - Data Warehouse Issues, 439
  - Data Extraction, Transformation, and Transportation, 441
  - OLAP Concepts, 443
  - OLAP Database Design, 445
    - Snowflake Design, 446
    - Star Design, 447
  - OLAP Data Analysis, 448
    - Cube Browsers, 448
    - OLAP in SQL, 450
    - SQL Analytic Functions, 455
    - SQL OLAP Windows Partition, 456
  - Data Mining and Business Intelligence, 458
    - Data Configuration, 459
    - Classification, 460

- Association Rules/Market Basket Analysis*, 463
- Cluster Analysis*, 467
- Geographic Analysis*, 469
- Summary, 472
- Key Terms, 473
- Review Questions, 473
- Exercises, 474
- Additional Reading, 477

**Database Administration, 478**

**Database Administration, 479**

- Introduction, 480
- Two-Minute Chapter, 481
- Data Administrator, 482
- Database Administrator, 483
- Database Structure, 485
- Metadata, 486
- Database Tasks by Development Stages, 488
  - Database Planning*, 488
  - Database Design*, 489
  - Database Implementation*, 489
  - Database Operation and Maintenance*, 490
- Backup and Recovery, 492
- Physical Configuration, 494
- Security and Privacy, 496
  - Data Privacy*, 497
  - Threats*, 498
  - Physical Security*, 499
  - Managerial Controls*, 500
  - Logical Security*, 500
  - Division of Duties*, 506
  - Software Updates*, 507
- Encryption, 507
- Sally's Pet Store, 510
- Summary, 512
- Key Terms, 514
- Review Questions, 514
- Exercises, 515
- Web Site References, 519
- Additional Reading, 519

**Distributed Databases, 520**

- Introduction, 521
- Two-Minute Chapter, 522
- Distributed Databases, 523
  - Goals and Rules*, 524
  - Advantages and Applications*, 525
  - Creating a Distributed Database System*, 526
  - Network Speeds*, 527

- Query Processing and Data Transfer*, 529
- Data Replication*, 530
- Generating Keys with Replicated Data*, 532
- Concurrency, Locks, and Transactions*, 533
- Distributed Transaction Managers*, 535
- Distributed Design Questions*, 536
- Client/Server Databases, 536
  - Client/Server versus File Server*, 537
  - Three-Tier Client/Server Model*, 539
  - The Back End: Server Databases*, 540
  - The Front End: Windows Clients*, 540
  - Maintaining Database Independence in the Client*, 541
- Centralizing with a Web Server, 542
  - Web Server Database Fundamentals*, 543
  - Browser and Server Perspectives*, 545
- Data Transmission Issues in Applications, 546
- Cloud Databases, 548
  - Cloud Computing Basics*, 548
  - Data Storage in the Cloud*, 549
  - Sally's Pet Store*, 550
- Summary, 551
- Key Terms, 553
- Review Questions, 553
- Exercises, 554
- Web Site References, 557
- Additional Reading, 557

**Physical Database Design, 558**

- Introduction, 559
- Two-Minute Chapter, 560
- Physical Data Storage, 560
- Table Operations, 561
  - Retrieve Data*, 562
  - Store Data*, 563
  - Reorganize the Database*, 564
  - Identifying Problems*, 565
- Data Storage Methods, 565
  - Sequential Storage*, 566
  - Pointers and Indexes*, 567
  - Linked Lists*, 569
  - B+ Trees*, 571
  - Direct or Hashed Access*, 574
  - Bitmap Index*, 575
  - Comparison of Access Methods*, 575
- Storing Data Columns, 576
  - Text and Numbers*, 576
  - Image and Binary Data*, 578
  - Transferring Data with Delimited Files*, 578
- Data Clustering and Partitioning, 579
  - Data Clustering*, 579
  - Data Partitioning*, 580

Managing Tablespaces, 582  
Sally's Pet Store, 582  
Summary, 583  
Key Terms, 584  
Review Questions, 584  
Exercises, 585  
Web Site References, 586  
Additional Reading, 587

## **Non-Relational Databases, 588**

Introduction, 589  
Two-Minute Chapter, 591  
Non-Relational Databases: Background, 592  
    *Key-Value Pairs, 594*  
    *Sparse Data and Flexible Columns, 595*  
    *Distributed Data, 597*  
    *Consistency and Integrity, 598*  
    *Optimizing Data Storage for Queries, 600*  
Cassandra, 601  
    *Installation Issues, 601*  
    *Pet Store Web Example, 602*  
    *Database Design, 603*  
    *Primary Keys, 605*  
    *Initial Queries, 607*  
    *Indexes, 609*  
    *Querying Tables with Compound Keys, 612*  
    *INSERT and UPDATE, 613*  
Cloud Databases, 614  
Summary, 616  
Key Terms, 618  
Review Questions, 618  
Exercises, 619  
Web Site References, 621  
Additional Reading, 621